

NBER WORKING PAPER SERIES

A STRUCTURAL APPROACH TO IDENTIFYING THE SOURCES OF LOCAL-CURRENCY
PRICE STABILITY

Pinelopi K. Goldberg
Rebecca Hellerstein

Working Paper 13183
<http://www.nber.org/papers/w13183>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2007

We are grateful to Charles Engel for very useful conversations at an early stage of this project. We have also benefitted from the comments of seminar participants at Harvard, NYU, Princeton, Yale, 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, the Federal Reserve System, or the National Bureau of Economic Research.

© 2007 by Pinelopi K. Goldberg and Rebecca Hellerstein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Structural Approach to Identifying the Sources of Local-Currency Price Stability
Pinelopi K. Goldberg and Rebecca Hellerstein
NBER Working Paper No. 13183
June 2007, Revised August 2011
JEL No. E30,F10,F30,L10

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 structural model to identify the sources of this local-currency price stability and applies it to micro data from the beer market. The empirical procedure exploits manufacturers' and retailers' first-order conditions in conjunction with detailed information on the frequency of price adjustments following exchange-rate changes to quantify the relative importance of local non-traded cost components, markup adjustment by manufacturers and retailers, and nominal price rigidities in the incomplete transmission of such changes to prices. We find that, on average, approximately 60% of the incomplete exchange rate pass-through is due to local non-traded costs; 8% to markup adjustment; 30% to the existence of own-brand price adjustment costs, and 1% to the indirect/strategic effect of such costs, though these results vary considerably across individual brands according to their market shares.

Pinelopi K. Goldberg
Yale University
Department of Economics
37 Hillhouse Ave.
P.O. Box 208264
New Haven, CT 06520-8264
and NBER
penny.goldberg@yale.edu

Rebecca Hellerstein
International Research
Federal Reserve Bank of New York
33 Liberty Street
New York, NY 10045
Rebecca.Hellerstein@ny.frb.org

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 global imbalances, it is more important than ever to understand why import prices do not respond fully to exchange-rate changes, especially since different explanations have very different implications for exchange rate policy.

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 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 identified interesting empirical patterns that are surprisingly robust across markets, time, and specific modeling assumptions.

The general structure of the approach in this 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 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 not 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. 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, 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 this 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 for 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 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 as well as aggregated across product categories, one observes prices changing every period. This makes it impossible to identify potential costs of repricing, which by nature imply that prices should sometimes remain fixed. Hence, to the extent that such price rigidities are present, they 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 because 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 approach) it could be due to the existence of 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 is to modify the standard framework of profit maximization to allow firms to deviate from their first-order conditions when faced with fixed costs of repricing. 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 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. To identify the potential role of nominal price rigidities we propose using higher frequency (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 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 followed by earlier papers by constraining the estimation to the periods in 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 the 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 a 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 periods when prices adjust and periods when prices remain unchanged to derive bounds on the adjustment costs associated with a price change. This derivation draws on a “revealed-preference-approach”: Our method is based on the insight that in periods in which prices change, the costs of price adjustment must be lower than the additional profit the firm makes by changing its price; we use this insight to derive an upper bound to this price adjustment cost. Similarly, in periods when prices do not change, the costs of adjustment must exceed the extra profit associated with a price change; we use this insight to derive a lower bound for the price adjustment cost.

The costs of price adjustment is 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. As such, they are not directly comparable to estimates obtained in earlier studies using different methods (e.g., direct measurement or firm surveys).¹ More importantly,

¹Several studies attempt to measure price adjustment costs directly. 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

the adjustments costs alone do not allow for 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 firms' pricing behavior with price rigidities to their behavior given fully flexible prices. The differential response of prices across these two scenarios is attributed to the effect of nominal price rigidities. In the same procedure, we also identify the role of markup adjustment and non-traded costs in generating incomplete pass-through.

We apply the approach described above to weekly store-level data for the beer market. The beer market is well suited for investigating questions regarding 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) long-run exchange-rate pass-through to consumer prices is low, on the order of 5-10 percent; (3) highly disaggregate weekly data are available with both wholesale and retail prices, which enable us to examine how prices respond at each stage of the distribution chain; (4) The patterns of prices for beer are typical for the type of goods included in the CPI, as summarized by Klenow and Malin (2010); in particular, "regular" prices remain constant over several weeks, but there are periodic discounts from these regular prices (i.e., sales); (5) both non-traded local costs and price rigidities are *a-priori* plausible; as noted above, weekly data reveal price inertia, both at the wholesale and retail level, which is suggestive of price rigidities. While our particular assumptions regarding demand and supply are tailored to the beer market, the general features of our approach can be applied to any market for which high-frequency data are available to identify the points of price adjustment.

Perhaps the biggest caveat of the approach we propose is its static nature. Dynamic considerations may affect the analysis in two 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. On the demand side, Hendel and Nevo (2006a, 2006b) show that when consumers stockpile laundry detergents in response to temporary price reductions (sales), static demand estimates may overstate the long-run price elasticities of demand by a factor of 2 to 6. On the supply side, Aguirregabiria (1999) analyzes the pricing behavior of a retailer who holds inventories in a central store and delivers these goods from to its 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;² ignoring such dynamics may lead to biased estimates of the importance of

drugstores' revenue. A few detailed microeconomic studies have cast doubt on the importance of menu costs in price rigidity: Carlton (1986) and Midrigan (2010) find that firms change prices frequently and in small increments, which is not consistent with a menu-cost explanation of 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.

²A similar point is made in Alessandria, Kaboski, and Midrigan (2010).

nominal price rigidities. Fortunately, concerns that rely 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 consumer stockpiling is not a first-order concern in this market.³ On the supply side, state and local regulations concerning the distribution of all alcohol, including beer, in the market we study 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. 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 or the individual outlet, there is no inventory problem associated with beer, unlike most other products which *are* distributed by the central store.⁴ As the central store does not keep inventories of beer (indeed cannot by law), there is no relationship between inventory decisions and prices. And there is no incentive for individual outlets to maintain inventories, as they can get shipments each week from the distributor rather than bearing the costs of holding inventories themselves.

A second limitation of the static approach is that it fails to model explicitly the fact that with ongoing uncertainty and rational expectations there is an option value to not adjusting prices, which will magnify the effects of even small costs of adjustment - a point initially made by Dixit (1991). Failure to model this option value may bias adjustment cost estimates upwards. Our approach is similar to the models considered in Akerlof and Yellen (1985) and Mankiw (1985) in which agents have myopic expectations. While this is certainly a strong (and in many settings unreasonable) assumption, we note that the primary source of uncertainty in our setting is *exchange rates*, which are highly persistent. Indeed, the consensus among International economists is that exchange rates follows a random walk. Therefore, the assumption of myopic expectations seems more reasonable in the context of exchange rates than in the case of other cost shocks, which may exhibit less persistence. In addition, we run alternative versions of our model that show that our results are robust to relaxing this assumption of firm myopia. Unfortunately, characterizing firms' optimal behavior in a fully dynamic setup requires working with quadratic approximations to profit functions as in Dixit (1991) and Caplin and Leahy (1997), and abstracting from firm heterogeneity and product differentiation. In contrast to these papers, in which dynamics are key, our approach places the emphasis of the analysis on product differentiation, identification of the demand curvature, and firms' strategic interactions at the expense of dynamics. We reiterate, however, that within the static framework we interpret the derived "adjustment costs" in the broadest possible sense which includes the option value of the status

³See Anheuser-Busch, Inc. "Beer Shopper Poll," 2005.

⁴As Aguirregabiria (1999) notes, "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).

quo rather than the literal labor or material costs a firm has to pay to change prices. We hope that future research can make more progress in merging the current framework with an explicit modeling of dynamics. We discuss these issues in more detail in Section 3.4. In the same section, we explore the implications of using *weekly* prices, as opposed to adopting longer time horizons.

Our analysis yields several interesting findings. First, at the descriptive level, we document infrequent price adjustment both at the wholesale and retail level. However, this price inertia seems driven primarily by infrequent adjustment of wholesale rather than retail prices. In our data, there is no instance where a product’s retail price remains unchanged in response to a wholesale price change. Hence, it seems the primary reason retail prices do not change each period is that there is little reason for them to do so, as the underlying wholesale prices remain fixed.

As we discussed above, nominal price rigidities may affect the pricing decisions of a particular producer in two ways. First, they may prevent the 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 over 30 percent on average of 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 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. There is also an *indirect/strategic* effect at the wholesale stage of the distribution chain that accounts for approximately 1 percent of the overall incomplete pass-through, though over 30 percent for the foreign brand with the smallest market share among those in the counterfactuals, *Bass*. Our final decomposition attributes 60.0 percent of the incomplete pass-through to local non-traded costs, 8.3 percent to markup adjustment, 30.4 percent to the existence of own-brand price adjustment costs, and 1.3 percent to 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.

It is important to distinguish between the general features of our approach that can be adopted easily by other researchers and those specific to our application. These general features include: the use of product-level price and market-share data along with limited demographic information to estimate demand and markups; the use of a “revealed-preference approach” that combined with firms’ first-order conditions of profit maximization allows us to compute bounds on their price adjustment costs; and the use of an accounting procedure to assess the relative importance of the sources of incomplete exchange-rate pass-through. Similarly, as we noted above, we believe that the assumption of myopic expectations is – independent of the particular industry under consideration – defensible in the case of exchange-rate shocks, so that it can be adopted in other studies of *exchange-rate* pass-through.

But the specific assumptions of static demand and supply and of firms that are vertically-separated, Bertrand-Nash competitors, cannot be applied *as is* to any given industry. A researcher will always need to make specific assumptions to examine any given industry, e.g., some industries may require dynamics to be incorporated on the demand or supply side, while others may need a different model of vertical interaction. We do not claim to provide a model that can be applied “off the shelf” to *any* industry, but rather a general approach that can be integrated with the institutional details of specific markets to examine the sources of incomplete pass-through.

Along the same lines, it is instructive to point out which features of the data drive our results and how they generalize to other settings. In our approach demand parameters, and so markups, are identified from plausibly exogenous variation in relative prices across products over time (i.e., variation from changes in input prices and bilateral exchange rates). Total costs are in turn identified as the difference between prices and markups. Because markups are estimated to be relatively stable over time, costs, expressed in local currency (U.S. dollars), also appear remarkably stable over time. In the absence of repricing costs, this stability can only be rationalized by the existence of a local, non-traded cost component that is not affected by exchange rates. Put differently, the markup adjustment generated by a demand system that is estimated from the relative price changes observed in our data is not sufficient to generate the low pass-through observed in the data: Hence, in the absence of nominal price rigidities, local non-traded costs will emerge as the dominant residual explanation. This feature of the results is not unique to the beer market but has been consistently documented in many micro studies that employ flexible demand systems (e.g., Goldberg and Verboven, 2001; Hellerstein, 2008; Nakamura and Zerom, 2010). The contribution of our study is to allow for an additional – to local, non-traded costs – source of price (and hence derived cost) stability: repricing costs. Once we have estimated markups, we identify the role of price rigidities separately from that of non-traded costs, by distinguishing between periods of price adjustment and periods of non-adjustment. We find that even conditional on price adjustment, derived costs appear to be stable. This provides strong evidence in favor of local, non-traded costs. On the other hand, these non-traded costs cannot completely account for the fact that prices do not adjust at all in some periods: The latter can only be explained by the existence of nominal rigidities. Perhaps the most interesting aspect of our results is that we still find local non-traded costs to be the primary source of incomplete exchange rate pass-through, accounting on average for circa 60 percent of the incomplete response, despite the fact that we focus on a market with infrequent price adjustment and allow for nominal rigidities. We therefore believe that the large role attributed to local non-traded costs in studies of exchange-rate pass-through is a robust finding likely to apply to multiple consumer markets.

The remainder of the paper is organized as follows. To set the stage, we start by describing the market and the data and by examining the adjustment patterns in our retail and wholesale prices in the next section. Section 3 discusses the model, how it allows us to derive bounds for the price adjustment

costs, and the implications of its key assumptions, notably the static nature of our approach. Section 4 discusses the empirical implementation of the model and the estimation and simulation results, and Section 5 concludes.

2 Data

2.1 Overview

The beer market is particularly 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 (below ten percent); both non-traded local costs and price stickiness due to adjustment 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 the role of local non-traded costs and of adjustment costs in firms' incomplete pass-through of exchange-rate shocks to prices.

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 and wholesale prices for every major brand of beer sold in the U.S.⁵ 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.⁶ 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.⁷ Supermarkets

⁵The data can be found at <http://gsbwww.uchicago.edu/kilts/research/db/dominicks/>.

⁶Features of the *Dominicks'* wholesale-price data confirm that brewers control distributors prices to the supermarket. Across *Dominicks'* stores, which may each be served by a different distributor 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 industry lore that distributors pricing is coordinated by brewers and is not set separately by each distributor to each retail outlet.

⁷Canadian Trade Commissioner (1998).

sell approximately 20 percent of all beer consumed in the U.S.

In our data, we define a product as one six-pack serving of a brand of beer, quantity as the total number of servings sold per week, and a market as one of *Dominick's* price zones in one week. We aggregate data from each *Dominick's* store into one of two price zones (for more details about this procedure, see Hellerstein, 2008). 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, 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 and the characteristics data used in the demand estimation are provided in Table 1.

2.2 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 model. We estimate three price equations:

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

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

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

where the subscripts j , z , and t refer to product, zone, and week respectively; p^r is the product's retail price⁸; 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, including 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

⁸Imported beer is generally invoiced and paid for in U.S. dollars in this market. We do not believe this practice "explains" the observed incomplete pass-through: It simply restates the problem. A firm's decision to price in dollars to keep its local-currency prices constant still begs the question of why it wants its price constant, which our analysis addresses.

exporting country in our sample; and ε is a random error term. All the variables are specified in levels, not first differences, as our focus is on the long-run pass-through of exchange rate changes, and not the short-term dynamics.⁹

Table 2 reports results from OLS estimation of the pricing equations. Columns 2 and 4 report results from specifications that include the full set of controls specified above, while in columns 1 and 3 the cost controls are omitted (as 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 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 2 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 2 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? The 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.

2.3 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 what we expect with price stickiness. That said, by itself, the infrequent adjustment of prices is not 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

⁹It is not necessary, therefore, for import payments to be made on a weekly basis for us to identify the long-run impact of exchange-rate shocks on prices using this specification.

(sales). As we discuss in Section 3.4, these sales appear to be random in our sample, in the sense that we cannot find *anything* that predicts the timing of a sale. This has important implications for the demand estimation. Third, a striking feature of Figure 1 is that retail prices always adjust when wholesale prices adjust. Therefore, 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 contrasts 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 illustrate is that price adjustment is not synchronized across brands. Given the strategic interactions between firms, this asynchronous price adjustment may generate significant price inertia, even if the nominal price rigidities facing each individual manufacturer or retailer are small.

3 Model

We describe here 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 markup is added to the marginal cost of the product before it reaches the consumer.¹⁰ Our approach builds on Hellerstein’s (2008) work on the beer market, but makes one key modification to her model: 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.

The strategic interaction between the manufacturer and retailer is as follows. First, the manufacturer decides whether or not to change its 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 whether or not to change the same 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

¹⁰In modeling the industry’s vertical relationships we draw on previous work (Hellerstein, 2004) that used both industry lore and formal non-nested tests developed by Villas-Boas (2007) to choose the model that best fits the data.

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 in market t are given by:

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

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 non-traded¹¹ 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:

$$A_{jt}^r = 0 \text{ if } p_{jt}^r = p_{jt-1}^r; \quad A_{jt}^r > 0 \text{ if } p_{jt}^r \neq p_{jt-1}^r. \quad (5)$$

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 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 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: $p_{jt}^r \neq p_{jt-1}^r$. In this case the retailer solves the standard profit maximization problem to determine the new optimal price, so the observed

¹¹We use the term “non-traded” to indicate that these costs are paid in dollars regardless of the origin of the product, and so will not be affected by exchange-rate shocks.

retail price p_{jt}^r satisfies the first-order profit-maximizing condition:

$$s_{jt} + \sum_{k=1,2,\dots,J} (p_{kt}^r - p_{kt}^w - ntc_{kt}^r) \frac{\partial s_{kt}}{\partial p_{jt}^r} = 0. \quad (6)$$

This gives us a set of J equations, one for each product. One can solve for the retailer's markups by defining a $J \times J$ matrix Ω_{rt} , called the retailer reaction matrix (the matrix of demand substitution patterns), with element $S_{jk}^r = \frac{\partial s_{kt}(p_t^r)}{\partial p_{jt}^r}$ $j, k = 1, \dots, J$, that is the marginal change in the k th product's market share given a change in the j th product's retail price.¹² The stacked first-order conditions can be rewritten in vector notation and inverted together in each market to get the retailer's pricing equation:

$$p_t^r = p_t^w + ntc_t^r - \Omega_{rt}^{-1} s_t \quad (7)$$

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 if the retailer changes her price in the current period, the extra profits associated with the new price must be at least as large as the adjustment cost:

$$(p_{jt}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_t^r) + \sum_k (p_{kt}^r - p_{kt}^w - ntc_{kt}^r) s_{kt}(p_t^r) - A_{jt}^r \geq \\ (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{-jt}^r) + \sum_k (p_{kt}^r - p_{kt}^w - ntc_{kt}^r) s_{kt}^c(p_{jt-1}^r, p_{-jt}^r), j \neq k \quad (8)$$

where $s_{jt}^c(p_{jt-1}^r, p_{-jt}^r)$ denotes the *counterfactual* market share that product j *would* have, if the retailer had kept the price unchanged at p_{jt-1}^r , and p_{-jt}^r denotes the other products' prices that may or may not have changed from the previous period. This inequality states that the profits the retailer makes by adjusting product j 's price in the current period must be greater than the profits she *would have* achieved had she left the price unchanged (in which case the first-order condition of profit maximization would have been violated, but she would have saved on the adjustment cost A_{jt}^r). In this sense, it simply captures a "revealed-preference" argument. Note that this inequality is only one from a large set of inequalities that could potentially be used to infer adjustment costs, as a multi-product retailer may consider all possible permutations of price changes across the multiple products she offers. However, we do not need to exploit *all* the inequalities to infer the bounds. In fact, this would be infeasible. Using a subset, in this case containing a single inequality, is sufficient to obtain consistent estimates of the bounds given that our revealed preference approach must hold for every single permutation of price changes. (Note that the same applies to the derivation of the retailer's lower bounds as well as

¹²The Ω_{rt} matrix is pre-multiplied by $J \times J$ matrix T^r with the j th, k th element equal to 1 if both products j and k are sold by the retailer, and zero otherwise. In our application, T^r is all ones so we omit it for simplicity.

the derivation of the manufacturer's bounds.) By rearranging terms we can use the above inequality to derive an *upper* bound $\overline{A_{jt}^r}$ to the price adjustment costs of product j :

$$A_{jt}^r \leq \overline{A_{jt}^r} = (p_{jt}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_t^r) - (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{-jt}^r) + \sum_k [(p_{kt}^r - p_{kt}^w - ntc_{kt}^r) (s_{kt}(p_t^r) - s_{kt}^c(p_{jt-1}^r, p_{-jt}^r))], j \neq k \quad (9)$$

Case 2: The price remains unchanged from the previous period: $p_{jt}^r = p_{jt-1}^r$. In this case the first-order conditions of profit maximization do not necessarily hold. If the retailer does not adjust product j 's price in period t , then the profits she makes from keeping the price constant must be as large as the profits she would have made from adjusting the price according to her first-order condition less the adjustment cost:

$$(p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_{jt-1}^r, p_{-jt}^r) + \sum_k (p_{kt}^r - p_{kt}^w - ntc_{kt}^r) s_{kt}(p_t^r) \geq (p_{jt}^{rc} - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^r) + \sum_k (p_{kt}^r - p_{kt}^w - ntc_{kt}^r) s_{kt}^c(p_{jt}^{rc}, p_{-jt}^r) - A_{jt}^r, j \neq k \quad (10)$$

where p_{jt}^{rc} denotes the *counterfactual* price the retailer *would have* charged had she behaved according to her optimality conditions, and $s_{jt}^c(p_{jt}^{rc}, p_{-jt}^r)$ the *counterfactual* market share corresponding to this optimal price. As in Case 1, we rewrite this inequality to derive a *lower* bound $\underline{A_{jt}^r}$ to the adjustment costs:

$$A_{jt}^r \geq \underline{A_{jt}^r} = (p_{jt}^{rc} - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^r) - (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_{jt-1}^r, p_{-jt}^r) + \sum_k [(p_{kt}^r - p_{kt}^w - ntc_{kt}^r) (s_{kt}^c(p_{jt}^{rc}, p_{-jt}^r) - s_{kt}(p_t^r))], j \neq k \quad (11)$$

The heart 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 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 from the first-order conditions; however, we can derive estimates of them by imposing some additional structure on the problem, i.e. by modeling non-traded costs parametrically as a function of observables. With these in hand, we calculate the counterfactual price p_{jt}^{rc} 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 $s_{jt}^c(p_{jt}^{rc}, p_{-jt}^r)$. In the final step, we exploit inequalities (9) and (11) to derive upper and lower bounds to the adjustment costs A_{jt}^r .

Note that in this approach 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 a product's price if the adjustment cost associated with this price change exceeds the additional profit. Second, there is an indirect effect operating through the impact of price rigidities on competing products' prices. 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 size of the price change of product j may be smaller than if price rigidities were altogether non-existent. The existence of this indirect effect implies that relatively small adjustment costs can 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 that which would prevail with fully flexible prices.

3.1.2 Manufacturers

Let there be M manufacturers that each produce some subset S_w 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:

$$\Pi_{jt}^w = \sum_{j \in S_w} (p_{jt}^w - c_{jt}^w(tc_{jt}^w, ntc_{jt}^w)) s_{jt}(p_t^r(p_t^w)) - A_{jt}^w \quad (12)$$

where c_{jt}^w is the marginal cost incurred by the manufacturer to produce and sell product j , which in turn is 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 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 that for the retail adjustment cost; it is a discrete cost paid only when the manufacturer adjusts the price of product j :

$$A_{jt}^w = 0 \text{ if } p_{jt}^w = p_{jt-1}^w; \quad A_{jt}^w > 0 \text{ if } p_{jt}^w \neq p_{jt-1}^w \quad (13)$$

We use the same procedure we applied to the retailer's problem to derive upper and lower bounds to the manufacturer adjustment cost. The derivation of the manufacturer bounds is, however, more

complicated as he must account for the possibility that the retailer does not adjust her price due to the presence of a retail adjustment cost. As with the retailer, in the data we observe one of two cases:

Case 1: The wholesale price changes from the previous period: $p_{jt}^w \neq p_{jt-1}^w$. 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. However, in our data we do not observe a single instance of this occurring. 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, each wholesale price p_{jt}^w must satisfy the first-order profit-maximizing conditions if it was adjusted from the previous period:

$$s_{jt} + \sum_{k \in S_w} (p_{kt}^w - c_{kt}^w) \frac{\partial s_{kt}}{\partial p_{kt}^w} = 0. \quad (14)$$

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 market share with respect to a change in each product's wholesale price.¹³ 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 matrix are described in greater detail in Villas-Boas (2007) for the case with no adjustment costs.¹⁴ For the case *with* adjustment costs, note that we are able to compute this matrix even though we, as researchers, do not observe the retailer adjustment costs. The reason is that for Ω_{pt} to be evaluated, all that is required is for the manufacturer to know if the retailer will adjust the prices of her products. In our complete information setup, the manufacturer knows the adjustment costs of the retailer (even though we do not) and, thus, can solve the model to compute the optimal response of the retailer. At the equilibrium, the observed response of the retailer in the data has to correspond to the response assumed by the manufacturer. Otherwise, the observed responses would not be consistent with an equilibrium. Accordingly, we compute Ω_{pt} by setting the manufacturer's expectation regarding the retailer's price adjustment equal to the adjustment of the retailer observed in the data. Note that we do not need to know the retailer adjustment costs to do this.¹⁵ The manufacturers' marginal costs

¹³ Ω_{wt} is preceded by an ownership matrix, $J \times J$ matrix T^w with the j th, k th element equal to 1 if both products j and k are produced by the same manufacturer, and zero otherwise, which we omit for simplicity. Two manufacturers are modeled as multiproduct firms, *Heineken N.V.* which owned *Amstel* and *Heineken* over the sample period, and *Guinness PLC*, which owned *Guinness* and *Harp*.

¹⁴To 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 .

¹⁵For example, suppose that the retail price of a particular product does not adjust in our data. Then in our computation of Ω_{pt} we take into account that equilibrium implies that the manufacturer had correctly anticipated that the retail price of that product would not adjust. Suppose that this were not true, and that the manufacturer had instead assumed that the retail price of that product would change. Then, the manufacturer would have computed the corresponding Ω_{pt} and would have set the optimal wholesale price given the assumed retailer response. But given that the manufacturer knows the retailer adjustment costs, the manufacturer could then calculate whether the additional profits the retailer

(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:

$$p_t^w = c_t^w - \Omega_{wt}^{-1} s_t \quad (15)$$

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

If the manufacturer changed her price from the previous period, the profits she made from changing the price (net of the price adjustment cost A_{jt}^w) must have exceeded the profits that she *would have* made had she left the price unchanged at p_{jt-1}^w :

$$(p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) + \sum_{k \in S_w} (p_{kt}^w - c_{kt}^w) s_{kt}(p_t^r(p_t^w)) - A_{jt}^w \geq \\ (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) + \sum_{k \in S_w} (p_{kt}^w - c_{kt}^w) s_{kt}^c(p_{jt}^{rc}, p_{-jt}^{rc}), j \neq k \quad (16)$$

This condition is similar to inequality (9) for the retailer with one 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 , as well as the counterfactual retail prices p_{-jt}^{rc} of the other products offered by the retailer. The same applies for the counterfactual market shares s_{kt}^c of the other products sold by this manufacturer, $k \in S_w$. This poses the following problem: Given the existence of retailer adjustment costs, the counterfactual retail prices may or may not change from the previous period. To assess the manufacturer's beliefs about the retailer's behavior, we need to know the retailer adjustment costs; however, the procedure described in the previous subsection allows us to derive only bounds and not point estimates of these costs. Note that because the scenarios involved in the computation of p_{jt}^{rc} and p_{-jt}^{rc} are *counterfactual*, and hence by definition not observed in practice, we can no longer invoke the equilibrium conditions as we did above when computing the matrix Ω_{pt} to infer the retailer adjustment behavior that the manufacturer uses in her optimization from the actual retail prices changes observed in the data.

We deal with this issue in the following way. Consider p_{jt}^{rc} first. As noted earlier, in our data we never observe that the retail price of a product adjusts when its wholesale price does not. We interpret this pattern as indirect evidence that the manufacturer assumes that whenever she does not adjust the

would make by adjusting her price were high enough to justify the retail price change. If these adjustment costs were too high, then the retailer would not adjust her price, and the manufacturer would know that the response she had assumed was wrong. In short, for the observed price adjustments to be consistent with an equilibrium, it has to be the case that the retailer response anticipated by the manufacturer is the actual one observed in the data.

wholesale price, the retail price will not change either; accordingly, we set $p_{jt}^{rc} = p_{jt-1}^r$ if $p_{jt}^w = p_{jt-1}^w$. We emphasize that the justification of this approach is *not* based on an equilibrium notion; rather, it reflects a conjecture that is justified purely on *empirical* grounds, namely based on the observation that in the data retail price changes occur only when wholesale prices change.¹⁶ Given that this approach is based on a conjecture, we examine the robustness of our results to alternative ways of computing counterfactual prices as follows. Note that in the case under consideration, the retail price for product j has adjusted. Accordingly, we have an upper bound for the retailer adjustment cost for that period from solving the retailer’s problem in the previous subsection. We can therefore solve the retailer’s problem again under the premise that the retailer adjusts her price, even though the wholesale price did not change, use the retailer FOC to determine the optimal retail price in this case, and then insert our estimate of the upper bound to the retailer adjustment cost from inequality (9) to determine whether it is indeed optimal for the retailer to adjust her price. The drawback of this approach is that the use of the *upper* bound of the retailer adjustment cost biases the results towards finding that the retailer will not adjust (which is what we want). Therefore, we also conduct an alternative robustness check, in which we use an average of the lower bounds (obtained for different periods) to conduct the same calculations. In both cases, we find that the retailer’s optimal response is *not* to adjust her price if the wholesale price did not change from the previous period. Intuitively, this happens because when comparing the retailer profits for a product across two periods, t and $t - 1$, there are hardly any changes in the variables affecting profits, unless the wholesale price of the product has changed. If the wholesale price is the same, the change in the profits is very small, and it is optimal for the retailer not to adjust her price even when the adjustment cost is low.

To compute the counterfactual prices of the other products, p_{-jt}^{rc} , we again rely on the observation that *empirically* the retail price of a product almost never changes when only the wholesale prices of other products change; it only changes if the wholesale price of the same product has changed from the previous period. Accordingly, if the retail price of another product did not change from the previous period, we do not change the price of that product in the counterfactual either. However, if the price of that product has changed from the previous period, we compute a new counterfactual price based on the retailer’s optimization conditions.¹⁷ Once these counterfactual prices and shares have been

¹⁶In other words, our premise here is that the manufacturer, who observes all adjustment costs, solves the model and computes the equilibrium associated with the counterfactual scenario. We, as researchers, cannot compute this equilibrium, but based on the empirical observation that retailer prices adjust only when wholesale prices adjust, we conjecture that the equilibrium computed by the manufacturer also displays this feature evident in the data; that is, it implies that if the wholesale price did not adjust, the retail price will not adjust either.

¹⁷For example, suppose we have three products, j , m , and n . Suppose also that in reality at time t , their retail prices are: $p_{jt}^r, p_{mt}^r, p_{nt}^r$. Then, in the counterfactual, we use: $p_{jt-1}^r, p_{mt-1}^r, p_{nt-1}^r$. The idea here is that the prices of products m and n would not change if the wholesale (and retail) price of product j had been different, since in the data we only see the retail price of a product changing when its own wholesale price changes. But if the prices in period t are $p_{jt}^r, p_{mt}^r, p_{nt-1}^r$, then in the counterfactual we use $p_{jt-1}^r, p_{mt}^{rc}, p_{nt-1}^r$. So, in this case we compute the new optimal p_{mt}^{rc} , based on the retailer response matrix (the idea here being that the retail price of product m had changed from the previous period, almost certainly because its wholesale price had changed too, so that the retailer had paid the associated

computed, the upper bound to the *manufacturer's* adjustment cost $\overline{A_{jt}^w}$ can be derived based on the inequality:

$$A_{jt}^w \leq \overline{A_{jt}^w} = (p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) - (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) + \sum_{k \in S_w} [(p_{kt}^w - c_{kt}^w) (s_{kt}(p_t^r(p_t^w)) - s_{kt}^c(p_{jt}^{rc}, p_{-jt}^{rc}))], j \neq k \quad (17)$$

Case 2: The wholesale price does not change from the previous period: $p_{jt}^w = p_{jt-1}^w$. The lack of price adjustment in this case implies that the wholesale price is not necessarily determined by the manufacturer first-order condition. Regarding the retail price, it is again possible that the retailer adjusts its price in periods when the wholesale price remains unchanged. However, in practice we do not 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, the profits she makes at p_{jt-1}^w must be at least as large as the profit she *would have* made had she changed the price to a counterfactual wholesale price p_{jt}^{wc} according to her profit maximization condition and paid the associated adjustment cost:

$$(p_{jt-1}^w - c_{jt}^w) s_{jt}(p_{jt-1}^r, p_{-jt}^r) + \sum_{k \in S_w} (p_{kt}^w - c_{kt}^w) s_{kt}(p_{jt-1}^r, p_{-jt}^r) \geq (p_{jt}^{wc} - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) + \sum_{k \in S_w} (p_{kt}^w - c_{kt}^w) s_{kt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) - A_{jt}^w, j \neq k \quad (18)$$

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

$$A_{jt}^w \geq \underline{A_{jt}^w} = (p_{jt}^{wc} - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) - (p_{jt-1}^w - c_{jt}^w) s_{jt}(p_{jt-1}^r, p_{-jt}^r) + \sum_{k \in S_w} [(p_{kt}^w - c_{kt}^w) (s_{kt}^c(p_{jt}^{rc}, p_{-jt}^{rc}) - s_{kt}(p_{jt-1}^r, p_{-jt}^r))], j \neq k \quad (19)$$

The main complication in deriving the lower bounds is that to compute counterfactual prices and market shares, we need to know the manufacturer's belief about the retailer's reaction if the wholesale price changes. Given the existence of retailer adjustment costs, it is in principle possible that a product's retail price will not adjust 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 again rely on the patterns observed in our data and assume that whenever the wholesaler adjusts, the retailer adjusts too and sets a counterfactual retail price p_{jt}^{rc} that is determined based on the FOC of the retailer. In

adjustment cost; if the price of another product had been different, the optimal price calculated by the retailer would have been different too).

this case p_{jt}^{wc} is determined according to equation (15) 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. As with the derivation of the upper bounds, we emphasize that the justification for this strategy is empirical, but the patterns in the data are striking. *In principle*, given the presence of retailer adjustment costs, it is conceivable that the retailer would find it optimal not to change the retail price when the manufacturer sets a new wholesale price. However, we do not have *a single instance* in our data set where we observe such a pattern. Accordingly, we believe that it is reasonable to assume that the manufacturer believes that whenever she adjusts the price of a product, the retailer will adjust her price too.

Finally, regarding the prices of the other products sold by the retailer, the p_{-jt}^r ’s, we employ the same procedure described earlier under the derivation of upper bounds. If a product’s retail price did not change from the previous period, then we keep its price constant in the counterfactuals (the justification again being that we never observe retail prices changing in response to wholesale price changes of *other* products). If, on the other hand, a product’s retail price differs from its previous period’s price, then we allow the retailer to adjust its price optimally in response to the counterfactual wholesale price set by the manufacturer (again, the idea being that the retailer, who already adjusted the product’s price, would have set a different price had she been confronted with a different wholesale price for product j .)

3.2 Pass-through

To assess the overall impact of these adjustment costs on firms’ pricing behavior we employ simulations.¹⁸ We first compute the industry equilibrium that would emerge if a firm faced an exchange-rate shock and prices were fully flexible and, thus, all adjustment costs were equal to zero. In a second set of simulations, we compute 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.

Our first counterfactual simulates the effect of a shock to foreign firms’ marginal costs (i.e., an exchange-rate shock) on all firms’ wholesale and retail prices by computing a new Bertrand-Nash equilibrium. Formally, suppose that an exchange-rate shock hits the traded component of the j th

¹⁸As noted above, the price-adjustment-cost 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.

product's marginal cost (the component 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 solves the system:

$$p_t^{w*} = c_t^{w*}(tc_t^{w*}, ntc_t^w) - \Omega_{wt}^{-1} s_t^*(p_t^{r*}(p_t^{w*})) \quad (20)$$

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:

$$p_t^{r*} = p_t^{w*} + ntc_t^r - \Omega_{rt}^{-1} s_t^*(p_t^{r*}) \quad (21)$$

After running each simulation, we use the new equilibrium wholesale and retail prices to compute pass-through elasticities. The pass-through elasticity of the exchange-rate shock to the wholesale price after accounting for manufacturer nontraded costs is $(d \ln(ntc_j^w + tc_j^w)/d \ln tc_j^w)$ and after accounting for manufacturer markup adjustment is $(d \ln p_j^w/d \ln tc_j^w)$. The pass-through elasticity to the retail price after accounting for its non-traded costs is $(d \ln(p_j^w + ntc_j^r)/d \ln tc_j^w)$ and after accounting for its markup adjustment is $(d \ln p_j^r/d \ln tc_j^w)$. The decomposition then computes the contributions of the manufacturers' and retailer's non-traded costs, markup adjustment, and adjustment costs to the $1 - (d \ln p_j^r/d \ln tc_j^w)$ part of the original shock not passed through to the retail price, as we describe in Appendix A.

3.3 Demand

The estimation of marginal costs, markups, and adjustment costs requires consistent estimates of the demand system as a first step. Market demand is derived from a standard discrete-choice model of consumer behavior. Given that the credibility of our results ultimately depend on the credibility of the demand system, it is imperative to adopt as general and flexible an approach 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 curvature of demand. We provide here 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} of consumer i from consuming product j at time t take the quasi-linear form:

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

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, consumer, and 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: $\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i$ where D_i is a vector of demographics for consumer i , and Π is a matrix of coefficients that characterize how consumer tastes vary with demographics. 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 $\mu_{ijt} = [\Pi D_i] * [p_{jt} \ x_{jt}]$:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}. \quad (23)$$

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 (or not to purchase at all). The price of the outside good is assumed to be set independently of the prices observed in the sample.¹⁹ The mean utility of the outside good is normalized to be zero and constant over markets. 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: $s_{jt} = \int_{\zeta \in A_j} P^*(d\zeta)$ where $P^*(d\zeta)$ is the density of consumer characteristics $\zeta = [D]$ in the population. To compute this integral, one must make assumptions about the distribution of consumer characteristics. We report estimates from two models. For diagnostic purposes, we initially restrict heterogeneity in consumer tastes to enter only through the random shock ε_{ijt} , which is assumed to be i.i.d. with a Type I extreme-value distribution. In the full random-coefficients model, we make the same assumptions about ε_{ijt} but also allow heterogeneity in consumer preferences to enter through an additional term μ_{ijt} . This allows for a more flexible specification of the demand curvature and substitution patterns than permitted under the restrictions of the multinomial logit model. In the full random-coefficients model, the market share function 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:

$$s_{jt} = \frac{1}{N} \sum_{i=1}^N \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} \quad (24)$$

¹⁹The 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 local market. The effect of local conditions is accounted for by the presence of this outside good. When computing the price elasticities for each brand, we consider that consumers have the option of going to other retail outlets to purchase beer. In equilibrium, the retailer and manufacturer decide how much to raise the price of a brand following a foreign-cost shock after taking these elasticities into account. If consumers switch to domestic brands not in our sample or purchase beers in another supermarket following a rise in *Dominick’s* prices, our model will produce consistent estimates of pass-through elasticities. One limitation of this method of deriving the market’s aggregate demand curve is that one must assume the price of the outside good remains constant, which does not allow for strategic interactions between our retailer and other retailers in the same market.

To estimate the demand parameters, we equate these predicted market shares to the observed market shares and solve for the mean utility across consumers using a nonlinear generalized method-of-moments procedure, following Berry et al (1995) and Nevo (2000), as we discuss further in Section 4.1.

3.4 Discussion of the Main Assumptions

We discuss here several key assumptions of our model that simplify both the consumer’s and the firm’s problems by abstracting from various dynamic considerations.

Demand: Inventories and Sales An important premise of our analysis is that consumers do not hold inventories of beer and that the price reductions (sales) documented in Figures 1 and 2 are not systematically related to stockpiling behavior. Otherwise, our static demand estimates would overstate the long-run price elasticities of demand (see Hendel and Nevo, 2006a and 2006b). As we note in the Introduction, the industry wisdom is that the typical buyer consumes beer within hours after purchase, so stockpiling is not a first-order concern. To investigate further if sales exhibit a systematic pattern in our data, 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 elapsed since a previous sale, holiday dummies, or sales for other brands could predict a current sale for a particular brand. Our results contrast with those of Chevalier, Kashyap, and Rossi (2003, hereafter, CKR) who analyzed category-level (e.g. “beer”) price indexes constructed from the same *Dominick’s* dataset we use and concluded that the average prices of beer and other seasonal products declined during holidays and other peak demand periods due to more frequent retail promotions. Our findings appear more supportive of the explanation offered by Nevo and Hatzitaskos (2005) who, like us, analyze the *Dominick’s* data at a more disaggregated (UPC) level than do CKR, and find that shifts in the relative demand for individual brands combined with changes in overall price sensitivities account for the lower average prices in peak demand periods, rather than increased promotions.

This apparent randomness of sales in our sample is also 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 (Chung, 2009). We emphasize that we by no means claim that the patterns we observe in our beer data generalize to all other product categories. In fact, there is substantial evidence from Pesendorfer (2002) and Hendel and Nevo (2006a; 2006b) that intertemporal considerations, such as those emphasized in models by Sobel (1984) and Pesendorfer (2002), are important in determining sales for 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.

Supply: Inventories A key assumption we make on the supply side is that firms set prices to maximize profits on a weekly basis in a static framework. This in turn presumes they do not hold inventories. To the extent that the retailer does hold inventories of beer, the supply-side estimates obtained by this static approach may be biased. As we note in the Introduction, local and state regulations regarding the distribution of alcohol make this concern second-order in our case. 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 observe the sharp “round-trip” drops in our wholesale price graphs.

Supply: Myopic Expectations Perhaps the strongest assumption in our model is that firms maximize profits myopically, period by period. Even with static demand and supply, this assumption may seem implausible given fixed costs of repricing, which create an option value to *not* adjusting prices in the presence of ongoing uncertainty, as rational firms understand they will be stuck at the same price for a while, and so the value in waiting to set a price that is optimal over a longer horizon. And as shown by Dixit (1991), rational expectations imply that even very small costs of adjustment can generate significant price inertia.

While the myopia assumption may seem implausible in a general setting, we believe that it is rather inconsequential in the context of *exchange-rate* pass-through and more generally, in other markets with very persistent cost shock processes. The consensus among International economists is that exchange rates follow a random walk. This finding suggests that cost shocks associated with exchange-rate movements are unlikely to be anticipated in advance, either in magnitude or direction, and will be highly persistent. In a market where exchange-rate movements are the primary source of uncertainty, the assumption that firms believe the world will stay the same, at least for the near future, is not as important for our conclusions as it may at first seem. Of course, persistence is not the same as permanence, and in that sense, our assumption that firms set prices under the premise that shocks will be the same in future periods, not just in expectation, but in realization, is strong, and one that we make primarily out of necessity (estimation of a full dynamic model with rational expectations is infeasible).²⁰

Leaving aside the particular features of our market, the most straightforward way to provide support for our model is to relax the myopia assumption explicitly to show that it does not matter, in the sense that it does not affect the bottom line of our results. While estimation of a dynamic model

²⁰In addition to exchange rates, firms face other sources of cost variation (energy prices, the price of barley, etc.) but the variation induced by these factors pales in comparison to that from exchange rates, as we show in Figure 1 of Goldberg and Hellerstein (2007).

with rational expectations is not feasible, we can consider the extreme opposite of myopia, namely, perfect foresight about the future. In Appendix B, we report the results from a set of robustness checks that allow firms to optimize over longer horizons while maintaining demand as in our baseline weekly model. We assume full information and provide firms with perfect foresight about the realized values of all variables in the model, including their own and competitors' adjustment costs. We run versions of this alternative model in which firms optimize over two, four, or six weeks, finding results quite similar to our baseline results. This is not surprising, as the market does not change much over four to six weeks, and the source of most of the change, the exchange rate, is very difficult to predict, and quite persistent in any realized shocks. In this sense, we believe that it is exchange rates' unusual degree of persistence that underlies this robustness of our baseline results to relaxing the firm myopia assumption. This robustness indicates that our myopia assumption is a reasonable approximation of a finite-horizon dynamic problem, which we believe in turn approximates the true infinite horizon problem.

Finally, we have computed the ex-post errors firms make following a period of price adjustment by not adjusting their prices (that is the losses firms incur in one week by not changing the price they have set in an earlier period) and examined the correlation between these ex-post errors and observables (exchange rates, local wages, etc.) at the time the price was last changed. We find no evidence that the ex-post errors are correlated with observables at the time of the price change. This result provides support for our main premise that shocks firms experience subsequent to a price change were not forecastable at the time the price was set and therefore did not affect the choice of price.²¹

4 Empirical Implementation and Results

Our empirical approach has two components: estimation and simulation. In 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 on the price-adjustment costs. In the simulation stage, we use these estimates to perform counterfactual simulations and to decompose the incomplete transmission of exchange rate shocks to prices. This section describes each component in turn.

4.1 Estimation

4.1.1 Demand estimation

The estimation of the demand system follows Hellerstein (2008). The demand parameters are identified from plausibly exogenous variation in relative prices across products over time, generated though

²¹Dynamics may also enter the demand estimation from intertemporal effects stemming, for example, from habit formation, as explored by Froot and Klemperer (1989), Slade (1998), and Ravn, Schmitt-Grohe and Uribe (2010). We discuss this issue further in Appendix C.

changes in input prices and bilateral exchange rates. Consumers choose between individual products over time, where a product is defined as a bundle of characteristics, one of which is price. As prices are not randomly assigned, we use input price changes that are significant and exogenous to unobserved changes in product characteristics to instrument for prices. Formally, to estimate our model’s demand parameters we equate the predicted market shares from equation (24) to the observed shares, and solve for the mean utility across all consumers using a nonlinear generalized method-of-moments (GMM) procedure, following Berry et al (1995) and Nevo (2000). We model the mean utility associated with product j at time t as follows:²²

$$\delta_{jt} = \beta d_j - \alpha p_{jt} + \Delta\xi_{jt} \quad (25)$$

where the product fixed effects d_j proxy for both the observed characteristics x_{jt} as defined 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, a set of variables z_{jt} that are correlated with the product price p_{jt} but 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 independently to a given supply shock.

Table 3 reports the results from the demand estimation using the multinomial logit model. Following the specification in equation (25), we regress the mean utility, δ_j , which for the logit model is defined as $\ln(s_{jt}) - \ln(s_{Ot})$, on prices and product dummies. Due to its restrictive functional form, the multinomial logit model will not produce credible estimates of pass-through. It is helpful, however, to gauge how well the instruments for price perform in the logit demand estimation before turning to the full random-coefficients logit demand model. The table’s first two columns report the ordinary least squares (*OLS*) estimate for the mean price coefficient, α , and columns 3 and 4 for two instrumental variables (*IV*) specifications, using as instruments for prices manufacturer factor prices interacted with exchange rates and product fixed effects. 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 therefore promising that the mean price coefficient falls from about -0.92 in each *OLS* estimation to -2.43 in the two *IV* estimations. Columns 2 and 4 report results when a dummy for major holidays is included as an additional right-hand-side

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

variable, as one might expect the coefficient on price to differ in periods of high demand for beer, such as during major holidays. Including the holiday dummy as a right-hand-side variable does not affect the demand coefficients in either the *OLS* or the *IV* estimation. Finally, Table 3 reports that the first-stage partial *F*-stat of the two *IV* specifications is high (over 34 in both cases), and accordingly the *F*-test for zero coefficients associated with the instruments is rejected, suggesting that factor costs interacted with exchange rates are valid instruments.

Table 4 reports results from the estimation of the full random-coefficients logit demand system, in which 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 (bitterness and percentage alcohol). The coefficients on the characteristics generally appear reasonable. The random coefficient on income, at 38.45, is significant at the five-percent level, which implies that higher-income consumers are less price sensitive. Consistent with industry lore, the mean preference in the population is amenable to a bitter taste in beer, which has a positive and significant coefficient. As the percentage alcohol rises across brands, the mean utility in the population also rises, an intuitive result, though higher-income consumers' are less keen on this characteristic than the average consumer, which is evident in the negative and significant random coefficient of -50.36, and consistent with industry lore regarding the typical consumer of light beer.²³

The random-coefficients logit demand system is very flexible in the dimension that matters most for a pass-through analysis, the curvature of demand. This flexibility is *not* available with the multinomial logit, and in theory can accommodate a range of elasticities and super-elasticities from CES to the Kimball (1995)-style kinked demand curve used in some of the macroeconomic literature (e.g. Klenow and Willis, 2006; Dotsey and King, 2005). The model's parameters that characterize heterogeneity in consumers' tastes for various product characteristics, particularly their price sensitivity, also contribute to the curvature of demand. As a firm raises its product's price, more price-sensitive consumers will respond by not purchasing the product or by dropping out of the market altogether (purchasing the outside good), meaning the firm will retain only its less price-sensitive consumers and aggregate demand will appear more inelastic.

One common measure of demand curvature is the "super-elasticity of demand," the percentage change in the demand elasticity for a given percentage-point change in prices (Klenow and Willis, 2006; Kimball, 1995). A Dixit-Stiglitz demand model has a super-elasticity of zero which generates constant markups under monopolistic competition. A positive super-elasticity of demand implies a concave demand curve: As a firm increases its price, it faces increasingly elastic demand. We estimate

²³We show in Appendix D that our demand system is robust to the Knittel and Metaxoglou (2008) critique that in highly nonlinear models such as random-coefficients demand systems, the objective function may exhibit many local minima.

the super-elasticity of demand to be 0.8 in our random-coefficients model, in other words, we compute that a 10-percent increase in prices leads to an 8-percent increase in the absolute value of the price elasticity of demand. We have also conducted robustness checks for the super-elasticity measure, using alternative versions of the demand system with additional random coefficients, and found the super-elasticity coefficient to be generally close to 1, which generates a reasonable incentive for firms to adjust their markups.

4.1.2 Computation of total retail costs, non-traded retail costs ntc_{jt}^r , and retail markups

Once we've estimated the parameters of the demand system, we know the market share function $s_{jt}(p_t^r)$ as well as the own- and cross-price derivatives $\frac{\partial s_{kt}}{\partial p_{jt}^r}$. We then use the retailer's first-order conditions for each product j in equation (6) to back out each product's markups, which in turn enables us to calculate the total marginal costs, including the non-traded retail costs, of product j . Retail non-traded costs are given by the difference between the retail marginal costs and the (observed) wholesale prices. Under the assumption of no adjustment costs, these 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 only for those products whose prices adjust. 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 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 incur some cost every time they change the posted price); yet, given that firms charge *exactly* the same price as before the promotion, assuming these post-promotion prices are determined by firms' first-order conditions seems implausible. To be safe, we conducted the empirical analysis both ways, first applying the FOC's to all periods in which the price changed (including 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, as we discuss in Appendix B, but the second approach significantly reduces the number of observations associated with a price change that are available for 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 periods when a firm charges a price genuinely different from that charged in earlier periods.

Table 5 reports retail and wholesale prices and our estimates of markups and costs for selected foreign brands. The markups appear reasonable and consistent with industry wisdom, on the order of 5-10 percent of the retail price, which coincides with industry estimates of supermarkets' average markups on beer (Tremblay and Tremblay, 2005).

4.1.3 Estimation of non-traded retail cost function

The procedure described above allows us to back out the retailer’s non-traded costs for the periods in which we observe a product’s price adjust, so we can reasonably assume that the retailer sets the new price according to its first-order conditions. However, this approach does not work for periods in which the price does not change. To get estimates of the non-traded costs for these periods, we collect the non-traded costs ntc_{jt}^r from the periods in which the price of product j adjusted, and model them 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 this regression using data from the periods with price adjustments, with the coefficients reported in Table 6, and then use the parameter estimates to compute predicted non-traded costs for the periods without price adjustments. The results reported in Table 5 indicate that the retailer’s average non-traded costs for each product ranges from about 30 to 40 cents for both the backed-out and the fitted series, which is a fairly narrow band given its average price of \$5.52 for a six-pack of imported beer. One exception is *Bass*, with an average retailer non-traded cost closer to 25 cents. As the retailer’s non-traded costs are computed as the difference between the structural model’s derived total costs for the retailer and the observed wholesale price, we expect that these cross-product differences primarily reflect measurement error in our model’s estimates of the retailer’s total marginal costs.²⁴ In addition, Tremblay and Tremblay (2005) report significantly lower advertising expenditure per barrel by *Bass* relative to *Corona* and *Heineken* (one to two orders of magnitude difference) over the sample period. To the extent that more promotional activity requires more labor inputs from the supermarket as a complement, this would explain the different estimated sizes of non-traded costs across the brands.

4.1.4 Derivation of bounds for the retailer price adjustment costs A_{jt}^r

With the demand parameters and non-traded cost estimates in hand, we employ equations (9) and (11) to derive the upper and lower bounds of the retailer adjustment costs A_{jt}^r . The computation of the upper bound is straightforward: in equation (9) all the variables are observed, except for the counterfactual market share $s_{jt}^c(p_{jt-1}^r, p_{-jt}^r)$ that product j would have attained had the retailer *not* changed her price from the previous period. This counterfactual share can easily be evaluated once the demand parameters are estimated, given that the market share function is known. Note that when computing profits in each of our counterfactual scenarios (for both the retailer and manufacturer upper and lower bounds), we assume that demand shocks are unchanged. Computing the lower bound based

²⁴Some products may incur higher per-unit energy costs if they are refrigerated instead of “dry shelf” (not refrigerated in industry parlance) though we do not have data on this for our market.

on equation (11) requires calculating the counterfactual optimal price p_{jt}^{rc} the retailer would charge if she changed the retail price from the previous period and the associated market share $s_{jt}^c(p_{jt}^{rc}, p_{-jt}^r)$. These are computed using equation (7) which reflects the first-order condition of the retailer.

Table 7 reports the mean estimates of the upper and lower bounds on the retailer’s adjustment costs for selected foreign brands. The upper and lower bounds generally are consistent for each brand as well across the brands. The mean lower bounds on adjustment costs range from \$30.18 for *Bass* to \$148.30 for *Heineken*, with a mean lower bound across foreign brands of \$61.52, and a mean upper bound of over \$230. The entries in the third and fourth columns report the sum of the upper or lower bounds for each brand’s price-adjustment costs divided by the retailer’s total revenue from that brand over the full sample period. These numbers are more comparable to those of the Levy et al (1997), Dutta et al (1999), and Klenow and Willis (2006) studies, which divide the 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 not. The sum of the upper bounds to repricing costs across all foreign brands is 3.16 percent of total revenue and for the lower bounds is 0.84 of total revenue.²⁵ 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 model, and are accordingly meant to be used only in the context of our model to perform simulations. Nevertheless, we find it interesting that the order of magnitude of our estimates is similar to that obtained in other studies that used a completely different methodology.

4.1.5 Computation of manufacturer marginal costs, c_{jt}^w , and manufacturer markups

This procedure is similar to that used to derive the retailer’s non-traded costs. In periods when the wholesale price changes, manufacturers act according to their first-order conditions, so we can use equation (15) to back out their marginal costs c_{jt}^w . Manufacturers’ mean markup from our model is 47 cents, which given the mean retail price of \$5.52 across our sample’s imported brands, is between 5 and 10 percent of the retail price, which matches precisely the industry estimates reported in Tremblay and Tremblay (2005) and Consumer Reports (1996).

4.1.6 Estimation of manufacturer marginal cost function

The first-order conditions allow us to back out the manufacturers’ *total* marginal costs, but do not tell us how to decompose it into a traded and non-traded component. Further, it is not possible to back out the marginal manufacturer costs in periods when wholesale prices do not adjust, given that

²⁵As our repricing costs are defined in the most general sense to include *all* factors that may prevent firms from changing their prices (*not just the literal labor and material costs* of changing prices), our numbers are not directly comparable to Levy et al’s (1997). Still, it is interesting that despite these differences the two sets of numbers are of similar order of magnitude.

firms' first-order conditions do not necessarily hold. To do this, we model total manufacturer costs parametrically as a function of observables, and estimate this function using data from the periods of wholesale price adjustment *only*. We assume manufacturers' marginal costs c_{jt}^w take the form:

$$c_{jt}^w = \exp(\theta_j + \omega_{jt})(w_t^d)^{\theta_{dw}}(e_{jt}w_t^f)^{F_j*\theta_{fw}}(p_{bjt})^{D_j*\theta_{dp}}(e_{jt}p_{bjt})^{F_j*\theta_{fp}} \quad (26)$$

or, in log-terms:

$$\ln c_{jt}^w = \theta_j + \theta_{dw} \ln w_t^d + F_j * \theta_{fw} \ln(e_{jt}w_t^f) + D_j * \theta_{dp} \ln(p_{bjt}) + F_j * \theta_{fp} \ln(e_{jt}p_{bjt}) + \omega_{jt} \quad (27)$$

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 (27) can be easily estimated by Least Squares which serves two purposes. It allows us to, first, decompose the total marginal cost into a traded and a non-traded component, and second, to use its parameter estimates to construct predicted values for the manufacturer traded and non-traded costs for the periods when wholesale prices do not change.

Recall that by definition the traded component refers to the part of the marginal cost that is paid in foreign currency and hence 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 equation (26) 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, if the local non-traded cost component is nonzero (so that $\theta_{fw} + \theta_{fp} < 1$), 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 prices to exchange-rate shocks.

Our estimate of the "local content" of foreign manufacturers' marginal cost is reflected in the magnitude of the "domestic U.S. wages" coefficient that captures the cost share accounted for by domestic labor. As reported in Table 8, the highly significant coefficient of 0.60 indicates the share of local costs appears to be substantial and is consistent with external evidence on the share of local

costs in U.S. consumer goods.²⁶ This estimate implies that a big part of foreign manufacturers’ costs of selling in the U.S. market is 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.

4.1.7 Derivation of bounds for the wholesale price adjustment costs A_{jt}^w

The final step is to use the parameter estimates obtained in the previous steps to compute the upper and lower bounds of the manufacturer price-adjustment costs based on equations (17) and (19). The estimates of the manufacturers’ adjustment-cost bounds, reported in Table 9, are roughly the same order of magnitude as those for the retailer, with upper bounds ranging from \$36.76 for *Bass* to \$310.44 for *Beck’s*, and lower bounds from \$14.42 for *Bass* to \$186.14 for *Corona*. The entries in the third and fourth columns report the sum of the upper or lower bounds for each brand’s price-adjustment costs divided by the manufacturer’s total revenue from that brand over the full sample period. The sum of the upper bounds of repricing costs across all foreign brands is 1.28 percent of total revenue and for the lower bounds is 0.64 percent of total revenue.

Recall that in our approach, adjustment costs are a “catch-all” term that captures anything that may induce a firm not to change its price in a particular period - including concerns that it may lose customers in the future (in this sense this adjustment cost could be interpreted as a shorthand for dynamic considerations that are not accounted for in the static framework) or the option value of not changing the price. The advantage of thinking about adjustment costs this way is that it allows us to remain agnostic about their nature. As we state above, we consider them in a sense a “residual” explanation that is needed to justify periods with no adjustment. That said, one interesting and robust feature of our results is that we consistently find that the price-adjustment costs associated with promotional price changes are substantially lower than those associated with changes in the regular price, which suggests that to a certain extent these costs proxy for the managerial costs incurred to figure out new prices. Table 10 reports the results of a fixed-effects panel regression of the derived retail repricing costs on a dummy for a level change in a brand’s price as well as a dummy for sales, that is, temporary price reductions. Adjustment costs appear to be significantly higher for level changes in prices, averaging about \$850 for a “permanent” level price change, and closer to \$50 for a temporary price reduction. This finding is consistent with Kehoe and Midrigan (2010) who argue that the fixed cost of changing a regular price is larger than that of a temporary reduction.

²⁶Burstein, Neves, and Rebelo (2003), using data from the Bureau of Economic Analysis’s input-output tables and the 1992 Census of Wholesale and Retail Trade, and Goldberg and Campa (2010), using data from the OECD’s input-output tables, find that local distribution services account for roughly half of the retail price of the average consumer good in the U.S. Similarly, Goldberg and Verboven (2001) find that local costs account for circa 65 percent of the incomplete pass-through in the European auto industry.

4.2 Simulations

4.2.1 Counterfactuals

Using the full random-coefficients model and the derived measures of traded, non-traded, and repricing costs, we conduct a series of counterfactual experiments to assess how firms react to exchange-rate shocks. We consider the effect of a five-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 first compute the industry equilibrium that would emerge if a foreign 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 these cases as a measure of the impact of nominal price rigidities.

In all of the counterfactuals involving price-adjustment costs we use the mean of each brand's estimated upper bounds as our measure of the price-adjustment costs. We use the upper bounds because we want to "bias" in some sense our findings towards assigning the largest possible effect to nominal price rigidities. This is because, as we explained in the Introduction, we developed the current approach to address the criticism that earlier work, by assuming that firms always set prices optimally, had ignored nominal price rigidities (notably Engel, 2002). Hence, we want to attribute to nominal price rigidities the largest possible role, and see how this affects our findings. The bottom line is that in the end, despite using the upper bounds on the price-adjustment costs, we still find that local non-traded costs are the most significant source of price inertia.

The counterfactual experiments consider the effect of a five-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 11. There are three panels in the table, each one corresponding to one of the simulations we describe below. For each counterfactual, we report the median pass-through elasticity across the 404 markets in the sample. The first column of the table reports for each counterfactual 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 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 pass-through of the original shock to the retail price due to the retailer's markup adjustment. Given the Cobb-Douglas specification for the manufacturers' marginal costs in equation (27), 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 manufacturers' pass-through elasticity and that attributed to non-traded costs will reflect 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 of such costs 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 retailer’s pass-through elasticity and the elasticity attributed to its non-traded costs will capture the markup adjustment by the retailer. (We discuss the expressions used in the decomposition formally in Appendix A.)

Simulation 1: Simulate the effect of a 5-percent appreciation of the relevant foreign currency assuming that all prices are fully flexible. The first counterfactual experiment examines the manufacturers’ and the retailer’s pass-through following a 5-percent appreciation of the relevant foreign currency when they face no repricing costs. Its results are reported in the top panel of Table 11. The median pass-through of the exchange rate shock to manufacturer’s total marginal cost is 40 percent, which is determined by the coefficient on local wages (0.60) from the regression results reported in Table 8. As the average non-traded cost incurred by a foreign manufacturer is over 50 percent of her total costs, a nontrivial amount of non-traded value is added at this stage of the distribution chain. Next, manufacturer markup adjustments are substantial in this counterfactual. Once these are accounted for, the median pass-through elasticity of the exchange-rate shock to the wholesale price ranges from 30.3 percent for *Bass* to 38.0 percent for *Beck’s*. It is 31.7 percent across all brands. It is striking that our median wholesale pass-through elasticity across foreign brands is almost identical to that of Hellerstein (2008) at 32.0 percent, which uses a similar dataset on beer from *Dominick’s*, but aggregated up to a monthly frequency. Accounting for retail non-traded costs, the median pass-through falls to 28.4 percent and ranges from 27.4 percent for *Bass* to 33.4 percent for *Beck’s*. Finally, the median retailer pass-through elasticity across all brands is 26.1 percent. This counterfactual reveals that the curvature of demand is such that the retailer passes through the bulk of the cost shocks it experiences to its prices, rather than adjusting its markups. The median pass-through of the original shock to the retail price ranges from 24.4 percent for *Bass* to 29.3 percent for *Beck’s*. Note that these results are larger than the 5-10 percent retail pass-through elasticities estimated in Section 2.2.

We turn next to 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 its competitors will adjust their prices, and a second in which the firm under consideration assumes that its competitor prices will remain fixed due to adjustment costs. These cases correspond to the following two simulations:

Simulation 2: Simulate the effect of a 5-percent appreciation of the relevant foreign currency assuming that the foreign brand experiencing the exchange-rate shock faces no price adjustment costs but assumes that its competitors’ prices will remain fixed. In this

case, the new industry equilibrium is computed as in Simulation 1, but by additionally imposing that all other products' 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. In this simulation, competitive pressure from other manufacturer repricing costs for *Bass* prevent the wholesale price of this brand from adjusting, which translates to a zero median manufacturer pass-through elasticity. Another interesting finding from this counterfactual are the lower pass-through elasticities in those cases where prices do adjust (*Becks*, *Corona*, and *Heineken*) compared to Simulation 1. This additional reduction in the pass-through elasticities also 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. The largest reduction of this nature is for *Heineken*, whose median manufacturer pass-through elasticity falls from almost 38 percent in Simulation 1 to 30.9 percent in Simulation 2, and whose median retail elasticity falls from 28.2 percent in Simulation 1 to under 6 percent in Simulation 2, a 22 percentage-point decline attributable entirely to the strategic effects of repricing costs. The median retail pass-through elasticity across all four brands falls from 26.1 percent in Simulation 1 to 18 percent in Simulation 2, a modest but still notable 8 percentage-point drop.

Simulation 3: Simulate the effect of a 5-percent appreciation of the relevant foreign currency assuming that the foreign brand affected by the exchange-rate shock also incurs fixed repricing costs. This final counterfactual experiment considers how manufacturers and the retailer adjust their prices following a 5-percent appreciation of the relevant foreign currency if they must incur fixed repricing costs to alter their prices. As discussed earlier, we use the derived upper bounds on manufacturers' and the retailer's price-adjustment costs in this final set of counterfactuals, whose results are reported in the bottom panel of Table 11. The median pass-through of the exchange-rate change to manufacturers' total marginal costs is again 40 percent as the share of non-traded costs is unaffected by the nature of the counterfactual. But the manufacturer pass-through elasticities are now zero across brands. Thus, accounting for a brand's own price-adjustment costs reduces the median manufacturer pass-through elasticity from 30.4 percent in Simulation 2 to 0 percent in Simulation 3. This reduction is due to the zero transmission of the exchange-rate shock to the *wholesale* prices of *Becks*, *Corona*, and *Heineken* due to these three brands' *manufacturer* repricing costs. In contrast, retail repricing costs do not contribute directly to the reduction in the pass-through elasticities, though their indirect effects played a role in Simulation 2. These results are consistent with the patterns we documented earlier suggesting that retail prices always adjust whenever wholesale prices adjust in this market. Simulation 3's results may overstate the effects of repricing costs in lowering firms' pass-through, as the counterfactuals use our derived upper bounds as their measures of repricing costs.

In addition, Table 11 reports the median pass-through elasticities across the 404 markets used in the counterfactual, as these are generally more robust than are means in this type of structural analysis. It may be instructive, however, to compare our mean pass-through elasticities to the exchange-rate pass-through elasticities estimated in Section 2.2. The mean retail pass-through elasticity across the four brands in this final counterfactual is 10 percent, which is more comparable, and quite close to the (mean) retail pass-through elasticity of 7 percent estimated via simple regressions in Section 2.2.

4.2.2 Decomposition of the Incomplete Transmission

We next decompose the sources of the incomplete transmission of the exchange-rate shock to retail prices that is documented in Table 11. The first column of Table 12 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 effect that the fixed costs of repricing faced by competitors have on the manufacturer and retailer's pricing behaviors (the indirect or strategic effect). Columns four and eight 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). 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 5-percent appreciation of the relevant foreign currency, it is responsible for roughly half, or 60 percent, of the observed retail-price inertia. Manufacturers' markup adjustment accounts for 8.3 percent of the remaining adjustment, their competitors' price rigidities for 1.3 percent, and their own repricing costs for another 30.4 percent. As we noted above, the decomposition varies across brands by their market share. The brand with the smallest market share, *Bass*, exhibits the greatest impact from other brands repricing costs on its pass-through, at 30.2 percent. The results across the other three brands are quite similar: After accounting for manufacturers' own repricing costs, the retailer's markup adjustment and own repricing costs can only play a negligible role in explaining the incomplete transmission. 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. It is important to emphasize that without incorporating manufacturer repricing costs into our approach, however, we would conclude that local costs accounted for about 85 percent of the incomplete pass-through at the wholesale level.

We find, then, that the presence of a local, non-traded component in firms' total costs seems to be the primary source of the incomplete transmission; local costs however cannot explain why prices remain completely unchanged in several periods. To explain the latter, *complete* inertia, we need to incorporate nominal price rigidities into the model. Our model does therefore a good job in generating the pass-through patterns observed in the data.

5 Conclusion

This paper set out to develop and estimate a model that can be used to identify the determinants of local-currency price stability in the face of exchange-rate fluctuations. The empirical model we develop 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 to the fixed costs of price adjustment facing retailers and manufacturers, we find these costs are roughly of the same order of magnitude for manufacturers and retailers in absolute terms, though smaller for manufacturers as a share of their total revenue. Third, the counterfactual simulations we conduct to decompose the incomplete transmission of exchange-rate shocks into its sources suggest that both local non-traded costs and a firm's own repricing costs are important in generating local-currency price stability. We find little evidence of (indirect) strategic effects of price adjustment costs: Their aggregate effects are fairly subtle, though they can have dramatic effects on individual brands' pricing behavior. Markup adjustment appears more important at the manufacturer than the retail level, accounting on average for circa 8% of the incomplete price response. Intuitively, these results are driven by the fact that in the data we observe many periods during which prices remain completely unchanged; this *complete* inertia can only be accounted for by nominal price rigidities. But the data also indicate that conditional on prices changing, the response of prices to exchange rates is small. This incomplete response, conditional on adjustment, is attributed primarily to local non-traded costs. Markup adjustment is present, but not sufficient to rationalize the small size of price adjustments. Repricing costs affect primarily the adjustment of wholesale prices; their direct effect on retail prices is very minor. Why nominal price rigidities operate 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 on 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.
- [3] Aguirregabiria, V. 1999. "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, vol. 66, pp. 275-308.
- [4] Akerlof, G. and J. Yellen. 1985. "A Near-Rational Model of the Business Cycle with Wage and Price Inertia," *The Quarterly Journal of Economics*, vol. 100, pp. 823-38.
- [5] Alessandria, G., J. Kaboski, and V. Midrigan. 2010. "Inventories, Lumpy Trade, Large Devaluations," *The American Economic Review*, vol. 100, pp. 2304-39.
- [6] Asker, J. 2004. "Diagnosing Foreclosure Due to 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., 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] Caplin, A. and J. Leahy. 1997. "Aggregation and Optimization with State-Dependent Pricing," *Econometrica*, vol. 65, pp. 601-26.
- [11] Carlton, D. 1986. "The Rigidity of Prices," *The American Economic Review*, vol. 76, pp. 637-58.
- [12] Chevalier, J., A. Kashyap, and P. Rossi. 2003. "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *The American Economic Review*, vol. 93, pp. 15-37.
- [13] Chung, B. 2009. "Menu Costs and Sale Strategies of Multiproduct Retailers," Working Paper, Chinese University of Hong Kong.
- [14] Dixit, A. 1991. "Analytical Approximations in Models of Hysteresis," *Review of Economic Studies*, vol. 58, pp. 141-51.
- [15] Dotsey, M. and R. King. 2005. "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics*, vol. 52, pp. 213-42.

- [16] Dutta, S., M. Bergen, D. Levy, and R. Venable. 1999. "Menu Costs, Posted Prices, and Multi-product Retailers," *Journal of Money, Credit, and Banking*, vol. 31, pp. 683-703.
- [17] Engel, C. 2002. "Expenditure Switching and Exchange-Rate Policy," *NBER Macroeconomics Annual*, pp. 231-72.
- [18] Froot, K.A., and P.D. Klemperer. 1989. "Exchange Rate Pass-through When Market Share Matters," *The American Economic Review*, vol. 79, pp. 637-54.
- [19] Goldberg, L. and J. Campa. 2010. "The Sensitivity of the CPI to Exchange Rates," *The Review of Economics and Statistics*, vol. 92, pp. 392-407.
- [20] Goldberg, P.K. and R. Hellerstein. 2007. "Sticky Prices: Why Firms Hesitate to Adjust the Price of Their Goods," Federal Reserve Bank of New York *Current Issues in Economics and Finance*, vol. 13.
- [21] 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.
- [22] 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*, vol. 76, pp. 14-32.
- [23] Hellerstein, R. 2004. *Empirical Essays on Vertical Contracts, Exchange Rates, and Monetary Policy*. U.C. Berkeley Ph.D. Dissertation.
- [24] Hendel, I. and A. Nevo. 2006a. "Sales and Consumer Inventory," *Rand Journal of Economics*, vol. 37, pp. 543-61.
- [25] Hendel, I. and A. Nevo. 2006b. "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, vol. 74, pp. 1637-73.
- [26] Kehoe, P. and V. Midrigan. 2010. "Prices Are Sticky After All," Federal Reserve Bank of Minneapolis Staff Report No. 413.
- [27] Kimball, M. S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, vol. 27, pp. 1241-77.
- [28] Klenow, P. and B. Malin. 2010. "Microeconomic Evidence in Price Setting," in B. Friedman and M. Woodford, eds., *Handbook of Monetary Economics, Vol. 3*, pp. 231-84. Amsterdam: Elsevier.
- [29] Klenow, P. and J. Willis. 2006. "Real Rigidities and Nominal Price Changes," Research Working Paper 06-03, Federal Reserve Bank of Kansas City.

- [30] Knittel, C. and K. Metaxoglou. 2008. "Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings," Working Paper, MIT Sloan School.
- [31] Levy, D., M. Bergen, 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.
- [32] Mankiw, N.G. 1985. "Small Menu Costs and Large Business Cycles," *The Quarterly Journal of Economics*, vol. 100, pp. 529-39.
- [33] Midrigan, V. 2010. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," *Econometrica*, forthcoming.
- [34] Nakamura, E. and D. Zerom. 2010. "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, vol. 77, pp. 1192-1230.
- [35] Nevo, A. 2000. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics and Management Strategy*, vol. 9, pp. 513-58.
- [36] Nevo, A. 2001. "Measuring Market Power in the Ready-To-Eat Cereal Industry," *Econometrica*, vol. 69, pp. 307-42.
- [37] Nevo, A. and K. Hatzitaskos. 2005. "Why Does the Average Price of Tuna Fall During Lent?" NBER Working Paper No. 11572.
- [38] Pesendorfer, M. 2002. "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, vol. 75, pp. 33-66.
- [39] Ravn, M., S. Schmitt-Grohe, and M. Uribe. 2010. "Incomplete Cost Pass-through Under Deep Habits," *Review of Economic Dynamics*, vol. 3, pp. 7-35.
- [40] Slade, M. 1998. "Optimal Pricing with Costly Adjustment and Persistent Effects: Empirical Evidence from Retail-Grocery Prices", *Review of Economic Studies*, vol. 65, pp. 87-107.
- [41] Sobel, J. 1984. "The Timing of Sales," *Review of Economic Studies*, vol. 51, pp. 353-68.
- [42] Tremblay, V.J. and C.H. Tremblay. 2005. *The U.S. Brewing Industry*. Cambridge: MIT Press.
- [43] Varian, H. R. 1980. "A Model of Sales," *The American Economic Review*, vol. 70, pp. 651-59.
- [44] Villas-Boas, S.B. 2007. "Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data," *Review of Economic Studies*, vol. 74, 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*.

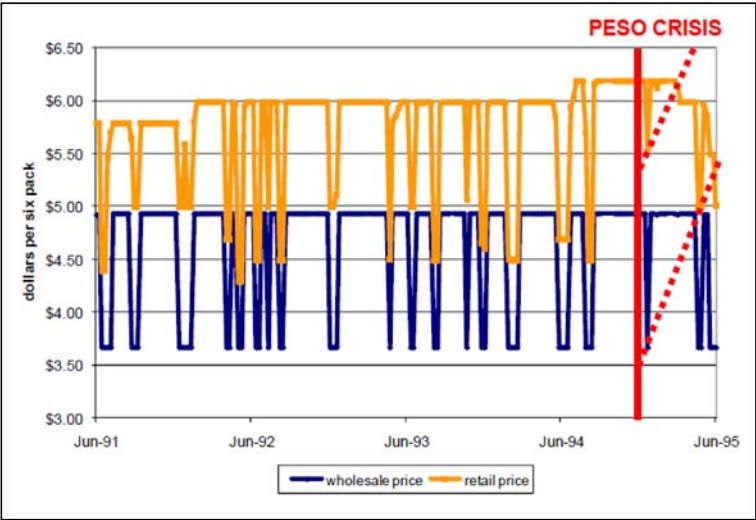


Figure 2: Weekly retail and wholesale prices for Corona. Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's*.

Category	Description	Mean	Median	Std Dev
Prices	Retail prices (\$ per six-pack)	5.44	5.79	1.28
	Wholesale prices (\$ per six-pack)	4.36	4.61	1.00
	Dummy for retail-price change (=1 if yes)	.20	0	.40
	Dummy for wholesale-price change (=1 if yes)	.08	0	.27
Product characteristics	Percent Alcohol	4.62	4.54	.51
	Bitterness	2.87	2.50	1.36

Table 1: *Summary statistics for prices and product characteristics for the 16 products in the sample.* 6464 observations. Source: Dominick's; "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)**
R^2	.65	.65	.81	.81	.80

Table 2: *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 U.S. (dollars per unit of foreign currency). Includes brand, price-zone, and week fixed effects. The second and fourth columns report results with controls for domestic and foreign costs. Robust standard errors in parentheses, those starred significant at the *5-percent or **1-percent level. 3636 observations. 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 Partial F-Statistic			34.45	34.24

Table 3: *Diagnostic results from the multinomial logit model of demand.* The dependent variable is $\ln(s_{jt}) - \ln(s_{0t})$. Regressions include brand fixed effects. Huber-White robust standard errors reported in parentheses: Those starred significant at the *5- or **1-percent level. Instruments: domestic wages in beverage industry interacted with brand fixed effects and with weekly nominal exchange rates for foreign brands. 6464 observations. Source: Authors' calculations.

Variable	Mean in Population	Interaction with Income
Price	-2.48 (.13)*	38.45 (12.20)*
Bitterness	0.81 (.13)*	
Percent Alcohol	0.74 (.01)*	-50.36 (9.14)*
Minimum-Distance Weighted R^2	0.84	

Table 4: *Results from the full random-coefficients logit model of demand.* Asymptotically robust standard errors in parentheses, those starred significant at the 5-percent level. Model includes a constant whose coefficient is not reported. 6464 observations. Source: Authors' calculations.

	Bass	Becks	Corona	Heineken	All Imports
Retailer					
Price	6.36 (0.64)	5.27 (0.48)	5.10 (0.62)	5.66 (0.70)	5.52 (0.90)
Markup	0.44 (0.005)**	0.43 (0.003)**	0.42 (0.004)**	0.42 (0.004)**	0.43 (0.003)**
Nontraded costs					
Backed out	0.25 (0.06)**	0.39 (0.05)**	0.37 (0.08)**	0.46 (0.07)**	0.40 (0.04)**
Fitted	0.22 (0.07)**	0.35 (0.05)**	0.35 (0.09)**	0.44 (0.07)**	0.36 (0.04)**
Manufacturer					
Price	5.76 (0.24)	4.44 (0.16)	4.27 (0.43)	4.99 (0.28)	4.69 (0.77)
Markup	0.54 (0.008)**	0.42 (0.003)**	0.42 (0.003)**	0.55 (0.008)**	0.47 (0.003)**
Total costs					
Backed out	5.25 (0.04)**	4.02 (0.01)**	3.81 (0.08)**	4.57 (0.07)**	4.23 (0.05)**
Fitted	5.18 (0.04)**	4.03 (0.01)**	3.80 (0.07)**	4.47 (0.07)**	4.38 (0.02)**

Table 5: *Mean prices, markups, and costs for selected foreign brands.* In the upper panel of the table, each entry reports the mean across weeks and zones of the retailer's prices, derived markups, and derived backed-out or fitted non-traded costs by brand in dollars per six-pack. In the lower panel of the table, each entry reports the mean across weeks and zones of the manufacturer's prices, derived markups, and derived backed-out or fitted total costs by brand in dollars per six-pack. The markups are price less marginal cost with the marginal costs derived from the structural model. The numbers in parentheses under the prices are standard deviations over the sample, and under the other variables are standard errors from bootstrap simulations with 400 draws. Those starred are significant at the *5- or **1-percent level. Source: Dominick's; Authors' calculations.

Variable	Coefficient
Chicago-Area Grocery Wages	.49 (.05)**
R^2	.13
Observations	805

Table 6: *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. Those starred are significant at the *5- or **1-percent level. 805 observations. Source: Authors’ calculations.

Brand	Mean Cost		Share of Brand’s Revenue	
	Upper Bounds (\$)	Lower Bounds (\$)	Upper Bounds (%)	Lower Bounds (%)
Bass	\$76.04 (31.24)**	\$30.18 (6.22)**	3.05 (1.58)*	0.89 (0.27)**
Beck’s	\$463.24 (214.29)**	\$142.52 (36.27)**	4.25 (2.02)**	1.31 (0.32)**
Corona	\$131.36 (62.04)**	\$62.24 (26.47)**	1.09 (0.43)**	0.52 (0.07)**
Heineken	\$621.06 (239.46)**	\$148.30 (39.53)**	3.80 (1.83)**	0.91 (0.23)**
All	\$230.60 (70.47)**	\$61.52 (11.39)**	3.16 (0.98)**	0.84 (0.15)**

Table 7: *Bounds for the retailer’s adjustment costs for selected foreign brands.* The entries in the first two columns report the mean over time of the dollar value of adjustment costs, and in the third and fourth columns the mean over revenue for that brand over all markets whether the price changed or not. Note that the upper and lower bounds are not obtained over the same sample periods. Standard errors from bootstrap simulations with 400 draws reported in parentheses under each coefficient. Starred coefficients are significant at the *10 or **5-percent level.

Variable	Coefficient
Domestic U.S. wages	.60 (.11)**
Price foreign barley	.17 (.05)**
Foreign wages	.23 (.06)**
Observations	188

Table 8: *Results from constrained linear regression of foreign manufacturer total backed-out costs on determinants.* Dependent variable is manufacturers’ total marginal costs for periods when the wholesale price changes which varies by week. Includes brand and price-zone fixed effects. Starred coefficients are significant at the *5- or **1-percent level. Source: Authors’ calculations.

Brand	Mean Cost		Share of Brand's Revenue	
	Upper Bounds (\$)	Lower Bounds (\$)	Upper Bounds (%)	Lower Bounds (%)
Bass	\$36.76 (20.73)*	\$14.42 (3.56)**	0.80 (0.43)*	0.31 (0.13)**
Beck's	\$310.44 (87.45)**	\$113.08 (36.25)**	2.43 (0.73)**	0.89 (0.32)**
Corona	\$213.48 (42.51)**	\$186.14 (41.36)**	1.41 (0.41)**	1.23 (0.19)**
Heineken	\$122.60 (39.02)**	\$116.66 (12.05)**	0.22 (0.06)**	0.21 (0.07)**
All	\$178.60 (29.31)**	\$89.46 (14.79)**	1.28 (0.24)**	0.64 (0.11)**

Table 9: *Bounds for manufacturers' adjustment costs for selected foreign brands.* The entries in the first two columns report the mean over time of the dollar value of adjustment costs, and in the third and fourth columns the mean over revenue for that brand over all markets whether the price changed or not. Note that the upper and lower bounds are not obtained over the same sample periods. Standard errors from bootstrap simulations with 400 draws reported in parentheses under each coefficient. Starred coefficients are significant at the *10- or **5-percent level.

Variable	Coefficient
Dummy for level change in retail price	853.04 (29.88)**
Dummy for all retail price changes	52.09 (9.21)**
Overall R^2	.19

Table 10: *Regression of retailer's fixed adjustment costs on a level-price-change dummy.* The regression includes brand and price zone fixed effects. 3636 observations. Starred coefficients are significant at the *5- or **1-percent level. Source: Authors' calculations.

	Manufacturer		Retailer	
	Traded	Markup Adjustment	Traded	Markup Adjustment
No repricing costs				
Bass	40.0 (0.0)**	30.3 (7.5)**	27.4 (6.5)**	24.4 (7.4)**
Beck's	40.0 (0.0)**	38.0 (3.4)**	33.4 (4.2)**	29.3 (6.2)**
Corona	40.0 (0.0)**	31.8 (11.8)**	28.0 (7.7)**	25.7 (8.5)**
Heineken	40.0 (0.0)**	36.4 (5.6)**	30.8 (6.7)**	28.2 (5.9)**
All	40.0 (0.0)**	31.7 (2.9)**	28.4 (3.3)**	26.1 (4.1)**
Competitor-brand repricing costs				
Bass	40.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (3.9)**
Becks	40.0 (0.0)**	35.2 (9.2)**	32.5 (5.7)**	18.0 (8.7)**
Corona	40.0 (0.0)**	31.5 (8.0)**	27.7 (9.1)**	18.0 (7.6)**
Heineken	40.0 (0.0)**	30.9 (6.9)**	25.5 (6.5)**	5.7 (2.3)**
All	40.0 (0.0)**	30.4 (3.3)**	25.5 (3.9)**	18.0 (4.1)**
Own-brand repricing costs				
Bass	40.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Becks	40.0 (0.0)**	0.0 (0.0)	0.0 (0.1)	0.0 (0.0)
Corona	40.0 (0.0)**	0.0 (0.0)	0.0 (0.1)	0.0 (0.0)
Heineken	40.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
All	40.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)

Table 11: *Counterfactual experiments: median pass-through of a 5-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. Standard errors from bootstrap simulations with 400 draws reported in parentheses under each coefficient. Starred coefficients significant at the *10- or **5-percent level. Source: Authors' calculations.

Brand	Manufacturer				Retailer				Total
	Local Costs	Markup Adjustment	Costs of Repricing		Local Costs	Markup Adjustment	Costs of Repricing		
			Other	Own			Other	Own	
Bass	60.0	9.7	30.2	0.0	0.0	0.0	0.0	0.0	100.0
Beck's	60.0	2.0	2.9	35.2	0.0	0.0	0.0	0.0	100.0
Corona	60.0	8.2	0.3	31.5	0.0	0.0	0.0	0.0	100.0
Heineken	60.0	3.6	5.5	30.9	0.0	0.0	0.0	0.0	100.0
All	60.0	8.3	1.3	30.4	0.0	0.0	0.0	0.0	100.0

Table 12: *Counterfactual experiments: Decomposition of the incomplete transmission of a 5-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, other: The effect of competitors' costs of price adjustment on the manufacturer or retailer's own price adjustment behavior. Repricing costs own: Fixed costs of price adjustment incurred by the manufacturer or retailer to change its own price. Source: Authors' calculations.

A Decomposing Pass-through Elasticities

In this appendix, we describe the expressions used to decompose the pass-through elasticities into the part due to the presence of local costs and the part reflecting markup adjustment. The pass-through elasticity of the exchange-rate shock to the wholesale price after accounting for the manufacturers' nontraded costs is given by $(d \ln(ntc_j^w + tc_j^w)/d \ln tc_j^w)$ and after accounting for manufacturers' markup adjustment is given by $(d \ln p_j^w/d \ln tc_j^w)$. The pass-through elasticity to the retail price after accounting for the retailer's non-traded costs is $(d \ln(p_j^w + ntc_j^r)/d \ln tc_j^w)$; and after accounting for its markup adjustment is $(d \ln p_j^r/d \ln tc_j^w)$.

The decomposition then computes the contributions of the manufacturers' and retailer's non-traded costs and markup adjustment to the $1 - (d \ln p_j^r/d \ln tc_j^w)$ part of the original shock not passed through to the retail price. Given the Cobb-Douglas specification used to model the manufacturers' marginal costs in Section 4.1, the contribution of its non-traded costs to generating incomplete pass-through will be captured by the coefficient on domestic wages θ_{dw} , which should equal the following expression: $(1 - (d \ln(ntc_j^w + tc_j^w)/d \ln tc_j^w)) / (1 - d \ln p_j^r/d \ln tc_j^w)$. The contribution of manufacturers' markup adjustment to the incomplete retail pass-through will then be given by $(d \ln p_j^w/d \ln(ntc_j^w + tc_j^w)) / (1 - d \ln p_j^r/d \ln tc_j^w)$. Similarly, the contribution of the retailer's non-traded costs to the incomplete retail pass-through is given by $(d \ln(p_j^w + ntc_j^r)/d \ln p_j^w) / (1 - d \ln p_j^r/d \ln tc_j^w)$. And finally, the contribution of the retailer's markup adjustment to the incomplete retail pass-through is $(d \ln p_j^r/d \ln(p_j^w + ntc_j^r)) / (1 - d \ln p_j^r/d \ln tc_j^w)$.

To assess the contribution of repricing costs to the incomplete pass-through, we compare the pass-through elasticities across Simulations 1 and 2 for the strategic effect of the repricing costs, and then across Simulations 2 and 3 for the direct effect of the repricing costs. Given that in our application the value of $(1 - d \ln p_j^r/d \ln tc_j^w)$ is the same across all three simulations, we only need to calculate the difference in the wholesale price pass-through across the three simulations to quantify the effects of manufacturer repricing costs on the total incomplete pass-through. The contribution of other firms' costs of price adjustment is given by $((d \ln p_{j1}^w - d \ln p_{j2}^w)/d \ln(ntc_{j1}^w + tc_{j1}^w)) / (1 - d \ln p_{j1}^r/d \ln tc_{j1}^w)$ where the second subscript (1 or 2) denotes which simulation produced the variable. Finally, the contribution of firms' own costs of price adjustment is given by $((d \ln p_{j2}^w - d \ln p_{j3}^w)/d \ln(ntc_{j1}^w + tc_{j1}^w)) / (1 - d \ln p_{j1}^r/d \ln tc_{j1}^w)$.

B Robustness Checks

This appendix reports the results of a number of robustness checks to our baseline model. We discuss results for variants of the model in which we lengthen the horizon over which firms and/or consumers optimize their behavior, as well as alter the types of price changes used to compute the structural model's cost parameters. We find that for each of these variants of our baseline model, the results for the pass-through counterfactuals and the decomposition do not differ significantly from our baseline results.

Relaxing Assumption of Firm Myopia We report here results from a variant of our baseline model that lengthens the horizon over which firms set prices and optimize profits relative to the baseline weekly model, while consumers continue to optimize their purchases on a weekly basis. The underlying assumption in this setup is that firms have perfect foresight: They accurately predict the realized values of all variables in the model in the near future, including their own and the competitors' adjustment costs. We report results for a model where firms optimize profits over one month: We set their discount factor equal to 1 because this is weekly data. The results, reported in Tables B.1 and B.2, do not differ significantly from our baseline results. We also ran versions of this model in which firms optimize profits over periods from two weeks to six weeks, finding no significant differences from our baseline results in any of these cases. These additional results are available on request from the authors.

The results from these additional robustness checks imply that our assumption of firm myopia is reasonable in the sense that it does not impact the bottom line of our findings. Intuitively, this happens because in most periods there is little change in the main variables affecting profits. The main variation comes from exchange rates, which exhibit little variation from week to week - except for the few periods when they abruptly change. Our model's results thus appear reasonably close to the results from a finite horizon dynamic model, which we in turn expect to approximate an infinite horizon model fairly well.

Additional Robustness Checks We also ran a number of additional robustness checks whose results are available on request from the authors. We first ran a bi-weekly model, a monthly model, and a six-weekly model in which we allowed *both* firms and consumers to optimize over longer time horizons. We found that the model's demand parameters do not vary significantly when we allow consumers to optimize their purchases over a longer time horizon which supports our argument that demand dynamics are not a first-order concern in this market. We also found that the results from the counterfactuals do not differ significantly from those of the baseline model, which, together with our robustness checks reported above, is consistent with our argument that the use of static first-order conditions to approximate firms' optimal pricing is a reasonable assumption in this market.

We also ran a series of robustness checks that vary the criteria for constructing the sample of price changes, which in turn is used to compute the structural model's cost parameters. The first sample

includes promotional price changes only (sales), the second, level price changes only, and the third, all price changes including the post-promotion price changes we exclude in the baseline model, as we discuss in Section 4.1. The results from the counterfactuals using bounds derived from each of these samples do not differ significantly from those of the baseline model.

Table B.1 Counterfactual Experiments: Model Relaxes Myopia Assumption

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
No repricing costs	<i>bass</i>	40.0%	31.9%	27.4%	27.1%
	<i>becks</i>	40.0%	40.0%	36.5%	32.5%
	<i>corona</i>	40.0%	32.2%	29.5%	32.2%
	<i>heineken</i>	40.0%	40.1%	35.4%	32.3%
	<i>all</i>	40.0%	38.1%	33.9%	31.7%
Other repricing costs	<i>bass</i>	40.0%	0.0%	0.0%	0.0%
	<i>becks</i>	40.0%	39.8%	36.4%	32.4%
	<i>corona</i>	40.0%	32.0%	29.4%	18.0%
	<i>heineken</i>	40.0%	32.4%	28.2%	18.0%
	<i>all</i>	40.0%	31.7%	28.4%	18.0%
Own repricing costs	<i>bass</i>	40.0%	0.0%	0.0%	0.0%
	<i>becks</i>	40.0%	0.0%	0.0%	0.0%
	<i>corona</i>	40.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	40.0%	0.0%	0.0%	0.0%
	<i>all</i>	40.0%	0.0%	0.0%	0.0%

Counterfactual experiments: Median pass-through of a 5-percent depreciation of the relevant foreign currency. Median over 398 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.

Table B.2 Decomposition: Model Relaxes Myopia Assumption

<i>brand</i>	Manufacturer				Retailer				Total
	local costs	markup adjustment	costs of repricing other	costs of repricing own	local costs	markup adjustment	costs of repricing other	costs of repricing own	
<i>bass</i>	60.0%	8.1%	31.9%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	60.0%	0.0%	0.2%	39.8%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	60.0%	7.8%	0.2%	32.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	60.0%	0.0%	7.8%	32.4%	0.0%	0.0%	0.0%	0.0%	100.2%
<i>all</i>	60.0%	1.9%	6.4%	31.7%	0.0%	0.0%	0.0%	0.0%	100.0%

Counterfactual experiments: Decomposition of the incomplete transmission of a 5-percent appreciation of the relevant foreign currency to consumer prices. Median over 398 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, other: The effect of competitors' costs of price adjustment on the manufacturer or retailer's own price adjustment behavior. Repricing costs, own: Fixed costs of price adjustment incurred by the manufacturer or retailer to change its own price. Source: Authors' calculations.

C Demand Estimation and Habit Formation

This appendix discusses one additional way dynamics may enter the demand estimation from intertemporal effects stemming, for example, from habit formation. The implications of such dynamics are explored in Froot and Klemperer (1989), Slade (1998) and in a recent theoretical paper by Ravn, Schmitt-Grohe and Uribe (2010). While we believe that such dynamics are potentially important, it is not feasible to incorporate them into our current approach in a way that would allow us to estimate the model. In addition, the reduced-form evidence presented in Froot and Klemperer (1989) is not conclusive regarding the importance of these dynamics in explaining incomplete exchange-rate pass-through. To our knowledge, the only paper that has succeeded in estimating a model similar to ours, but with dynamics on the demand side, is 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 quite flexible estimation of the curvature of demand. We believe that the latter is particularly important in the context of the question we are trying to address, as one of the explanations for incomplete pass-through is markup adjustment by firms which is quite responsive to the estimated curvature of demand. The price we pay for this flexibility, however, is not to offer an explicit treatment of dynamics. We note however that our demand framework allows us to control for habit formation and the resulting “brand loyalty” in a reduced-form way: All our specifications include brand fixed effects, which capture, among other things, consumer loyalty for a particular brand. One indication that this reduced-form treatment of dynamics may be a reasonable shortcut is the fact that the markup estimates we obtain based on the demand system appear plausible and consistent with industry wisdom.

D Knittel-Metaxoglou Critique

In the estimation of highly nonlinear models such as random-coefficients demand systems, the objective function may exhibit many local minima. Knittel and Metaxoglou (2008) (hereafter, KM) find that in random-coefficients demand systems convergence may occur at local minima, saddle points, and in regions of the objective function where the first-order conditions are not satisfied. To check that our estimates are not stuck at a local minimum, we examine their sensitivity to: 1) the tolerance criterion used to terminate the contraction mapping process that computes the mean utility; 2) different starting values for the nonlinear optimization; 3) the type of nonlinear search algorithm used. We also check that the objective function's first- and second-order conditions for optimality are satisfied at our identified optimum to ensure it has converged to a local minimum, not a saddle point. To verify that we have identified the global minimum, we confirm that our estimates are associated with the lowest value of the objective function we observe with a zero gradient and a positive-definite Hessian. Finally, we make sure that our estimates appear reasonable from an economic-theoretic standpoint.

1. Strictness of the tolerance criterion used to terminate the contraction mapping process that computes the mean utility

In the contraction mapping process that computes the mean utility, we use a fixed tolerance level of $1e^{-9}$. This tolerance level is much stricter than that used in some random-coefficients demand models in the literature. For example, Berry et al (1995) use a tolerance of $1e^{-4}$ to terminate their contraction mapping process. As illustrated by the results in Section 6 of KM, the strictness of the criteria used to terminate the contraction mapping process can help ensure that the GMM objective function does not converge at a saddle point or other unacceptable point. We find that the objective function does not converge to the same set of estimates when we use a lower tolerance level to terminate the contraction mapping (e.g., $1e^{-4}$ - $1e^{-6}$). This may indicate that the objective function is fairly flat in the neighborhood of the global optimum. (Note that the termination tolerance on the GMM objective function value and the parameter estimates for all our estimates is fixed at 0.01.) When we use a stricter convergence tolerance for the contraction mapping, such as $1e^{-12}$, the objective function converges to the same set of estimates but only with a limited set of starting values, as there are other local minima in the neighborhood of the global minima. We settle on a fixed tolerance level of $1e^{-9}$ for the estimates reported in the paper as it appears robust to the use of different starting values but also locates the global minimum.

2. Starting values

We verify that a number of different starting values converge to the region of our identified optimum and we document the presence of other local minima as the starting value for each random coefficient varies from a large negative number to a large positive number. We check that for each of these local minima the value of the GMM objective function is higher than at our identified global minimum.

3. Robust optimization algorithm

To ensure that our estimates are not overly sensitive to our starting values, we use a very robust optimization algorithm, the Nelder-Mead non-derivative simplex search method, although it takes much longer to converge than the gradient-based quasi-Newton search method. We find our results are robust to using the quasi-Newton search method provided the starting values are not too far from the global optimum.

4. Verification of first- and second-order conditions for optimality

KM emphasize the importance of verifying first- and second-order conditions for an identified optimum. Following KM's recommendation, we confirm that our identified optimum satisfies first and second-order conditions for optimality. (Recall that if the norm of the gradient is close to zero and the Hessian is positive definite (its eigenvalues are all positive) at parameter estimates x , then the GMM objective function attains a local minimum at x . If the Hessian has both positive and negative eigenvalues at parameter estimates x , then the GMM objective function attains a saddle point at x .)

The Nelder-Mead non-derivative simplex search method we use identifies an optimum without keeping track of the gradient or Hessian along the way. To compute first- and second-order conditions for optima identified via non-derivative search methods, KM calculate numerical gradients and Hessians using the Matlab routines `fminunc.m` and `eig.m`. Although `fminunc` is an optimization routine, it provides gradients and Hessians as by-products. Like KM, we use `fminunc` to compute the norm of the gradient and the eigenvalues of the Hessian for the optimum associated with our parameter estimates.

We find the maximum of the absolute value of the gradient for our identified optimum is $5.5 * e^{-10}$, which is quite close to zero and certainly below KM's cut-off point of 30 to identify a critical point. The Hessian's eigenvalues are both positive indicating it is positive definite. They are $0.0441 * e^{-4}$ and $0.5530 * e^{-4}$. Finally, to ensure that our identified optimum is the global minimum, we confirm that our estimates are associated with the lowest value of the objective function with a zero gradient and a positive-definite Hessian. The value of the GMM objective function at the global optimum is $2.0771 * e^{-18}$.

5. Reasonable results from an economic-theoretic standpoint

Finally, one way to be confident that a global optimum has been identified is to check that the parameter values appear reasonable from an economic theoretic viewpoint. As we write in Section 4.1.1, "Table 4 reports results from estimation of the full random-coefficients logit 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 (bitterness and percentage alcohol). The coefficients on the characteristics generally appear reasonable. The random coefficient on income, at 38.45, is significant at the five-percent level, which implies that higher-income consumers are less price sensitive. Consistent with industry lore, the mean preference in

the population is amenable to a bitter taste in beer, which has a positive and significant coefficient. As the percentage of alcohol rises across brands, the mean utility in the population also rises, an intuitive result, though higher-income consumers' utility falls, evident in the negative and significant random coefficient of -50.36, and consistent with industry lore regarding the typical consumer of light beer.”

E An Alternative Approach to Computing Price Adjustment Costs

A natural question is the relationship of our “revealed preference” approach to alternative approaches to estimate adjustment costs. This appendix discusses whether and how alternative modeling approaches might be applied to compute adjustment costs in our set-up, and reports results from one such model as a robustness check to our baseline results. We find, reassuringly, that the point estimates for the price adjustment costs from this alternative model lie between the means of the upper and lower bounds from our baseline results for both the retailer and the manufacturers. Moreover, if we use these alternative estimates to calibrate the size of adjustment costs in our model, we find it produces the same results in the pass-through counterfactuals and decomposition as does our baseline model.

Adjustment costs (or switching costs as they are alternatively called) appear in several literatures in Economics (e.g., in the estimation of entry models where fixed costs of entry play the same role as adjustment costs in our framework; in the international trade literature where firms need to incur a fixed cost to enter export markets; in the labor and public finance literature where adjustment costs may prevent workers from adjusting their labor supply optimally to exogenous shocks, or from moving across firms, sectors, or regions; and in the marketing literature where adjustment/switching costs may induce consumers to stick with previously purchased products). The particular way these costs are modeled depends on the questions the researchers try to address in each case; given the large number of studies addressing different topics, we do not attempt to provide an overview of this literature. However, a general characteristic of several of these studies is that they model adjustment costs as an additive linear term in an individual’s utility function or a firm’s (reduced form) profit function. For example, in the marketing literature, this term is interpreted as a loyalty parameter, whereby a consumer’s utility is higher when she sticks with a previously purchased product. In the entry literature, this term is interpreted as the fixed costs of entry. In the trade literature, the term is interpreted as fixed costs of exporting (that are zero if the firm keeps selling for the domestic market only). In addition, the specifications include an additive error term that represents factors affecting the utility or profits that are observed by the decision maker, but not the econometrician. By making additional assumptions about the distribution of this error term (e.g., the extreme value distribution), one can estimate the parameters of the corresponding model and obtain estimates of the adjustment costs.

In our case, an analogous approach for a firm facing a binary decision whether or not to change its price would produce a specification consisting of a term representing the observable part of the profit function, an additive adjustment cost and an additive error term. Using single-product notation for our retailer for simplicity, for product j , at time t , this specification would look like:

$$\Pi_{jt,\xi}^r = (p_{jt,\xi}^r - c_{jt}^r) s_{jt,\xi}(p_{t,\xi}^r) - A_j^r + \varepsilon_{jt,\xi}, \quad \xi = \{c, nc\} \quad (1')$$

where ξ is realized as alternative c (change) if $p_{jt} \neq p_{jt-1}$ and as alternative nc (no change) if $p_{jt} = p_{jt-1}$,

and A_j^r is an adjustment cost parameter that is paid if and only if $p_{jt} \neq p_{jt-1}$. For a given product j , the value of an alternative, $\xi = \{c, nc\}$, is hence decomposed into two parts: the first, V_ξ known by the researcher up to some parameters and the second, an unknown part ε_ξ treated by the researcher as random (and which may include measurement error in V_ξ). This specification is similar to the one we use in the paper, except for the presence of the error term. Estimating this specification would enable us to obtain point estimates, and not just upper and lower bounds, of the adjustment costs.

Note however, that there is an important difference between the specification above and those typically used in the adjustment cost literature. In the latter, the observable part of the profit function, the V_ξ , is typically modelled in reduced form (for example in the entry literature, the V_ξ represents reduced form profits, which are a function of a market's demographics). In contrast, in our framework, the V_ξ is modelled structurally. This has two implications. The first one is that in the structural modeling of these profits, we are explicit at each step of the analysis about the information structure and the interpretation of the error terms in the analysis; against this background, adding the error term in the end, seems arbitrary and hard to justify (in other words, it is not clear what it is that the researcher does not observe about the profits given that we have already made specific assumptions about what is observed about demand and costs in the earlier part of the analysis). More importantly, as the notation above makes clear, the structural modeling of the profits implies that the variables entering the first part of the profit function are themselves a function of the adjustment costs. In order to estimate the parameters in the above specification (i.e., the marginal cost and adjustment cost parameters) we would need to fully solve the model to figure out what the prices and market shares would be, had the firm's price adjustment followed a different scenario, and we need to do this for all possible counterfactual scenarios (note that the counterfactual prices and shares in each case are functions of all model parameters and adjustment costs, which however are not known at this stage). The presence of strategic interactions among manufacturers makes this a particularly daunting task. It is not clear how one would implement the estimation in this case, or if this would be even feasible. Even if it were feasible, the computational cost would be considerable. Note that this problem does not arise in studies of consumer choice where the observable part of the utility function does not depend on strategic interactions and the market equilibrium, and where the value of counterfactual scenarios (e.g., the choice of an alternative product) can easily be computed – at least as long as the choice set is not exceedingly large.

The approach we currently adopt allows us to separate the marginal cost from the adjustment cost estimation; we use standard estimation techniques to back out marginal costs using data from periods in which prices adjust (without any information on adjustment costs), and then use data from all periods in the sample to derive bounds on the adjustment costs. The advantage of this approach is that we do not need to fully solve the model to obtain estimates of the marginal costs and bounds - by relying on revealed preference, we exploit the idea that the observed scenario must be more profitable for the firm

than some alternative scenario, without having to specify the full set of these alternative scenarios. Of course, the price we pay is that we can derive bounds only, and not point estimates.

To further explore the possibility of adapting the approach commonly used in the switching cost literature to our problem, we conducted the following robustness check. As noted above, estimating marginal and adjustment costs using the specification discussed above is exceedingly hard, if at all feasible. However, we can take the marginal cost estimates we obtained using the standard procedure, use those to derive counterfactual prices and market shares for a subset of possible scenarios (as we currently do in the paper), and then plug those counterfactual prices and market shares into specification (1'). Given the interpretation of our bounds as the additional amount of profit available from changing a given product's price each period, this is equivalent to collecting the residuals from each of the inequalities considered in the paper's Section 3.1, both those associated with the upper and the lower bounds, collating them by product across periods with and without price adjustment, and then using the resulting dataset (with one observation for each product and market) to estimate a simple binary choice model of whether or not a firm adjusts its price in each period.

Note if one adopts the conventions of the most commonly used model in this literature, the logit model, and assumes that the unobserved component of net profit is distributed i.i.d. extreme value for each alternative, c and nc , then the parameter of interest, A_j^r , would be identified by the difference in profits between the two options for each observation – what we'll call relative profits, the observed part of which is the bound estimated for that observation, $V_j^* = V_{j,c} - V_{j,nc}$, and the relative error term $\varepsilon_j^* = \varepsilon_{j,c} - \varepsilon_{j,nc}$, which is distributed logistic.

As a robustness check, we estimate one such discrete choice model, a binary logit model, with results that are similar to, and consistent with, our estimated upper and lower bounds on the adjustment costs, for both the retailer and the manufacturers. To do so, we assume a single fixed adjustment cost parameter across all the products and time periods in our sample, to ensure sufficient observations to identify such a highly nonlinear model. We also include brand fixed effects in the estimation, to control for brand-specific pricing idiosyncrasies, although the results are robust to simply pooling the data as well.

We find that at the level of relative profits at which the marginal effects rise rapidly, a common measure of this type of threshold effect, an increase in the relative profits as a share of revenue of roughly 0.74 percent is associated with a price change by each manufacturer, as reported in Table E.1. This coefficient, which is statistically significant at the 1-percent level, lies between the mean manufacturer lower and upper bounds as a share of revenue of 0.64 and 1.28 percent, respectively, as reported in the paper's Table 9. Similarly an increase in relative profits as a share of revenue of roughly 1.33 percent is associated with a price change by the retailer. This coefficient, which is also statistically significant at the 1-percent level, lies between the mean retailer lower and upper bounds as a share of revenue of 0.84 and 3.16, respectively. Finally, when we use these estimates to calibrate the size of

adjustment costs in our model’s counterfactuals, we get results for the pass-through counterfactuals and the decomposition that are nearly identical to those of our baseline model. Note that the table reports results for the brands used in our counterfactuals, *Bass*, *Becks*, *Corona*, and *Heineken*, as these are the relevant brands for the counterfactuals. The results are also robust to including the other foreign brands in the sample.

Variable	Retailer	Manufacturer
Relative Profit as a Share of Revenue	1.33 (0.10)**	.74 (0.10)**
Pseudo R^2	0.08	0.11
Observations	1616	1616

Table E.1: *Adjustment costs as a share of revenue from structural choice model.* Coefficients represent the threshold level of relative profits as a share of revenue required to rapidly increase the probability of a price change. The model includes brand fixed effects. Includes observations for Bass, Becks, Corona, and Heineken. Source: Authors’ calculations.

Although this simple logit model produces results similar to our baseline model, it is difficult to justify the assumptions that underlie it, in particular the source and properties of the additive error term and the fact that in order to implement the logit we consider only a subset of all possible counterfactual scenarios (which is not theoretically consistent). In our view, it would require a completely different set-up to use such a framework appropriately to derive firms’ adjustment costs.