

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 or the Federal Reserve System.

© 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 2010
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 58% of the incomplete exchange rate pass-through is due to local non-traded costs; 9% to markup adjustment; 32% 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
Professor of Economics
Princeton University
Department of Economics
306 Fisher Hall
Princeton, NJ 08544-1021
and NBER
pennykg@princeton.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.¹

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

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

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

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

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

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

The current paper attempts to overcome this shortcoming by explicitly introducing price rigidities into the model and suggesting an approach for quantifying their importance in explaining the docu-

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

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

The basic idea behind our approach is as follows. First, even with nominal price rigidities, we can estimate the demand and cost parameters of the model along the lines followed by earlier papers by constraining the estimation to the periods for which we observe price adjustment; the underlying premise is that once a firm decides to incur the adjustment cost associated with a price change, it will set the product’s price according to the first-order conditions of its profit maximization problem. This of course does not imply that this firm’s behavior will be unaffected by the existence of price rigidities. Such rigidities may still have an indirect effect on the pricing behavior of firms that adjust their prices, as in any model of oligopolistic interaction firms take the prices (or quantities) of their competitors into account; if the competitor prices do not change in 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 the periods in which prices adjust and periods in which prices remain unchanged to derive bounds on the adjustment costs associated with a price change. The derivation of the bounds is based on a “revealed-preference-approach”: In particular, our method is based on the insight that in periods in which prices change, it has to be the case that the costs of price adjustment are lower than the additional profit the firm makes by changing its price; we can use this insight to derive an upper bound of this price adjustment cost. Similarly, in periods in which prices do not change, it has to be the case that the costs of adjustment exceed the extra profit associated with a price change; based on this insight, we can derive a lower bound for the price adjustment cost.

The costs of price adjustment are a concept that has a precise meaning *within the context of our model*; they are defined in the broadest possible sense as everything that prevents a firm from adjusting its price in a particular period. As such, they are not directly comparable to estimates obtained in earlier studies using different methods (e.g., direct measurement or firm surveys).² More importantly, the adjustments costs alone do not allow a full assessment of the impact of nominal price rigidities on exchange-rate pass-through; because such rigidities have both a direct and an indirect (operating through the competitor prices) effect on firms' pricing behavior, it is possible that very small rigidities induce significant price inertia. To provide an overall assessment of the impact of price adjustment costs we therefore perform simulations that compare firms' pricing behavior with price rigidities to their behavior given fully flexible prices. The differential response of prices across the 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 related to exchange-rate pass-through and price rigidities for several reasons: (1) a significant fraction of brands are imported and hence affected by exchange-rate fluctuations; (2) long-run exchange-rate pass-through onto consumer prices is low, on the order of 5-10 percent; (3) there exist highly disaggregate, weekly data on both wholesale and retail prices; this allows us to examine how prices respond at each stage of the distribution chain; (4) The patterns of prices for beer are typical for the type of goods included in the CPI, as described in the recent work of Bils and Klenow (2004), Klenow and Kryvtsov (2008), and Nakamura and Steinsson (2008); 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 more generally applied to any market for which high-frequency data are available so that the points of price adjustment can be identified.

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

²There have been several attempts in the literature to measure price adjustment costs. Levy et al (1997) find menu costs to equal 0.70 percent of supermarkets' revenue from time-use data. Dutta et al (1999) find menu costs to equal 0.59 percent of drugstores' revenue. Levy et al have four measures of menu costs: 1. the labor cost to change prices; 2. the costs to print and deliver new price tags; 3. the costs of mistakes; 4. the costs of in-store supervision of the price changes. Some detailed microeconomic studies have cast doubt on the importance of menu cost in price rigidity. Klenow and Willis (2006) estimate price adjustment costs to be 1.4 percent of total firm revenue. Blinder et al (1998) find in a direct survey that managers do not regard menu costs as an important cause of price rigidity. Both Carlton (1986) and Midrigan (2010) find that firms change prices frequently and in small increments, which is not consistent with a menu-cost explanation of price rigidity.

inventories of beer, the demand and supply side parameter estimates obtained by the static approach may be biased. Specifically, on the demand side, Hendel and Nevo (2006a, 2006b) have shown that when consumers stockpile in response to temporary price reductions (sales), static demand estimates may overstate the long-run price elasticities of demand by a factor of 2 to 6³. On the supply side, Aguirregabiria (1999) has analyzed the pricing behavior of a monopolistically-competitive retailer who holds inventories in a central store, and delivers goods from this store to individual outlets. He shows that in the presence of fixed ordering costs and nominal price rigidities inventory dynamics have a significant effect on the retailer’s decision to change a brand’s price⁴; ignoring such dynamics may hence lead one to biased estimates of the importance of nominal price rigidities. Fortunately, these concerns that both build on the importance of inventories appear to be less relevant in our case. The industry wisdom is that consumers typically consume beer within a few hours after its purchase, so that consumer stockpiling is not a first-order concern.⁵ On the supply side, state and local regulations concerning the distribution of all alcohol, including beer, in Illinois stipulate that it is illegal for the central store of a retail chain to maintain inventories of beer and to deliver them to individual outlets.⁶ 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. As the central store does not keep inventories of beer (indeed cannot by law), there is no relationship between inventory decisions and prices.⁷ And there is no incentive for individual outlets to maintain inventories, as they can get a new shipment 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 explicitly model the fact that with ongoing uncertainty and rational expectations there is 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 result in estimates of adjustment costs that are biased upwards. In this sense our approach is most similar to the static models considered in Akerlof and Yellen (1985) and Mankiw (1985). These models are based on the assumption that either agents have static expectations, or shocks are perceived as permanent. While this is certainly a strong (and

³These numbers refer to laundry detergents.

⁴A similar point is made in Alessandria, Kaboski and Midrigan (2010).

⁵See Anheuser-Busch, Inc. “Beer Shopper Poll,” 2005.

⁶For more on Illinois liquor regulations, see the homepage of the Illinois Liquor Control Commission at <http://www.state.il.us/LCC/>.

⁷As 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).

in many settings unreasonable) assumption, we note that in our setting the main (in fact in the formal model, *only*) source of uncertainty is *exchange rates*, which are highly persistent⁸. Therefore, the assumption that shocks are perceived as permanent seems more reasonable in the context of exchange rates than in the case of other cost shocks, which may exhibit less persistence. Unfortunately, characterizing the 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 should emphasize however that within the static framework we interpret the derived “adjustment costs” in the broadest possible sense as a concept that includes everything that prevents a firm from adjusting its price in a particular period, including the option value of the status 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 to be primarily driven by the infrequent adjustment of wholesale rather than retail prices. In our data, there is not a single instance where a product’s retail price remains unchanged in response to a wholesale price change. Hence it seems that the primary reason that retail prices do not change from period to period is that there is little reason for them to change, as the underlying wholesale prices remain fixed.

As we discussed above, nominal price rigidities may affect the pricing decisions of a particular producer in two ways. First, they may prevent this producer from adjusting her price, because her *own* costs of repricing exceed the benefits, even when all other competing producers adjust their prices (direct effect). Second, such costs may induce *other* competing producers to keep their prices fixed, which may make price adjustment less profitable for the producer under consideration (indirect/strategic effect). Our simulations indicate that the direct effect is significant at the wholesale level, accounting for over 30 percent on average for the incomplete pass-through. Interestingly, there is substantial variation in this estimate across brands; the *own* costs of price adjustment appear to be more important for brands with 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

⁸The consensus among International economists is that the exchange rate follows a random walk.

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 58.0 percent of the incomplete pass-through to local non-traded costs, 9.2 percent to markup adjustment, 31.4 percent to the existence of own-brand price adjustment costs, and 1.4 percent which represent the indirect/strategic effect of such costs. As suspected at the beginning, costs of price adjustment appear to be substantially more important at the wholesale than retail level.

It is important to distinguish between those features of the approach we propose that are general, i.e., can be adopted easily by other researchers, and those specific to our application. The general features of our approach 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 static expectations (or, equivalently, shocks that are perceived as permanent) 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 what features of the data drive our results and which of these features are likely to 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 generated through 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 through 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 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 explicitly 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 through the existence of nominal rigidities. Perhaps the most interesting aspect of our results is that we still find that local non-traded costs are the primary source of incomplete exchange rate pass-through, accounting on average for circa 58 percent of the incomplete response – this, despite the fact that we focus on a market with infrequent price adjustment and that we explicitly 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 providing a brief description of the market and the data in the next section; in the same section, we also provide some descriptive statistics and discuss the price adjustment patterns evident in the retail and wholesale price data. Section 3 discusses the model, 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 The Market and the Data

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

2.1 Market

Beer is an example of one type of imported goods: packaged goods imported for consumption. Such imports do not require any further manufacture before reaching consumers and make up roughly half of the non-oil goods imports to the U.S. over the sample period. The beer market is well suited for an exploration of the sources of local-currency price stability for the reasons discussed in the Introduction: a significant fraction of brands are imported; exchange-rate pass-through to prices is generally low (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 fixed adjustment costs in firms' incomplete transmission of exchange-rate shocks to prices.

2.2 Data

Our data come from *Dominick's Finer Foods*, the second-largest supermarket chain in the Chicago metropolitan area in the mid 1990s with a market share of roughly 20 percent. The data record the retail and wholesale prices for each product sold by *Dominick's* over a period of four years. They were gathered by the *Kilts Center for Marketing* at the University of Chicago's Graduate School of Business and include aggregate retail volume market shares and retail prices for every major brand of beer sold in the U.S.⁹ 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.¹⁰ 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.¹¹ Supermarkets sell approximately 20 percent of all beer consumed in the U.S.¹²

We aggregate data from each *Dominick's* store into one of two price zones. For more details about this procedure, see Hellerstein (2008).¹³ We define a product as one six-pack serving of a brand of beer and quantity as the total number of servings sold per week. We define a market as

⁹The data can be found at <http://gsbwww.uchicago.edu/kilts/research/db/dominicks/>.

¹⁰Features of the *Dominick's* wholesale-price data confirm that brewers control distributors prices to the supermarket. Across individual *Dominick's* stores, which may each be served by a different distributor, each with an exclusive territory, the variation in UPC-level wholesale prices is less than one cent. Asker (2004) notes that one cannot distinguish distributors by observing the wholesale prices they charge to individual *Dominick's* stores. This supports the industry lore that distributors pricing is coordinated by brewers and is not set separately by each distributor to each retail outlet.

¹¹Canadian Trade Commissioner (1998).

¹²As our data focus on one metropolitan statistical area, we do not need to control for variation in retail alcohol sales regulations. Such regulations can differ considerably across states.

¹³The zones are defined by *Dominick's* mainly on the basis of customer demographics. Although they do not report these zones, we identify them through zip-code level demographics (with a few exceptions, each *Dominick's* store in our sample is the only store located in its zip code) and by comparing the average prices charged for the same product across stores. We classify each store according to its pricing behavior as a low- or high-price store and then aggregate sales across the stores in each pricing zone. This aggregation procedure retains some cross-sectional variation in the data which is helpful for the demand estimation. Residents' income covaries positively with retail prices across the two zones.

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

2.3 Preliminary Descriptive Results

We begin the analysis by documenting in several simple regressions whether *Dominick's* imported-beer prices are systematically related to movements in bilateral nominal exchange-rates. These results can provide a benchmark against which we can measure the performance of the structural 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 price¹⁴; p^w is the product's wholesale price; c_j , θ_t , and ζ_z are product, week and zone dummies respectively that proxy among other things for demand shocks that may affect a brand's price independent of exchange rates; e is the bilateral nominal exchange rate (domestic-currency units per unit of foreign currency); co_{jt} denotes a set of variables that proxy for cost shocks that again may affect prices; such variables include measures of domestic (U.S.) wages, the price of barley in each country producing beer in our sample, the price of electricity in the Chicago area, and - for foreign brands - wages in each beer exporting country in our sample; ε is a random error term. All variables are specified in levels, and not first differences, as our focus is on the long-run pass-through of exchange

¹⁴Imported beer is generally invoiced and paid for in U.S. dollars in this market. We do not believe that this practice in itself "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 prices to be constant, which our analysis addresses.

rate changes, and not the short-term dynamics.¹⁵

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 (since the latter do not vary at the weekly level). The results across the two specifications are remarkably similar. The average pass-through elasticity α for the retail price is - based on column 2 - 6.7 percent and is significant at the one-percent level. The regression establishes a roughly 7-percent benchmark for the retailer's pass-through elasticity, that we will try to explain within the framework of the structural model. The fourth column of Table 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.

2.4 Patterns of Price Adjustment in the Data

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

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

Second, a substantial fraction of the price variation reflects temporary price reductions (sales). As

¹⁵It 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.

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

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

3 Model

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

3.1 Supply

We model the supply side of the market using a linear-pricing model in which manufacturers, acting as Bertrand oligopolists with differentiated products, set their prices followed by retailers who set their prices taking the wholesale prices they observe as given. Thus, a double margin is added to the marginal cost of the product before it reaches the consumer.¹⁶ 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 manufacturer and retailer is as follows. First, the manufacturer decides whether or not to change the product's price taking into account the current period's observables (costs, demand conditions, and competitor prices), and the anticipated reaction of the retailer. If she decides to change the price, then the new price is determined based on the manufacturer's first-order conditions. Otherwise the wholesale price is the same as in the previous period. Next, the

¹⁶In modeling the industry's vertical relationships we draw on previous work that used both industry lore and formal non-nested tests to choose the model that best fits the data (Hellerstein, 2004). The non-nested tests, developed by Villas-Boas (2007) compare observed and derived markups under different assumptions about the industry's vertical relationships. The industry features that generate the observed markups thus determine the choice of model for the industry's vertical relationships, which in turn determines the pass-through results from the counterfactuals.

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

3.1.1 Retailer

Consider a retail firm that sells all of the market's J differentiated products. Let all firms use linear pricing and face constant marginal costs. The profits of the retail firm 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-traded¹⁷ 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:

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

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 this cost is independent of the sales volume; it is a discrete cost that the retailer pays every time the price adjusts from the previous period. The indexing of A by product j and time t in our notation corresponds to the most flexible specification, in which the price adjustment cost is allowed to vary by product and time. One could potentially impose 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:

¹⁷We use the term "non-traded" to indicate that these costs are paid in dollars regardless of the origin of the product. Hence, non-traded costs will not be affected by exchange rate shocks.

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

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

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

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 (and, in theory, a $J \times J$ 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 equal to zero otherwise. In our case T^r is all ones so we omit it for simplicity going forward). The stacked first-order conditions can be rewritten in vector notation:

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

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

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

where the retail price for product j in market t will be the sum of its wholesale price, non-traded costs, and markup. The presence of the adjustment costs A_{jt}^r in the profit function implies that for the retailer to change her price in the current period, it will have to be the case that the extra profits associated with the new price are at least as large as the adjustment cost:

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

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

adjustment costs A_{jt}^r). In this sense, it simply captures a “revealed-preference” argument. Note that the inequality we consider above is only one possibility from a large set of inequalities that can be potentially used to infer adjustment costs. This is because 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* inequalities to infer the bounds - in fact, this would be infeasible. Using a subset of them - in this case a subset containing a single inequality - is sufficient for obtaining consistent estimates of the bounds, given that the revealed preference approach we described above has to hold for every single permutation of price changes¹⁸. By rearranging terms we can use the above inequality to derive an *upper* bound \overline{A}_{jt}^r for the price adjustment costs of product j :

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

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

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

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

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

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

¹⁸The same applies to the derivation of lower bounds and the derivation of bounds for the manufacturer.

The essence of our empirical approach to quantify the adjustment costs can be described as follows. First, we estimate the demand function. Once the demand parameters have been estimated, the market share function $s_{jt}(p_t^r)$ as well as the own and cross price derivatives $\frac{\partial s_{jt}}{\partial p_{jt}^r}$ and $\frac{\partial s_{kt}}{\partial p_{jt}^r}$ can be treated as known. Next we exploit the first-order conditions for each product j (6) to estimate the non-traded costs and markups of product j , but contrary to the approach typically employed in the Industrial Organization literature, we use *only the periods in which the price of product j adjusts*, to back out costs and markups. In periods when the price does not adjust, the non-traded costs are not identified based on the first-order conditions; however, we can derive estimates of the non-traded costs for these periods by imposing some additional structure on the problem, e.g., by modeling non-traded costs parametrically as a function of observables along the lines described in the next section. Once estimates of non-traded costs for these periods have been derived, we can calculate the counterfactual price p_{jt}^{rc} that the retailer would have charged if there were no price rigidities and she behaved according to the profit maximization conditions, as well as the associated counterfactual market share $s_{jt}^c(p_{jt}^{rc}, p_{kt}^r)$. In the final step, we can exploit inequalities (10) and (12) to derive upper and lower bounds of the adjustment costs A_{jt}^r .

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

3.1.2 Manufacturers

Let there be M manufacturers that each produce some subset of the market's J_t differentiated products.¹⁹ Each manufacturer chooses its wholesale price p_{jt}^w taking the retailer's anticipated behavior

¹⁹Two manufacturers are modelled as multiproduct firms, *Heineken N.V.* which owned *Amstel* and *Heineken* over the sample period, and *Guinness PLC*, which owned *Guinness* and *Harp*.

into account. Manufacturer w 's profit function is:

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

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

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

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

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

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

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

$$s_{jt} + \sum_{k=1,2,\dots,J} (p_{kt}^w - c_{kt}^w) \frac{\partial s_{kt}}{\partial p_{kt}^w} = 0. \quad (15)$$

This gives us another set of J equations, one for each product. Let Ω_{wt} be the manufacturer's reaction matrix with elements $\frac{\partial s_{jt}(p_t^r(p_t^w))}{\partial p_{jt}^w}$, the change in each product's share with respect to a change in each product's wholesale price. The manufacturer's reaction matrix is a transformation of the retailer's reaction matrix: $\Omega_{wt} = \Omega'_{pt} \Omega_{rt}$ where Ω_{pt} is a J -by- J matrix of the partial derivative of each retail

price with respect to each product's wholesale price. Each column of Ω_{pt} contains the entries of a response matrix computed without observing the retailer's marginal costs. The properties of this manufacturer response matrix are described in greater detail in Villas-Boas (2007).²⁰

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

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

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

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

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

This condition is similar to inequality (10) for the retailer, with a slight difference: the counterfactual market share s_{jt}^c that the manufacturer would face if she left the price of product j unchanged is a function of the counterfactual retail price p_{jt}^{rc} that the retailer would charge when faced with an unchanged wholesale price p_{jt-1}^w . But given the existence of the retail adjustment cost, this counterfactual price can follow one of two scenarios: the first one is that the retailer does not change the price from the previous period, so that $p_{jt}^{rc} = p_{jt-1}^r$; the second possibility is that the retailer adjusts her price according to the retailer's first-order conditions (6). Hence, before one can use the above inequality to infer the upper bound of the manufacturer's adjustment cost, it is necessary to solve the retailer's problem to determine the retailer's price response. Specifically, using notation for a single-product retailer for simplicity, if:

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

²⁰To 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 .

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

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

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

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

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

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

The lack of price adjustment in this case implies that the wholesale price is not necessarily determined based on the manufacturer first-order condition. Regarding the retail price, it is again possible that the retailer adjusts the retail price in periods when the wholesale price remains unchanged. However, in practice we rarely observe this case in the data. Hence, we concentrate on the case where both wholesale and retail prices remain unchanged, that is $p_{jt}^w = p_{jt-1}^w$ and $p_{jt}^r = p_{jt-1}^r$.

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

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

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

adjustment cost A_{jt}^w :

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

The determination of the counterfactual optimal wholesale price p_{jt}^{wc} and the associated counterfactual market share s_{jt}^c is however more involved in this case, as the manufacturer has to take into account the reaction of the retailer, who may or may not adjust her price in response to a wholesale price change. To find the price p_{jt}^{wc} the manufacturer would set if she were willing to incur the adjustment cost, we proceed as follows. First, we consider the case in which the retail price would have changed in response to the wholesale price change. In this case p_{jt}^{wc} would be determined according to equation (16) which reflects the manufacturer's first-order condition; the inverted manufacturer reaction matrix Ω_{wt}^{-1} in this equation incorporates the optimal pass-through of the wholesale price change onto the retail price.

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

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

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

The final step in determining the counterfactual optimal wholesale price p_{jt}^{wc} that the manufacturer would choose if she changed the wholesale price from the previous period is to compare the manufacturer profits for the case where the retailer adjusts the price, to the manufacturer profits for the case where the retailer does not pass-through the wholesale price change, in which case the wholesale price will be set according to (21). The manufacturer will pick the p_{jt}^{wc} that corresponds to the higher

profits. Once the wholesale price is found, the optimal retail price response and associated market share can be determined as well, and inserted in (20) in order to infer the manufacturer adjustment cost lower bound.

3.2 Implied Pass-through

To assess the overall impact of these adjustment costs on firms' pricing behavior we employ simulations²¹. We first compute the industry equilibrium that would emerge if a firm faced an exchange-rate shock and prices were fully flexible, that is, all adjustment costs were equal to zero. In a second set of simulations, we derive the industry equilibrium under the presence of nominal rigidities. We interpret the differential response of prices across the different cases as a measure of the impact of nominal price rigidities.

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 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:

$$p_{jt}^{w*} = c_{jt}^w(tc_{jt}^w, ntc_{jt}^w) - \sum_{k \in \Gamma_{mt}} (S_{wt} * T_w)^{-1} s_{kt} \text{ for } j = 1, 2, \dots, J_t. \quad (22)$$

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:

$$p_{jt}^{r*} = p_{jt}^{w*} + ntc_{jt}^r - \sum_{k \in \kappa^r} (S_{wt} * T_w)^{-1} s_{kt} \text{ for } j, k = 1, 2, \dots, J_t. \quad (23)$$

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

²¹As 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.

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 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, as we describe in Appendix A.

3.3 Demand

The estimation of costs, markups, and adjustment costs requires consistent estimates of the demand function as a first step. Market demand is derived from a standard discrete-choice model of consumer behavior. Given that the credibility of all our results will ultimately depend on the credibility of the demand system, it is imperative to adopt as general and flexible 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. In the following we provide a brief overview of the model, and refer the reader to Nevo (2001) and Hellerstein (2008) for a more detailed discussion of the implementation.

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

$$u_{ijt} = x_{jt}\beta_i - \alpha_i p_{jt} + \xi_{jt} + \varepsilon_{ijt} = V_{ijt} + \varepsilon_{ijt}, \quad i = 1, \dots, I, \quad j = 1, \dots, J, \quad t = 1, \dots, T. \quad (24)$$

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

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i + \Sigma v_i \quad (25)$$

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

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt} \quad (26)$$

Finally, consumers have the option of purchasing an “outside” good; that is, consumer i can choose not to purchase any of the products in the sample (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.²² 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:

$$s_{jt} = \int_{\zeta \in A_j} P^*(d\zeta) \quad (27)$$

where $P^*(d\zeta)$ is the density of consumer characteristics $\zeta = [D \ \nu]$ in the population. To compute this integral, one must make assumptions about the distribution of 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 is permitted under the restrictions of the multinomial logit model. The market share function is computed as the integral over the taste terms μ_{ijt} of the multinomial logit expression:

$$s_{jt} = \int_{\mu_{it}} \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} f(\mu_{it}) d\mu_{it} \quad (28)$$

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

$$s_{jt} = \frac{1}{N} \sum_{i=1}^N \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} \quad (29)$$

²²The existence of an “outside” good means that the focus on a single retailer (*Dominick’s*) does not imply that this retailer has monopoly power in the retail market. The effect of local conditions is accounted for by the presence of this outside good. When computing the price elasticities for each brand, we thus 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 included in our sample (such as micro-brews) or purchase beers in another supermarket following a rise in *Dominick’s* price for a foreign brand, our model will produce consistent estimates of pass-through elasticities. One limitation of this method of deriving the market’s aggregate demand curve, however, is that one must assume that the price of the outside good remains constant, which does not allow for strategic interactions between the retailer in our data 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 all consumers, using a nonlinear generalized method-of-moments (GMM) 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 now discuss several key assumptions of our model that simplify both the consumer’s and the firm’s problems by abstracting from various dynamic considerations. What are the possible costs of this simplicity?

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 be overstating the long-run price elasticities of demand (see Hendel and Nevo, 2006a and 2006b). As we noted earlier in the Introduction, the industry wisdom is that the typical buyer consumes beer within hours after the purchase, so that stockpiling is not a first-order concern in this particular market. To further investigate whether there is a systematic pattern in sales 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 that had elapsed since a previous sale, holiday dummies, or sales for other brands, could predict a current sale for a particular brand. This apparent randomness of sales in our sample is consistent with a Varian-type explanation of temporary price reductions. A potential problem with such an explanation is that in Varian’s model firms randomize prices each period, so that the notion of a “regular price” does not exist. However, in the Varian (1980) model there are no costs of price adjustment. Introducing costs of price adjustment (as we do in this paper) can explain the existence of “regular prices” that is evident in Figures 1 and 2 (see Chung, 2009). We should point out that we by no means claim that the patterns we observe in our data for beer generalize to all other product categories. In fact, there is substantial evidence, based on Pesendorfer (2002) and recent papers by 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.

Demand: Habit Formation A second way in which dynamics may enter the demand estimation is through 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 retailers and manufacturers which is quite responsive to the estimated curvature of demand. The price we pay for this flexibility, however, is that we do not 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 specifications include brand fixed effects, which capture – among other things – the loyalty that consumers may have developed towards 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.

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 noted 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: Static Expectations Perhaps the strongest assumption in our setup is that firms maximize profits myopically, period by period. Even with static demand and supply, this static (myopic) optimization may seem unreasonable in the presence of fixed costs of repricing. This is

because these costs imply that with ongoing uncertainty, there is an option value in not adjusting prices: Rational firms realize that they will be stuck with the same prices for a while, so there is value in waiting to set a price that is optimal over a longer horizon. As shown by Dixit (1991), the implication of rational expectations is that even very small costs of adjustment can generate significant price inertia.

The static model we bring to the data is based on the premise that firms have static expectations, or, equivalently, that they assume that the shocks they face in the current period (week) are permanent. In this sense, our model is similar to the static models considered by Akerlof and Yellen (1985) and Mankiw (1985). While this assumption may seem controversial - in fact unreasonable - in a general setting, we believe that it is rather inconsequential in the particular context of *exchange rate* pass-through. The reason is the following.

The main source of uncertainty firms face in our sample is exchange rates. In fact, in the formal model, exchange rates are the only source of uncertainty, as demand and adjustment costs are assumed to be deterministic and perfectly observable to the firms (but not the econometrician). However, exchange rates are highly persistent; the consensus in the International literature is that exchange rates follow a random walk. Given this persistence, it is not unreasonable to assume that firms view exchange rate shocks as permanent. In this case, the firm's dynamic optimization problem becomes equivalent to a static one (up to a discount factor)²³, as current values of the cost shocks are assumed to persist in the future. Of course, this assumption would not be attainable in the case of other shocks that exhibit less persistence.

In sum, we justify our static optimization framework on two grounds: (1) the static model is equivalent to a dynamic model with myopic expectations, that is to a model in which firms perceive current shocks as being permanent; and (2) myopic expectations are a reasonable approximation to reality in the particular context of exchange rate shocks, because the latter are highly persistent. Of course, in addition to exchange rates, there are also other sources of cost variation (energy prices, the price of barley, etc.), but the cost variation induced by these factors pales in comparison to that generated by exchange rates. In Figure 3, we plot the exchange rate facing beer producers over our sample period along with the prices for glass and energy (all series are normalized to start at 1 at the beginning of our sample). While there is some variation in the prices of energy and glass, the exchange rate variation trumps everything else.²⁴

²³We refer the reader to Appendix B for a more detailed discussion of this issue.

²⁴One of the few papers that provide a dynamic treatment of price rigidities is Nakamura and Zerom (2010). They adopt a dynamic framework to estimate menu costs, which are modeled as random i.i.d. variables. While the menu costs are derived within a dynamic framework, computational constraints force them to estimate the cost parameters based on static profit maximization conditions that do not take menu costs into account, so that there is a tension between the model used to obtain menu costs and conduct simulations, and the static model that is used to obtain the cost estimates. Our approach has the advantage that it enables us to maintain a theoretically consistent approach throughout the analysis. Despite these differences, it is interesting to note that the results of the two papers are similar in several

While the assumption of static (myopic) optimization is not directly testable, we have tried to get a sense of its implications for our results by estimating three alternative versions of our structural model, each with a longer horizon over which consumers make purchases and firms set prices than in the paper’s weekly model, and by examining the robustness of our results to this extension of the relevant time horizon. In particular, we have estimated a bi-weekly, monthly, and six-weekly model. The results from these alternative models (reported in Appendix C) are very similar to the baseline specification with a weekly model.

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 make 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 (such as 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 subsequent shocks that firms experience following a price change were not forecastable at the time the price was set and therefore did not affect the choice of the price.

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 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 through 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

dimensions. Interestingly, the dynamic approach of Nakamura and Zerom (2010) provides indirect support for the static approach adopted here, as they note that the high persistence of the cost shocks they consider (they consider commodity price shocks, which follow a random walk just like exchange rates) implies that menu costs have a sizeable impact on pricing behavior in the short-run, but a negligible effect in the long-run. This is because in the presence of unit roots, firms assume that the shock will last forever, and hence find it optimal to respond within a short period of time, by either adjusting or not adjusting prices, depending on the size of the shock. This is exactly the premise underlying our analysis.

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 (29) 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 follows²⁵:

$$\delta_{jt} = \beta d_j - \alpha p_{jt} + \Delta \xi_{jt} \quad (30)$$

where the product fixed effects d_j proxy for both the observed characteristics x_{jt} as defined in equation (24) 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 are orthogonal to the error term $\Delta \xi_{jt}$. Factor prices and exchange rates satisfy this condition as they are unlikely to have any relationship to promotional activities while they are by virtue of the supply relation correlated with product prices. To construct our instruments we interact hourly wages in each country’s beverage industry with weekly bilateral exchange rates and indicator variables for each brand; this allows each product’s price to respond differently to a given supply shock.

Table 3 reports the results from the demand estimation using the multinomial logit model. Following the specification in equation (30), 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. Note that the 95-percent confidence interval of the latter coefficient does not include the value of the former. Columns 2 and 4 report results when a dummy for major holidays is included as

²⁵The demand model is also indexed by price zone z . In each period we have observations for two separate price zones. To keep the exposition simple, we omit the subscript z from our notation.

an additional right-hand-side variable, as one might expect the coefficient on price to differ in periods of high demand for beer, such as during major holidays. They indicate that including the holiday dummy as a right-hand-side variable does not affect the demand coefficients in either the *OLS* or the *IV* estimation. The coefficient on the holiday dummy is significant at the 1-percent level in the *OLS* specification reported in column 2, but not in the *IV* specification reported in column 4. Finally, the first-stage part of Table 3 reveals 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. This suggests that factor costs interacted with exchange rates are valid instruments for the demand estimation.

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. As consumers' income rises, they become less price sensitive. The random coefficient on income, at 40.37, is significant at the five-percent level. The mean preference in the population is not amenable to a bitter taste in beer, which has a negative and significant coefficient. As the percent alcohol rises across brands, the mean utility in the population also rises, an intuitive result. There is heterogeneity in the population with respect to this characteristic: Those with higher incomes get less utility from a high percent of alcohol in their beer, that is, prefer light beers, as one can glean from the negative and significant random coefficient of -51.84. This is consistent with industry lore: Higher income individuals tend to prefer light and imported beers.²⁶

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

²⁶We have checked that the demand system is robust to the Knittel and Metaxoglou (2008) critique that in a highly nonlinear model such as a random-coefficients logit demand system, the objective function may exhibit many local minima, as described in detail in Appendix D. 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 Knittel and Metaxoglou'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: $0.0567 * e^{-5}$ and $0.6207 * e^{-5}$. We also confirm that our estimates are associated with the lowest value of the objective function with a zero gradient and a positive-definite Hessian, which is consistent with our identified optimum being the global minimum.

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 the super-elasticity of demand to be 0.8 in our random-coefficients model, that is, that a 10-percent increase in prices leads to an 8-percent increase in the absolute value of the price elasticity of demand. This generates a reasonable incentive for firms to adjust their markups.²⁷

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

Once the parameters of the demand system have been estimated, we compute the market share function $s_{jt}(p_t^r)$ as well as the own- and cross-price derivatives $\frac{\partial s_{kt}}{\partial p_{jt}^r}$. Then we use the retailer’s first-order conditions for each product j (6) to back out each product’s markups, which in turn allow 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 in only for those products whose prices adjust from the previous period. As discussed earlier, many of the price changes in our data reflect promotions, during which the price of a particular brand is reduced for a few weeks (see also Figures 1 and 2). A striking characteristic of these promotions is that product prices return to their exact pre-promotion level once the promotion is over. In theory, the transition from the discount price to the pre-promotion level is a price change that could be handled in the same manner as a level change in price (after all, firms do incur some cost every time they change the posted price); yet, given that firms seem to charge *exactly* the same price that they were charging before the promotion, we were skeptical about the plausibility of the assumption that the post-promotion prices are determined based on firms’ first-order conditions. To be safe, we conducted the empirical analysis both ways, first applying the FOC’s to all periods in which the price changed (including 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 show in Appendix E, but the second approach significantly reduces the number of observations associated with a price change that we can

²⁷We 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.

exploit in the empirical analysis. Still, in the remainder of the paper we report results based on this second, more conservative approach, as we are more comfortable with the assumption that FOC’s hold only in those periods during which a firm charges a price that is genuinely different from the price charged in earlier periods.

Table 5 reports retail and wholesale prices and our estimates of markups and costs for selected foreign brands. The markups that are derived based on this approach appear reasonable and consistent with industry wisdom: We estimate retail markups to be on the order of 5-10 percent of the retail price which coincides with estimates from industry reports 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 for which we observe the price of a product adjusting, so that we can reasonably assume that the retailer sets the new price according to its first-order conditions. However, this approach does not work in periods in which the price does not change. To get estimates of the non-traded costs for these periods, we 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 in Table 5 show 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 \$5.45 for a six-pack of imported beer. One exception is *Bass*, with an average retailer non-traded cost closer to 20 cents. Recall that 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. Thus, we expect that these cross-product differences primarily reflect measurement error in our model’s estimates of the retailer’s total marginal costs.²⁸ 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

²⁸In addition, some products may incur higher per-unit energy costs depending on their location (e.g. if they are exclusively refrigerated or also “dry shelf”, i.e., not refrigerated, in industry parlance), though we do not observe data on this in our market.

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 (10) and (12) to derive the upper and lower bounds of the retailer adjustment costs A_{jt}^r . The computation of the upper bound is straightforward: in equation (10) all variables are observed, except for the counterfactual market share $s_{jt}^c(p_{jt-1}^r, p_{kt}^r)$ that product j would have if the retailer did not change her price from the previous period. This counterfactual share can easily be evaluated once the demand parameters are estimated, given that the market share function is known. Computing the lower bound based on (12) 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 $s_{jt}^c(p_{jt}^{rc}, p_{kt}^r)$. These are computed using equation (8) 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 \$38.18 for *Bass* to \$112.54 for *Heineken*, with a mean lower bound across foreign brands of \$46.14, and a mean upper bound of almost \$255. The entries in the third and fourth columns of Table 7 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 of repricing costs across all foreign brands is 3.5 percent of total revenue and for the lower bounds is 0.63 of total revenue. 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 in order to perform simulations. Nevertheless, we find it interesting that the order of magnitude of our estimates is similar to the one obtained in other studies that used a completely different methodology.²⁹

²⁹This procedure involves summing up the estimated upper bounds for the price adjustment costs for those periods within a year in which prices *did* change, and then dividing this sum by the firm’s revenue over the entire year. We should emphasize that because our “price adjustment costs” are defined in the most general sense to include *all* factors that may prevent firms from changing their nominal prices (and *not just the literal labor and material costs* of changing prices), our numbers are not directly comparable to Levy et al’s (1997). Still, it is interesting that despite these differences the two sets of numbers are of similar order of magnitude.

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 (16) to back out their marginal costs c_{jt}^w . Manufacturers’ mean markup from our model is 40 cents, which given the mean retail price of \$5.45 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 used above 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 for periods when wholesale prices do not adjust, given that firms’ first-order conditions do not necessarily hold. To do this, we model the total manufacturer costs parametrically as a function of observables, and estimate this function using data from the periods of wholesale price adjustment *only*. We assume that manufacturers marginal costs c_{jt}^w take the form:

$$c_{jt}^w = \exp(\theta_j + \omega_{jt})(w_t^d)^{\theta_{dw}}(e_{jt}w_t^f)^{F_j*\theta_{fw}}(p_{bjt})^{D_j*\theta_{dp}}(e_{jt}p_{bjt})^{F_j*\theta_{fp}} \quad (31)$$

or, in log-terms:

$$\ln c_{jt}^w = \theta_j + \theta_{dw} \ln w_t^d + F_j * \theta_{fw} \ln(e_{jt}w_t^f) + D_j*\theta_{dp} \ln(p_{bjt}) + F_j*\theta_{fp} \ln(e_{jt}p_{bjt}) + \omega_{jt} \quad (32)$$

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 (32) 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 is subject to exchange-rate fluctuations. For domestic producers the traded component will be (by definition) zero. Foreign producers selling in the U.S. will generally have both traded and local non-traded costs. The latter are captured in the above specification by the

term $(w_t^d)^{\theta_{dw}}$ that indicates the dependence of foreign producers’ marginal costs on the local wages in the U.S. The specification in (31) can be used to demonstrate two important facts regarding foreign suppliers’ costs. First, foreign producers selling to the U.S. will typically experience substantially more volatility than domestic producers due to their exposure to exchange-rate shocks. Second, as long as the local non-traded cost component is nonzero (so that $\theta_{fw} + \theta_{fp} < 1$), 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.³⁰ As reported in Table 8, the highly significant coefficient of 0.58 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 reported by both Burstein, Neves, and Rebelo (2003) and Goldberg and Campa (2010).³¹ This estimate implies that a big part of foreign manufacturers’ costs of selling in the U.S. market are not affected by exchange-rate fluctuations. Hence it comes as no surprise that foreign producers do not fully adjust their U.S. dollar prices in response to exchange-rate changes. This finding implies that even without menu costs, the existence of local non-traded costs can generate a significant degree of inertia in local currency prices.

4.1.7 Derivation of bounds for the wholesale price adjustment costs \mathbf{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 (18) and (20).³² The estimates of the manufacturers’ adjustment-cost bounds are reported in Table 9, are roughly the same order of magnitude as those for retailers. Their upper bounds range from \$32.42 for *Bass* to \$516.10 for *Beck’s*, and their lower bounds from \$2.40 for *Bass* to \$97.42 for *Beck’s*. The entries in the third and fourth columns of Table 9 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 2.07 percent of

³⁰ As discussed earlier, because distributors’ pricing is coordinated by brewers, we treat the manufacturer and distributor as one entity, so that the “local” manufacturer costs include the marketing and distribution costs incurred by the distributor.

³¹ Burstein, 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 (expenditure on transport, wholesale and retail trade services, and marketing) 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 ca. 65 percent of the incomplete pass-through in their study of the European automobile industry.

³² We set out the computational details of this step in Appendix F.

total revenue and for the lower bounds is 0.12 percent of total revenue.

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 \$900 for a “permanent” level price change, and closer to \$50 for a temporary price reduction. This finding provides empirical support for a recent paper by Kehoe and Midrigan (2008) that argues that the fixed cost of changing a regular price is larger than the fixed cost of a temporary reduction.

4.2 Simulations

4.2.1 Description of 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 in order 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 (32), 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 on the part of the retailer³³.

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 42 percent, which is determined by the coefficient on local wages (0.58) from the regressions 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 31.4 percent for *Bass* to 40.0 percent

³³We discuss the expressions used in the decomposition formally in Appendix A.

for *Beck's*. It is 33.2 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.

With retailer local costs taken into account, the median pass-through falls to 30.5 percent and ranges from 27.2 percent for *Bass* to 35.6 percent for *Beck's*. Finally, the median retailer pass-through elasticity across all brands is 23.4 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 18.0 percent for *Heineken* to 24.7 percent for *Beck's*. Note that these results are somewhat larger than the 5-10 percent retail pass-through elasticities estimated in Section 2.3.

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 32 percent in Simulation 2, and whose median retail elasticity falls from 18 percent in Simulation 1 to under 6

percent in Simulation 2, a 12 percentage-point decline attributable entirely to the strategic effects of repricing costs. The median retail pass-through elasticity across all four brands falls from 23.4 percent in Simulation 1 to 18 percent in Simulation 2, a modest but still notable 5 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 42 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 31.4 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 2.3. The mean manufacturer pass-through elasticity across all four brands in this final counterfactual is 3.3 percent, which is more comparable, and quite close to the (mean) manufacturer pass-through elasticity of 4 percent estimated via reduced-form regressions in Section 2.3.

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 58 percent, of the observed retail-price inertia. Manufacturers’ markup adjustment accounts for 9.2 percent of the remaining adjustment, their competitors’ price rigidities for 1.4 percent, and their own repricing costs for another 31.4 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 31.4 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 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 Conclusions

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 for the fixed costs of price adjustment facing retailers and manufacturers, we find that 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 in order 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 9% 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 into this issue.

References

- [1] "Can You Judge a Beer by its Label? Beer Ratings," 1996. *Consumer Reports*, June, pp. 10-19.
- [2] "The Alcoholic Beverage Market in the United States," 1998. Canadian Trade Commissioner, June.
- [3] Aguirregabiria, V. 1999. "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, vol. 66, no. 2, pp. 275-308.
- [4] Akerlof, G. and J. Yellen. 1985. "A Near-Rational Model of the Business Cycle with Wage and Price Inertia," *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* (forthcoming).
- [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] Bills, M. and P. J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, vol. 112, no. 5, pp. 947-85.
- [9] Blinder, A., E. Canetti, D. Lebow, and J. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation.
- [10] 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.
- [11] Caplin, A. and J. Leahy. 1997. "Aggregation and Optimization with State-Dependent Pricing," *Econometrica*, vol. 65, pp. 601-26.
- [12] Carlton, D. 1986. "The Rigidity of Prices," *The American Economic Review*, vol. 76, pp. 637-58.
- [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, no. 1, pp. 213-42.

- [16] Dutta, S., M. Bergen, D. Levy, and R. Venable. 1999. "Menu Costs, Posted Prices, and Multiproduct Retailers," *Journal of Money, Credit, and Banking*, vol. 31, no. 4 (November), pp. 683-703.
- [17] Engel, C. 2002. "Expenditure Switching and Exchange-Rate Policy," *NBER Macroeconomics Annual*, pp. 231-72.
- [18] Engel, C. 2003. "On the Relationship between Pass-Through and Sticky Nominal Prices." Working Paper.
- [19] Froot, K.A., and P.D. Klemperer. 1989. "Exchange Rate Pass-through When Market Share Matters," *The American Economic Review*, vol. 79, no. 4, pp. 637-54.
- [20] Goldberg, L. and J. Campa. 2010. "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure," *The Review of Economics and Statistics*, vol. 92, no. 2, pp. 392-407.
- [21] Goldberg, P. K. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, vol. 63, no. 4, pp. 891-951.
- [22] Goldberg, P. K. and M. Knetter. 1997. "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature*, vol. 35, pp. 1243-73.
- [23] 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.
- [24] Goldberg, P.K. and F. Verboven. 2005. "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market," *Journal of International Economics*, vol. 65, pp. 49-73.
- [25] 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.
- [26] Hellerstein, R. 2004. *Empirical Essays on Vertical Contracts, Exchange Rates, and Monetary Policy*. U.C. Berkeley Ph.D. Dissertation.
- [27] Hendel, I. and A. Nevo. 2006a. "Sales and Consumer Inventory," *Rand Journal of Economics*, vol. 37, no. 3, pp. 543-61.
- [28] Hendel, I. and A. Nevo. 2006b. "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, vol. 74, no. 6, pp. 1637-73.

- [29] Kehoe, P. and V. Midrigan. 2008. "Temporary Price Changes and the Real Effects of Monetary Policy," Federal Reserve Bank of Minneapolis Staff Report 413.
- [30] Kimball, M. S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, vol. 27, no. 4, pp. 1241-77.
- [31] Klenow, P. and O. Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?," *The Quarterly Journal of Economics*, vol. 123, no. 3, pp. 863-904.
- [32] Klenow, P. and J. Willis. 2006. "Real Rigidities and Nominal Price Changes," Research Working Paper 06-03, Federal Reserve Bank of Kansas City.
- [33] Knittel, C. and K. Metaxoglou. 2008. "Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings," U.C. Davis Working Paper.
- [34] Krugman, P. 1987. "Pricing to Market When the Exchange Rate Changes," In J.D. Richardson and S. Arndt (eds.), *Real-Financial Linkages Among Open Economies*, Cambridge, MA: MIT Press.
- [35] 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.
- [36] Mankiw, N.G. 1985. "Small Menu Costs and Large Business Cycles," *The Quarterly Journal of Economics*, vol. 100, pp. 529-39.
- [37] Midrigan, V. 2010. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," New York University Working Paper.
- [38] Nakamura, E. and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, vol. 123, no. 4, pp. 1415-64.
- [39] Nakamura, E. and D. Zerom. 2010. "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, vol. 77, pp. 1192-1230.
- [40] Nevo, A. 2000. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics and Management Strategy*, vol. 9, no. 4, pp. 513-58.
- [41] Nevo, A. 2001. "Measuring Market Power in the Ready-To-Eat Cereal Industry," *Econometrica*, vol. 69, no. 2, pp. 307-42.
- [42] Pesendorfer, M. 2002. "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, vol. 75, no. 1, pp. 33-66.

- [43] 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.
- [44] Slade, M. 1998. "Optimal Pricing with Costly Adjustment and Persistent Effects: Empirical Evidence from Retail-Grocery Prices", *Review of Economic Studies*, vol. 65, no. 1, pp. 87-107.
- [45] Sobel, J. 1984. "The Timing of Sales," *Review of Economic Studies*, vol. 51, pp. 353-68.
- [46] Tremblay, V.J. and C.H. Tremblay. 2005. *The U.S. Brewing Industry*. Cambridge: MIT Press.
- [47] Varian, H. R. 1980. "A Model of Sales," *The American Economic Review*, vol. 70, pp. 651-59.
- [48] Villas-Boas, S.B. 2007. "Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data," *Review of Economic Studies*, vol. 74, no. 2, pp. 625-52.

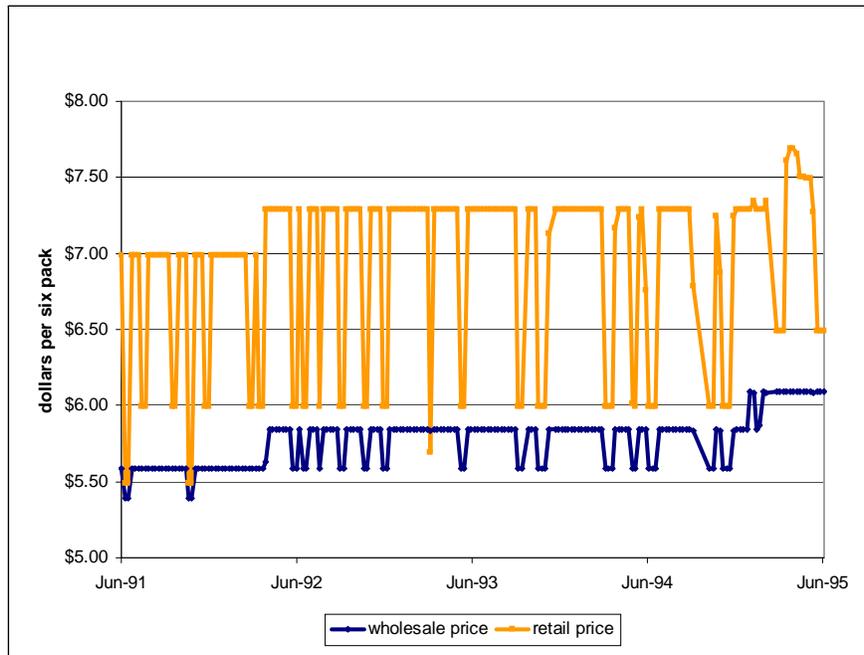


Figure 1: *Weekly retail and wholesale prices for Bass Ale.* Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's*.

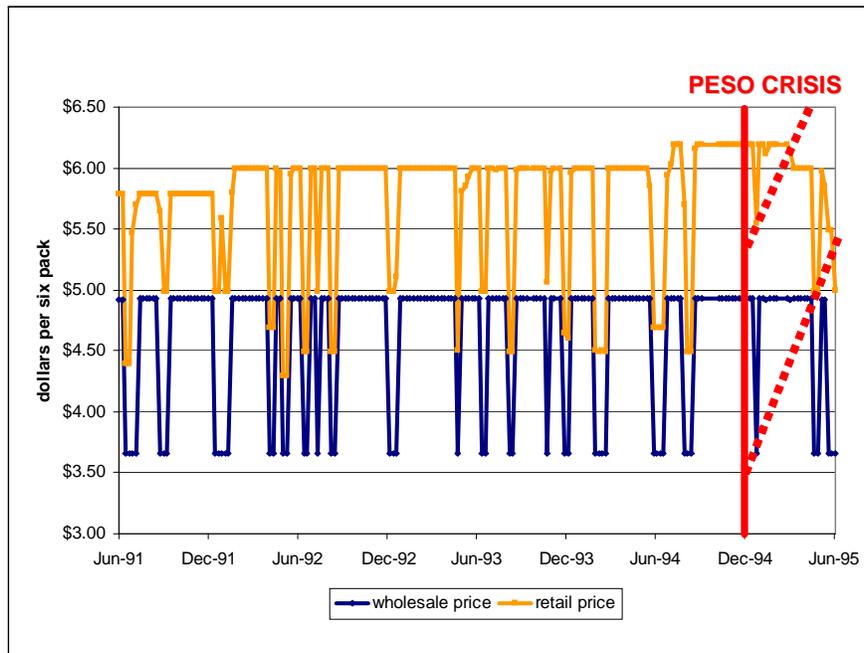


Figure 2: *Weekly retail and wholesale prices for Corona.* Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's*.

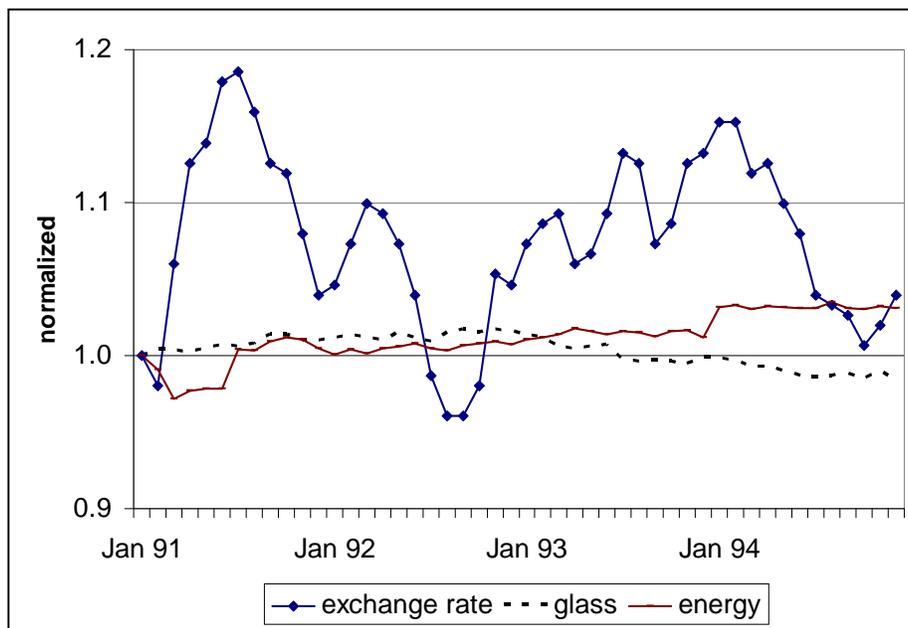


Figure 3: *The nominal exchange rate varies by more than do typical input prices for German brewers.* Source: *OECD*.

Description	Mean	Median	Standard Deviation
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	4.52	4.60	.68
Bitterness	2.50	2.10	1.08

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)**
Constant	1.83 (.02)**	1.79 (.16)**	1.66 (.08)**	1.67 (.10)**	.08 (.04)
Observations	3636	3636	3636	3636	3636
R^2	.65	.65	.81	.81	.80

Table 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 United States (dollars per unit of foreign currency). All regressions include brand, price-zone, and week fixed effects. The second and fourth columns of the table report results from regressions with controls for domestic and foreign costs. Robust standard errors in parentheses: Those starred are significant at the *5-percent or **1-percent level. Source: Authors' calculations.

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

Table 3: *Diagnostic results from the multinomial logit model of demand.* The dependent variable is $\ln(s_{jt}) - \ln(s_{Ot})$. Both regressions include brand fixed effects. Huber-White robust standard errors are reported in parentheses. Costs are domestic wages in the beverage industry interacted with weekly nominal exchange rates for foreign brands. Source: Authors' calculations.

Variable	Mean in Population	Interaction with Income
Constant	-16.39* (.20)	
Price	-2.48* (.05)	40.37* (18.17)
Bitterness	-5.57* (.20)	
Percent Alcohol	7.35* (.03)	-51.84* (20.32)
Minimum-Distance R^2	0.18	

Table 4: *Results from the full random-coefficients logit model of demand.* Based on 6464 observations. Asymptotically robust standard errors in parentheses. Starred coefficients are significant at the 5-percent level. Source: Authors' calculations.

	Bass	Becks	Corona	Heineken	All Imports
Retailer					
Price	6.36 (0.64)	5.28 (0.49)	5.04 (0.62)	5.66 (0.70)	5.45 (1.14)
Markup	0.41 (0.005)**	0.40 (0.003)**	0.40 (0.004)**	0.40 (0.004)**	0.40 (0.003)**
Nontraded costs					
Backed out	0.21 (0.070)**	0.37 (0.045)**	0.37 (0.091)**	0.43 (0.070)**	0.38 (0.038)**
Fitted	0.18 (0.071)**	0.33 (0.047)**	0.34 (0.091)**	0.41 (0.070)**	0.37 (0.041)**
Manufacturer					
Price	5.77 (0.24)	4.44 (0.16)	3.85 (0.46)	4.95 (0.27)	4.62 (0.83)
Markup	0.40 (0.005)**	0.40 (0.003)**	0.39 (0.003)**	0.39 (0.006)**	0.40 (0.003)**
Total costs					
Backed out	5.36 (0.04)**	4.02 (0.01)**	3.46 (0.07)**	4.60 (0.07)**	4.18 (0.05)**
Fitted	5.28 (0.03)**	4.02 (0.01)**	3.43 (0.05)**	4.51 (0.07)**	4.41 (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	.48 (.05)**
Constant	-4.11 (.51)**
R^2	.12
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. Source: Authors’ calculations.

Brand	Mean Cost		Share of Brand's Revenue	
	Upper Bounds	Lower Bounds	Upper Bounds	Lower Bounds
	(\$)	(\$)	(%)	(%)
Bass	\$101.92 (58.45)*	\$38.18 (9.40)**	3.01 (2.01)	1.13 (0.31)**
Beck's	\$531.26 (273.89)*	\$49.96 (9.21)**	2.74 (1.30)**	0.26 (0.06)**
Corona	\$139.24 (70.95)*	\$72.62 (41.99)*	1.16 (0.54)**	0.60 (0.0)**
Heineken	\$652.58 (249.73)**	\$112.54 (30.00)**	3.68 (1.77)**	0.63 (0.16)**
All	\$254.98 (78.96)**	\$46.14 (7.97)**	3.49 (1.08)**	0.63 (0.12)**

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. 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	.58 (.08)**
Price foreign barley	.15 (.04)**
Foreign wages	.28 (.05)**
Observations	336

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	\$32.42 (18.54)*	\$2.40 (0.35)**	0.67 (0.39)*	0.05 (0.02)**
Beck's	\$516.10 (121.45)**	\$97.42 (21.25)**	3.39 (0.85)**	0.64 (0.15)**
Corona	\$380.46 (35.55)**	\$24.72 (7.61)**	2.39 (0.23)**	0.16 (0.01)**
Heineken	\$231.00 (73.02)**	\$53.34 (4.83)**	0.31 (0.10)**	0.07 (0.02)**
All	\$306.22 (50.31)**	\$18.36 (3.03)**	2.07 (0.34)**	0.12 (0.02)**

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. 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	907.00 (34.03)**
Dummy for all retail price changes	55.44 (10.21)**
Overall R^2	.19
Observations	3636

Table 10: *Regression of retailer's fixed adjustment costs on a dummy for a level retail-price change.* The regression includes brand and price zone fixed effects. Source: Authors' calculations.

	Manufacturer		Retailer	
	Traded	Markup Adjustment	Traded	Markup Adjustment
No repricing costs				
Bass	42.0 (0.0)**	31.4 (7.4)**	27.2 (6.5)**	20.4 (7.9)**
Beck's	42.0 (0.0)**	40.0 (3.2)**	35.6 (3.7)**	24.7 (4.2)**
Corona	42.0 (0.0)**	33.2 (11.8)**	29.2 (7.0)**	24.5 (8.1)**
Heineken	42.0 (0.0)**	37.9 (5.6)**	32.4 (6.1)**	18.0 (5.9)**
All	42.0 (0.0)**	33.2 (2.9)**	30.5 (3.6)**	23.4 (4.3)**
Competitor-brand repricing costs				
Bass	42.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (7.9)**
Becks	42.0 (0.0)**	33.0 (10.0)**	27.1 (7.3)**	18.0 (8.7)**
Corona	42.0 (0.0)**	32.5 (8.0)**	28.6 (8.9)**	17.8 (7.6)**
Heineken	42.0 (0.0)**	32.0 (6.7)**	26.4 (6.2)**	5.7 (2.3)**
All	42.0 (0.0)**	31.4 (3.3)**	26.5 (3.9)**	18.0 (4.1)**
Own-brand repricing costs				
Bass	42.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Becks	42.0 (0.0)**	0.0 (0.0)	0.0 (0.1)	0.0 (0.0)
Corona	42.0 (0.0)**	0.0 (0.0)	0.0 (0.1)	0.0 (0.0)
Heineken	42.0 (0.0)**	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
All	42.0 (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 Other	Costs of Repricing Own	Local Costs	Markup Adjustment	Costs of Repricing Other	Costs of Repricing Own	
Bass	58.0	10.6	31.4	0.0	0.0	0.0	0.0	0.0	100.0
Beck's	58.0	2.2	6.8	33.0	0.0	0.0	0.0	0.0	100.0
Corona	58.0	9.1	0.4	32.5	0.0	0.0	0.0	0.0	100.0
Heineken	58.0	4.0	6.0	32.0	0.0	0.0	0.0	0.0	100.0
All	58.0	9.2	1.4	31.4	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. Similarly, the contribution of firms' own costs of price adjustment to the incomplete pass-through 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 Static Expectations

In a dynamic framework in which a firm maximizes the expected value of its present discounted future profits, the firm's problem is:

$$Max[E_0 \sum_{t=0}^{\infty} \delta^t \Pi_{jt}^w] = Max[E_0 \sum_{t=0}^{\infty} \delta^t \{p_{jt}^w - c_{jt}^w(tc_{jt}^w, ntc_{jt}^w) s_{jt}(p_t^r(p^w)) - A_{jt}^w\}]$$

where δ denotes the discount rate, assumed to be constant for simplicity, and as before:

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

In this framework, the only stochastic variable is traded costs (tc_{jt}^w), and uncertainty about these costs derives from uncertainty about exchange rates. Under the assumption that exchange-rate shocks, and hence cost shocks, are permanent, the firm's problem becomes:

$$Max[E_0 \sum_{t=0}^{\infty} \delta^t \{p_j^w - c_j^w(tc_j^w, ntc_j^w) s_j(p^r(p^w)) - A_j^w\}]$$

The absence of the subscript t in the cost and price variables in the above notation indicates that the firm assumes that the values of these variables are going to be permanent. Suppose the firm faces an unexpected exchange-rate shock at time $t = 0$. Then the firm considers whether it should change its price from p_j^w (the price it had set before the shock) to a new price p_j^w . It makes this decision assuming that the shock will be permanent, so that the new price it will set and the resulting market shares will be permanent as well. As before, the firm incurs an adjustment cost if it changes its price from p_j^w to p_j^w . If we use this (myopic) dynamic model to derive adjustment cost bounds, the single-product version of inequality (18) becomes:

$$A_j^w \leq \overline{A_j^w} = \frac{\delta}{1-\delta} \{ (p_j^w - c_j^w) s_j(p^r(p^w)) - (p_j^w - c_j^w) s_j^c(p_j^{rc}(p_j^w, p_k^r)) \}, \quad k \neq j$$

That is, the firm would compare the fixed cost of adjusting its price today to the present discounted value of all future additional profits that it would make if it changed its price. A similar argument applies to the lower bound $\underline{A_j^w}$. This implies that if we adopted a dynamic framework, in order to obtain the bounds for the price-adjustment costs consistent with the dynamic model we would have to multiply the static estimates by the factor $\frac{\delta}{1-\delta}$, which is of the same order of magnitude as the inverse of the interest rate. This would "blow up" the static estimates for the price-adjustment costs. In this sense, one could argue that the static estimates of the adjustment costs – if taken literally –

would lead one to understate the magnitude of these costs. Note however, that this would leave our simulations and decomposition results at the end of the paper *completely unchanged*: in the dynamic setup, *both* sides of the inequalities (18) and (20) that determine whether the firm will adjust its price in response to an exchange-rate shock or not, would be multiplied by the same factor ($\frac{\delta}{1-\delta}$), so that this factor cancels out. This is why we emphasize that the term “price-adjustment costs” is meaningful only within the context of the particular model we use as the basis of the empirical analysis, and that our results regarding pass-through are robust as long as this term is used consistently throughout the analysis.

C Robustness Checks for Longer Horizons

This appendix reports the results of robustness checks that lengthen the horizon over which firms set prices and optimize profits relative to the baseline weekly model. We report results for three alternative models: a bi-weekly model, a monthly model, and a six-weekly model. We show 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 show that the results from the counterfactuals do not differ significantly from those of the baseline model, which 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 report results for the repricing costs’ upper bounds.

Unfortunately, it is not feasible to go beyond a six-week horizon in these robustness checks, as we then have insufficient observations for the regressions used to compute the retailer’s nontraded costs and the manufacturers’ total costs. These already start to show some strain in the six-week model, as we ask the data to do quite a lot in our structural model. We note that lengthening the time horizon over which prices are set reduces our sample size from 6464 observations in the weekly model, to 3328 observations in the bi-weekly model, to 1568 in the monthly model, and to 1120 in the six-weekly model.

Bi-weekly Model Appendix Table C.1 reports the estimation results for the random coefficients demand system for the bi-weekly model. The results are quite similar to those of the baseline case, in particular the coefficient on price, which at -2.24, is quite similar to the -2.48 price coefficient reported in Table 4. The results from the counterfactuals are roughly the same as those for the baseline model, and are reported in Appendix Tables C.2 and C.3. Most of the incomplete pass-through is attributable to the manufacturer’s local costs, which account for 58 percent of the incomplete pass-through. Manufacturer markup adjustment accounts for 4.6 percent, while the strategic effects of repricing costs account for almost 7 percent, their direct effects for 30.5 percent.

Monthly Model Appendix Table C.4 reports the demand estimates for the monthly model. The coefficients are quite similar to those of the baseline model, in particular the coefficient on price, which at -2.32 is quite close that of the baseline weekly model. The counterfactuals’ results, reported in Tables C.5 and C.6, are almost identical to those of the baseline model. The manufacturer’s local costs, own repricing costs, and markup adjustment account for roughly 58, 29, and 7 percent of the incomplete pass-through, respectively. The strategic effects of repricing costs account for the remaining 6 percent.

Six-weekly Model Appendix Table C.7 reports coefficients from estimating the random-coefficients demand system using the six-weekly sample and reveals, again, little difference with the baseline weekly model. The coefficients are similar to those of the weekly model, in particular the coefficient on price, which at -2.20 is identical to that of the biweekly model, and close to the baseline estimate of -2.48. The decomposition is similar to that of the baseline model, with the manufacturer's local costs, indirect and direct effects of repricing costs accounting for 58, 13, and 28 percent of the incomplete pass-through, respectively. There is a slightly smaller role in this model for manufacturers' markup adjustment in the incomplete pass-through, which accounts for about 1 percent of the incomplete pass-through.

Appendix Table C.1 Biweekly Random-Coefficients Demand Model

	Mean in Population	Interaction with Income
Constant	-4.71 (0.14)*	
Price	-2.24 (0.21)*	43.09 (4.00)*
Bitterness	2.33 (0.14)*	
Percent Alcohol	1.79 (0.01)*	-52.51 (3.46)*
Observations	3328	

Asymptotically robust standard errors in parentheses. Starred coefficients are significant at the 5-percent level. Source: authors' calculations.

Appendix Table C.2 Counterfactual Experiments for Biweekly Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	31.3%	29.7%	23.9%
	<i>becks</i>	42.0%	37.8%	35.1%	28.2%
	<i>corona</i>	42.0%	37.6%	33.6%	25.9%
	<i>heineken</i>	42.0%	37.8%	32.5%	28.0%
	<i>all</i>	42.0%	37.4%	33.2%	26.7%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	36.3%	33.8%	16.7%
	<i>corona</i>	42.0%	30.9%	27.6%	18.3%
	<i>heineken</i>	42.0%	36.1%	30.4%	18.9%
	<i>all</i>	42.0%	30.5%	26.6%	16.5%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	0.0%	0.0%	0.0%
	<i>corona</i>	42.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	42.0%	0.0%	0.0%	0.0%
	<i>all</i>	42.0%	0.0%	0.0%	0.0%

Counterfactual experiments: Median pass-through of a 5-percent appreciation of the relevant foreign currency. Median over 208 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.

Appendix Table C.3 Decomposition for Biweekly Model

<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>	58.0%	10.7%	31.3%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	4.2%	1.5%	36.3%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	58.0%	4.4%	6.7%	30.9%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	58.0%	4.2%	1.7%	36.1%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	4.6%	6.9%	30.5%	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 208 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.

Appendix Table C.4 Monthly Random-Coefficients Demand Model

	Mean in Population	Interaction with Income
Constant	-13.22 (0.28)*	
Price	-2.32 (0.21)*	40.84 (14.31)*
Bitterness	0.40 (0.28)	
Percent Alcohol	2.52 (0.07)*	-49.71 (20.73)*
Observations	1568	

Asymptotically robust standard errors in parentheses. Starred coefficients are significant at the 5-percent level. Source: authors' calculations.

Appendix Table C.5 Counterfactual Experiments for Monthly Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	32.3%	29.5%	24.3%
	<i>becks</i>	42.0%	35.9%	33.1%	26.4%
	<i>corona</i>	42.0%	36.1%	33.7%	25.9%
	<i>heineken</i>	42.0%	36.0%	32.4%	26.0%
	<i>all</i>	42.0%	34.7%	32.0%	25.6%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	34.0%	31.5%	14.6%
	<i>corona</i>	42.0%	29.9%	26.9%	18.0%
	<i>heineken</i>	42.0%	30.4%	26.1%	23.8%
	<i>all</i>	42.0%	28.6%	25.1%	18.2%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	0.0%	0.0%	0.0%
	<i>corona</i>	42.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	42.0%	0.0%	0.0%	0.0%
	<i>all</i>	42.0%	0.0%	0.0%	0.0%

Counterfactual experiments: Median pass-through of a 5-percent appreciation of the relevant foreign currency. Median over 98 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.

Appendix Table C.6 Decomposition for Monthly Model

<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>	58.0%	9.7%	32.3%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	6.1%	1.9%	34.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	58.0%	5.9%	6.2%	29.9%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	58.0%	6.0%	5.6%	30.4%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	7.3%	6.1%	28.6%	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 98 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.

Appendix Table C.7 Six-weekly Random-Coefficients Demand Model

	Mean in Population	Interaction with Income
Constant	-6.79 (0.28)*	
Price	-2.20 (0.03)*	42.20 (7.54)*
Bitterness	0.81 (0.28)*	
Percent Alcohol	1.82 (0.02)*	-53.50 (1.17)*
Observations	1120	

Asymptotically robust standard errors in parentheses. Starred coefficients are significant at the 5-percent level. Source: authors' calculations.

Appendix Table C.8 Counterfactual Experiments for Six-Week Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	37.3%	34.7%	30.0%
	<i>becks</i>	42.0%	41.6%	38.9%	17.3%
	<i>corona</i>	42.0%	45.0%	42.5%	28.8%
	<i>heineken</i>	42.0%	42.0%	38.4%	12.3%
	<i>all</i>	42.0%	40.8%	38.1%	17.9%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	38.8%	36.3%	19.9%
	<i>corona</i>	42.0%	27.0%	24.4%	17.5%
	<i>heineken</i>	42.0%	35.3%	31.3%	12.3%
	<i>all</i>	42.0%	28.3%	25.2%	17.3%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	0.0%	0.0%	0.0%
	<i>corona</i>	42.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	42.0%	0.0%	0.0%	0.0%
	<i>all</i>	42.0%	0.0%	0.0%	0.0%

Counterfactual experiments: Median pass-through of a 5-percent appreciation of the relevant foreign currency. Median over 70 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.

Appendix Table C.9 Decomposition for Six-Week Model

<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>	58.0%	4.7%	37.3%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	0.4%	2.8%	38.8%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	58.0%	-3.0%	18.0%	27.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	58.0%	0.0%	6.7%	35.3%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	1.2%	12.5%	28.3%	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 98 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.

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 and we report the gradient's norm and the Hessian's eigenvalues in footnote 26 of the paper. (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: $0.0567 * e^{-5}$ and $0.6207 * e^{-5}$.

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: $2.0705 * e^{-18}$.

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. As consumers' income rises, they become less price sensitive. The random coefficient on income, at 40.37, is significant at the five-percent level. The mean preference in the population is not amenable to a bitter taste in beer, which has a negative and significant coefficient. As the percent alcohol rises across brands, the mean utility in the population also rises, an intuitive result. There is heterogeneity in the population with respect to this characteristic: Those with higher incomes get less utility from a high percent of alcohol in their beer, that is, prefer light beers, as one can glean from the negative and significant random coefficient of -51.84. This is consistent with industry lore: Higher income individuals tend to prefer light and imported beers."

E Robustness Checks for Different Types of Price Adjustments

This appendix reports the results of robustness checks that each vary the criteria for constructing the sample of price changes used to compute the structural model’s cost parameters. We report the counterfactual results for three alternative models: the first 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 4.1.2. Although we emphasize throughout the paper that the magnitudes of the adjustment costs are meaningful only in the context of the particular model we use, we report them for each of the alternative models we run to provide intuition for any departures from our baseline model in the counterfactuals.

Promotional Price Changes Appendix Table E.1 reports the upper and lower bounds on repricing costs when only promotional price changes are included in the sample used to compute the cost parameters. The retailer’s upper bounds are somewhat lower than in the baseline case, while the results for the other bounds are similar. The results from the counterfactuals for the upper bounds are roughly the same as those from the baseline model, and are reported in Appendix Tables E.2 and E.3.

Level Price Changes Appendix Table E.4 reports the upper and lower bounds on repricing costs when only level price changes are included in the sample used to compute the cost parameters. These bounds cannot be computed for several of the brands (*Corona* and *Molson Golden*) that do not exhibit any level price changes in either their retail or manufacturer price, or in both. The bounds are somewhat wider than in the baseline case with much higher upper bounds in particular, for both the manufacturers and the retailer. This is consistent with our conjecture that permanent adjustments in prices require more managerial time and effort to implement than do promotional adjustments. We run counterfactual experiments for those brands for which we have bounds, and find that the results from the counterfactuals do not differ significantly from the baseline model for the upper bounds, as reported in Tables E.5 and E.6.

All Price Changes Appendix Table E.7 reports the upper and lower bounds on repricing costs when all price changes are included in the sample used to compute the model’s cost parameters. The estimates of the bounds are somewhat wider than in the baseline case, which is consistent with our assumption that post-promotion prices may not be set optimally according to firms’ first-order conditions. The upper bounds associated with post-promotions price adjustments also appear to be similar in magnitude to those of level price changes, both of which tend to be price increases. The

results from the counterfactuals for the upper bounds are reported in Tables E.8 and E.9 and do not differ significantly from those of the baseline model.

Appendix Table E.1 Bounds for Promotions-Only Model

<i>brand</i>	Retailer		Manufacturer	
	upper bounds	lower bounds	upper bounds	lower bounds
<i>amstel</i>	42.02	7.56	94.18	10.62
<i>bass</i>	43.00	0.22	19.28	2.36
<i>beck's</i>	271.50	64.78	625.46	102.28
<i>corona</i>	151.20	91.84	376.74	25.26
<i>heineken</i>	241.14	95.28	211.10	63.86
<i>molson golden</i>	101.36	15.14	278.56	28.44
<i>all</i>	122.72	26.76	386.28	25.44

The entries in the first two columns report the mean over time of the dollar value of the retailer's price adjustment costs, and in the third and fourth columns the mean of the manufacturers' price adjustment costs. Source: Authors' calculations.

Appendix Table E.2 Counterfactual Experiments for Promotions-Only Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	31.4%	27.2%	20.4%
	<i>becks</i>	42.0%	40.0%	35.6%	24.7%
	<i>corona</i>	42.0%	33.2%	29.2%	24.5%
	<i>heineken</i>	42.0%	37.9%	32.4%	18.0%
	<i>all</i>	42.0%	33.2%	30.5%	23.4%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	33.0%	27.1%	18.0%
	<i>corona</i>	42.0%	32.5%	28.6%	17.8%
	<i>heineken</i>	42.0%	32.0%	26.4%	5.7%
	<i>all</i>	42.0%	31.4%	26.5%	18.0%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	0.0%	0.0%	0.0%
	<i>corona</i>	42.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	42.0%	0.0%	0.0%	0.0%
	<i>all</i>	42.0%	0.0%	0.0%	0.0%

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. Source: Authors' calculations.

Appendix Table E.3 Decomposition for Promotions-Only Model

<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>	58.0%	10.6%	31.4%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	2.2%	6.8%	33.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	58.0%	9.1%	0.4%	32.5%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	58.0%	4.0%	6.0%	32.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	9.2%	1.4%	31.4%	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 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.

Appendix Table E.4 Bounds for Level-Changes Model

<i>brand</i>	Retailer		Manufacturer	
	upper bounds	lower bounds	upper bounds	lower bounds
<i>amstel</i>	543.90	3.74	9.62	0.00
<i>bass</i>	558.54	4.52	4.24	0.00
<i>beck's</i>	4195.72	0.00	3291.98	0.24
<i>corona</i>	11.76	4.12	.	.
<i>heineken</i>	3409.20	1.56	235.72	0.32
<i>molson golden</i>	.	2.02	.	0.00
<i>all</i>	1935.50	3.72	571.72	0.14

The entries in the first two columns report the mean over time of the dollar value of the retailer's price adjustment costs, and in the third and fourth columns the mean of the manufacturers' price adjustment costs. Source: Authors' calculations.

Appendix Table E.5 Counterfactual Experiments for Level-Changes Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	31.4%	27.2%	20.4%
	<i>becks</i>	42.0%	40.0%	35.6%	24.7%
	<i>corona</i>	42.0%	33.2%	29.2%	24.5%
	<i>heineken</i>	42.0%	37.9%	32.4%	18.0%
	<i>all</i>	42.0%	33.2%	30.5%	23.4%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.0%	0.0%
	<i>becks</i>	42.0%	33.0%	27.1%	18.0%
	<i>corona</i>	42.0%	32.5%	28.6%	17.8%
	<i>heineken</i>	42.0%	32.0%	26.4%	5.7%
	<i>all</i>	42.0%	31.4%	26.5%	18.0%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.0%	0.1%	0.0%
	<i>becks</i>	42.0%	0.0%	0.1%	0.0%
	<i>corona</i>	na	na	na	na
	<i>heineken</i>	42.0%	0.0%	0.1%	0.0%
	<i>all</i>	42.0%	0.0%	0.1%	0.0%

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. Source: Authors' calculations.

Appendix Table E.6 Decomposition for Level-Changes Model

<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>	58.0%	10.6%	31.4%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	2.2%	6.8%	33.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	na	na	na	na	na	na	na	na	na
<i>heineken</i>	58.0%	4.0%	6.0%	32.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	9.2%	1.4%	31.4%	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 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.

Appendix Table E.7 Bounds for All-Price-Changes Model

<i>brand</i>	Retailer		Manufacturer	
	upper bounds	lower bounds	upper bounds	lower bounds
<i>amstel</i>	637.22	23.02	54.64	8.62
<i>bass</i>	487.36	4.70	26.50	2.10
<i>beck's</i>	2250.36	4.02	440.54	114.26
<i>corona</i>	2865.10	46.08	725.38	18.08
<i>heineken</i>	2908.64	38.40	578.98	38.46
<i>molson golden</i>	1764.72	0.88	175.24	11.60
<i>all</i>	1586.98	17.50	344.08	20.78

The entries in the first two columns report the mean over time of the dollar value of the retailer's price adjustment costs, and in the third and fourth columns the mean of the manufacturers' price adjustment costs. Source: Authors' calculations.

Appendix Table E.8 Counterfactual Experiments for All-Price-Changes Model

		Manufacturer		Retailer	
		traded	markup adj	traded	markup adj
<i>No repricing costs</i>	<i>bass</i>	42.0%	27.1%	27.2%	20.4%
	<i>becks</i>	42.0%	40.0%	35.6%	24.7%
	<i>corona</i>	42.0%	33.2%	29.2%	24.5%
	<i>heineken</i>	42.0%	37.9%	32.4%	18.0%
	<i>all</i>	42.0%	33.2%	30.5%	23.4%
<i>Other repricing costs</i>	<i>bass</i>	42.0%	0.1%	0.0%	0.0%
	<i>becks</i>	42.0%	33.0%	27.1%	18.0%
	<i>corona</i>	42.0%	32.5%	28.6%	17.8%
	<i>heineken</i>	42.0%	32.0%	26.4%	5.7%
	<i>all</i>	42.0%	31.4%	26.5%	18.0%
<i>Own repricing costs</i>	<i>bass</i>	42.0%	0.1%	0.0%	0.0%
	<i>becks</i>	42.0%	0.1%	0.0%	0.0%
	<i>corona</i>	42.0%	0.0%	0.0%	0.0%
	<i>heineken</i>	42.0%	0.1%	0.0%	0.0%
	<i>all</i>	42.0%	0.1%	0.0%	0.0%

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. Source: Authors' calculations.

Appendix Table E.9 Decomposition for All-Price-Changes Model

<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>	58.0%	10.6%	31.4%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>becks</i>	58.0%	2.2%	6.8%	33.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>corona</i>	58.0%	9.1%	0.4%	32.5%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>heineken</i>	58.0%	4.0%	6.0%	32.0%	0.0%	0.0%	0.0%	0.0%	100.0%
<i>all</i>	58.0%	9.2%	1.4%	31.4%	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 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.

F Computation of Manufacturers' Bounds

Consider inequality (18) that determines the upper bound on manufacturers' adjustment costs. Once steps 1-6 of Section 4.1 are completed, all variables in this inequality are known, except for the counterfactual retail price p_{jt}^{rc} that the retailer would charge if the manufacturer did not change her price that period. The counterfactual price p_{jt}^{rc} can take on one of two values: it is either equal to p_{jt-1}^r , or it is determined according to the retailer's first-order condition, *conditional on the retailer observing the wholesale price* p_{jt-1}^w . To determine which of the two prices the retailer will choose, we first solve for the optimal price that the retailer would pick if she behaved according to her profit maximization condition. Then we compare the retail profits evaluated at this retail price to the profits that the retailer would make if she kept the retail price unchanged at p_{jt-1}^r . The retailer will choose the price associated with the higher retail profits. Once the counterfactual retail price p_{jt}^{rc} has been determined this way, the associated counterfactual market share $s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r))$, $k \neq j$, can easily be evaluated.

Next consider inequality (20) that determines the lower bound on manufacturers' adjustment costs. Again, all variables in this inequality can be treated as known once steps 1-6 of Section 4.1 are completed, except for the counterfactual retail and wholesale prices, p_{jt}^{rc} and p_{jt}^{wc} respectively. To determine those, we consider two cases. In the first case the retail price changes from the previous period; the optimal prices p_{jt}^{wc} and p_{jt}^{rc} are then determined according to the manufacturer and retailer first-order conditions, equations (16) and (6) respectively, with the inverted manufacturer reaction matrix Ω_{wt}^{-1} reflecting the optimal pass-through of the wholesale price change onto the retail price. Let π_1^{wc} denote the manufacturer profits associated with the so-computed prices p_{jt}^{wc} and p_{jt}^{rc} . Next, consider the case in which the retail price does not change, even though the wholesale price does. As noted earlier, the optimal manufacturer pricing behavior in this case will involve changing the wholesale price up to the point where the retailer is just indifferent between changing her price and keeping it constant at p_{jt-1}^r . The optimal wholesale price will then be determined based on equation (21). Let π_2^{wc} denote the manufacturer profits associated with the prices p_{jt}^{wc} and p_{jt-1}^r in this case. If $\pi_1^{wc} > \pi_2^{wc}$, the manufacturer will set the wholesale price anticipating that the retailer will adjust her price too. Hence, the counterfactual prices p_{jt}^{wc} and p_{jt}^{rc} will satisfy the conditions described under the first case above. If $\pi_1^{wc} < \pi_2^{wc}$, the manufacturer will price the product anticipating that the retailer will not adjust her price. The resulting counterfactual wholesale price will then satisfy the indifference condition discussed under the second case, while the retail price will remain unchanged at p_{jt-1}^r . Once the counterfactual wholesale and retail prices have been determined, evaluation of the adjustment cost lower bound based on inequality (20) is straightforward.

G Optimization Routines Used To Compute Bounds on the Price-Adjustment Costs

To determine the upper bound on the retailer's price-adjustment costs in inequality (10), we calculate the counterfactual market share by plugging the previous period's price into the market share equation, so no additional optimization is necessary. To compute the retailer's lower bound as described in inequality (12), we used the *lsqnonlin* optimization routine in Matlab to compute the counterfactual price the retailer would have charged to satisfy its optimality conditions. The *lsqnonlin* command minimizes the sum of squares of a non-linear vector-valued function. In our model, this vector x is comprised of the prices that equilibrate supply and demand in each market. One can define a set of lower and upper bounds (lb and ub) on the variables in this vector, so that the solution is always in the range $lb \leq x \leq ub$, and *lsqnonlin* is run as a constrained optimization routine. To compute the upper and lower bounds on the manufacturers' price-adjustment costs as described in inequalities (18) and (20), respectively, we used the *lsqnonlin* optimization routine to compute both the retailer's response to potential counterfactual wholesale prices, as well as the optimal wholesale price in computing the lower bound, as we describe in more detail in Appendix F.