AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

Year 2010

Paper 1016R3

Outsourcing and Pass-Through

Rebecca Hellerstein

Federal Reserve Bank of New York

and

Sofia B. Villas-Boas

University of California at Berkeley

OUTSOURCING AND PASS-THROUGH

REBECCA HELLERSTEIN AND SOFIA B. VILLAS-BOAS*

February 26, 2010

Abstract

A large share of international trade occurs through intra-firm transactions. We show that this common cross-border organization of the firm has implications for the well-documented incomplete transmission of shocks across such borders. We present new evidence of an inverse relationship between a firm's outsourcing of inputs and its rate of exchange-rate pass-through. We then develop a structural econometric model with final assemblers and upstream parts suppliers to quantify how firms' organization of their activities across national borders affects their pass-through behavior.

Keywords: Exchange-rate pass-through; Intra-firm trade; Outsourcing; Vertical contracts.

JEL classifications: F14, F3, F4.

*Federal Reserve Bank of New York and University of California at Berkeley. Corresponding author: Rebecca Hellerstein, International Research Group, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045; email: Rebecca.Hellerstein@ny.frb.org. We thank the editor, Giancarlo Corsetti, two anonymous referees, Linda Goldberg, Pinelopi Goldberg, Gita Gopinath, Kala Krishna, Aviv Nevo, Maurice Obstfeld, Alexander Rose, Stephen Yeaple, Kei-Mu Yi and participants at the NBER Universities' Research Conference on "Firms and the Evolving Structure of Global Economic Activity," the CEPR's 2nd Workshop on Globalization and Contracts, and the Philadelphia Fed's 2006 Trade Workshop for their comments. This paper previously circulated as a working paper under the title "Arm's-Length Transactions as a Source of Incomplete Cross-Border Transmission: The Case of Autos." 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 of the Federal Reserve System.

1 Introduction

In recent years, the trade literature has produced several important new insights into the international organization of production. First, firms that engage in foreign trade are not a random sample of firms operating in the domestic economy but are larger, more productive, and more likely to be multinationals (See Helpman 2006, for a review). Second, over the half-century of globalization dominated by multinationals (Bordo, Eichengreen, and Irwin 1999), trade in intermediate inputs has risen dramatically (Hummels, Ishii, and Yi 2001). Third, over the past two decades, spurred by advances in computer-aided manufacturing, outsourcing has expanded both domestically and abroad.¹ A natural question, then, is, how do these stylized facts help us better understand some of the pricing puzzles in international macroeconomics, such as the incomplete transmission of exchange-rate shocks to domestic prices? This paper analyzes the extent to which how firms organize their production activities across borders affects their pass-through of exchange-rate shocks to their prices.

Although a substantial theoretical and empirical literature has made progress explaining *how* firms do not fully adjust their import prices following an exchange-rate shock, much less is understood about *why*. Yet assumptions about the *why* of this inertia shape economists' policy recommendations on some of the most basic issues in international goods and financial markets. Current research (most notably, Engel, 2002) points to three overarching sources of incomplete pass-through: the importance of local non-traded costs in the total costs of imported goods; markup adjustment by firms along the distribution chain; and the costs of nominal price adjustment.² Evidence on the relative importance and underlying drivers of these sources remains mixed as the most important variables identified by theory—such as markups, nontraded costs, or vertical contracts—generally remain unobservable in practice, particularly in aggregate data sets. Before macroeconomic models can grapple with the implications of each of these potential sources of incomplete transmission, they need stylized facts from the microeconomic literature about their relative importance and underlying drivers.

To identify the relative importance and underlying drivers of each potential source of incomplete transmission, a recent literature has drawn on new data sets for modeling the microeconomics of incomplete pass-through by exploiting the structural approach common to industrial organization (IO) models. The unique advantages of IO studies enable them to provide insights unattainable through other common approaches in the literature. First, IO studies use microeconomic data sets that generally have higher quality data than aggregate data sets, with more precise estimates of both prices and quantities. Although quantities are

¹Helpman (2006) reviews the relevant literature.

²See Burstein, Neves, and Rebelo (2003), Campa and Goldberg (2005), Corsetti and Dedola (2005), Corsetti, Dedola, and Leduc (2010), Gopinath, Itskhoki, and Rigobon (2010), Gopinath and Rigobon (2008), Gron and Swenson (2000), and Krugman (1987).

rarely available in aggregate data, they are needed to observe or derive variables of interest to pass-through studies, such as markups. Second, IO studies draw on industry lore to inform modeling assumptions, which, together with their higher-quality microeconomic data, can deliver credible estimates of markups, local costs, and the like. Third, regarding the sources of markup adjustment, IO studies provide possible mechanisms for formalizing the reasons that price does not equal costs and thus for why changes in costs may not be passed through into prices. These mechanisms may include long-term contracts (Bettendorf and Verboven 2000) or the possibility of strategic mark-up adjustment along the supply distribution chain (Goldberg and Verboven 2001; Bonnet, Dubois, and Villas-Boas 2009). Finally, IO studies generally use structural models that allow counterfactual experiments for testing the various hypotheses about why pass-through is low and thus enable researchers to ask deeper questions than are possible in modeling approaches that rely exclusively on aggregate data. Examples include Hellerstein (2008), Goldberg and Hellerstein (2008), and Nakamura and Zerom (2009), which all use structural models to decompose the relative importance of various sources of incomplete pass-through. This paper, based on the IO approach, analyzes the impact of vertical relationships on firms' markup adjustment – and the consequent incomplete pass-through. The disadvantage of IO studies is that to be generalizable, the results require confirmation from work on other industries. That said, the literature so far has begun to identify stylized microeconomic facts that are remarkably robust across markets and modeling assumptions.

Our first goal in this paper is to present evidence of a positive relationship between a firm's degree of vertical integration and its pass-through of foreign cost shocks. The second goal is to look more deeply at the sources of incomplete transmission by using a structural model to run counterfactual experiments which simulate variations in firm boundaries across national borders. For the structural model, our empirical approach has two components: estimation and simulation of counterfactuals. At the estimation stage, we estimate the demand parameters and then the marginal costs and markups of the original equipment manufacturer (OEM) and what in industry parlance is known as a tier-one parts supplier (TOPS) for each vertical contractual model. To assess the overall impact of each vertical contract on firms' pass-through behavior, we employ simulation.

A cross-border contract is by definition a vertical contract between an upstream (foreign) and a downstream (domestic) firm. Our structural model studies a continuum of contracts with two extreme points of integration: vertically integrated and outsourced supply chains. In a vertically integrated supply chain, all transactions are conducted in-house, within the boundaries of the firm. In an outsourced supply chain, all transactions are conducted at arm's length between firms with no formal ownership ties or control rights over one another's production and pricing decisions. In the structural model, we begin by computing the industry equilibrium that emerges if firms experience a 10-percent cost increase in the outsourced model

and compare it to the equilibrium that prevails when one firm, a multinational, controls pricing along the distribution chain. We interpret the different response of prices across the two models as a measure of the overall impact of firms' cross-border organization on their pass-through of shocks. We find that, on average, following a 10-percent cost increase, firms pass through 54 percent of a foreign-cost shock to their retail prices in a vertically integrated supply chain and only 13 percent in an outsourced supply chain, a 40 percentage-point difference. Next, we compute pass-through with a structural model calibrated to the average degree of vertical integration observed in the auto industry. We then compare the results to the outcome given a one-standard-deviation increase in the share of outsourced production. We find that a one-standard-deviation increase in outsourcing is associated with a 12 percentage-point decline in firms' average pass-through to prices. The counterfactual experiments thus confirm that pass-through is much higher in a vertically integrated than in an outsourced production chain.

The empirical pass-through literature generally attributes some part of local-currency price inertia to firms' markup adjustment but without analyzing its drivers, such as the role of different vertical contracts between suppliers and final-goods producers. An early paper by Rangan, Lawrence, and Cooper (1993) argues that U.S. multinationals pass-through exchange-rate shocks to export prices at a higher rate than do vertically separated firms.³ Several more recent papers, including Clausing (2003) and Neiman (2010), examine the empirical relationship between the boundaries of the firm and the transmission of shocks across national borders. Using the microeconomic price data used to construct the U.S. Bureau of Labor Statistics' international price indexes, Clausing (2003) and Neiman (2010) estimate pass-through elasticities for intrafirm and market transactions for U.S. imports and exports. Clausing (2003) finds a higher rate of pass-through for U.S. imports traded in market transactions than in intrafirm transactions, while Neiman (2010) finds a higher rate of pass-through for differentiated goods traded in intrafirm transactions than for those traded in market transactions, a very similar result to our findings. Our study is different from this previous work in several important respects. Although the microdata used in previous work encompass a number of industries, the intrafirm data are transfer prices, whose variation may not fully reflect variation in transaction prices because of firms' incentives to transfer income across borders to minimize taxes, a fact established by Clausing (2003).⁴ While our micro-data cover only the auto industry, we are able to use quantity data to derive intrafirm markups. We do not use transfer prices

³Bernard, Jensen, and Schott (2006) find in U.S. census data that multinationals adjust arm's-length prices by more than intrafirm prices after an exchange-rate shock, though in a very limited sample.

⁴Roughly half of the intra-firm prices included in the BLS's import price samples are cost-based rather than market-based: They are not set with reference to market trends but instead are constructed as a markup over the exporting subsidiary's observable costs. A study by Alterman (1997) finds that, for particular industries, cost-based prices change less frequently than market-based prices and, in particular, fluctuate less than market-based prices following macroeconomic shocks such as exchange-rate changes.

in our structural analysis.

We study the auto industry in our structural analysis for several reasons. First, because prices of manufactured goods tend to exhibit dampened responses to foreign shocks in aggregate data, autos are an appropriate choice to investigate the puzzling phenomenon of incomplete pass-through. Second, trade in autos and auto parts is quite large for most countries, which gives our empirical results direct policy relevance. For example, trade in autos and auto parts makes up 25 to 40 percent of U.S. goods imports in any given year, depending on how one concords the industry’s input-output tables with import data. Third, we have a rich panel data set with monthly transaction prices for 24 models from six major manufacturers over a period of 37 months, as well as the share of inputs sourced from abroad by model. We focus on the top-selling models and thus have a limited number of car models in our sample, each with a very large U.S. market share. This approach enables us to use product fixed effects to control for product-specific, time-invariant characteristics in our structural demand model, as we discuss further in the estimation section. These 24 models comprise roughly one-third of U.S. sales of passenger cars and sports utility vehicles (SUVs). Our sample includes models assembled in the United States with a significant share of parts sourced from abroad.

The next section outlines our conceptual framework, while section 3 presents stylized facts about the relationship between exchange-rate pass-through and intrafirm trade. Section 4 sets out our structural model, and section 5 discusses the estimation methodology and demand results. Section 6 presents results from the counterfactual experiments, and section 7 concludes.

2 Vertical Integration and Pass-Through: Conceptual Framework

The goal of this paper is to investigate the relationship between a firm’s degree of vertical integration and its pass-through of exchange-rate-induced cost shocks. While aware that a firm’s pass-through may relate to other factors previously identified in the literature, we highlight the role of its vertical contracts. These contracts, which could affect its pass-through behavior for many reasons, are especially difficult to examine empirically owing to their inframarginal components, transaction costs, and imperfect information issues. Our conceptual framework sidesteps these issues by focusing on how these contracts affect the traditional problem of double marginalization, drawing on the IO literature and its focus on the inefficiencies associated with different vertical contracts (see Villas-Boas 2007, for a review). Vertical contracts determine the vertical profit and so the size of the total producer surplus to be divided among firms along a supply chain. They also promote efficiency in the supply chain if they depart from

the linear pricing that leads to double marginalization.⁵ The vertical contracts in a particular market can thus significantly affect downstream price dynamics (Hastings 2004; Villas-Boas 2007).

Before we set out our conceptual framework, we define several terms used throughout the paper. Intrafirm trade refers to cross-border shipments between U.S. firms and their foreign subsidiaries and between foreign firms and their U.S. subsidiaries. A vertically integrated firm refers to a firm that controls pricing and production decisions at multiple points along a supply chain.⁶ Our conceptual framework focuses on whether firms are vertically integrated upstream. An upstream vertically integrated firm can be thought of as a final-goods producer that buys an upstream supplier from which it previously purchased goods in market transactions. In our structural model, a multinational refers to a single firm that controls production and pricing decisions along the supply chain and sources some of its inputs from abroad, while a nonmultinational refers to several firms that control production and pricing decisions along the supply chain but also source some of their inputs from abroad.

We begin our analysis by using traditional pass-through analyses (as in Feenstra 1989; Goldberg and Knetter 1997; Campa and Goldberg 2005) to establish a positive relationship between an industry's share of intrafirm imports and its pass-through of exchange-rate shocks to import prices. We find a similar relationship holds when we focus on the auto industry – the higher the degree of vertical integration of a firm, measured by the ratio of its total value added to total revenue, the larger its pass-through of exchange-rate shocks. While our regression analysis is indicative of the phenomenon, making counterfactual predictions of firms' out-of-sample vertical integration based on their pass-through behavior may be problematic. The regression analysis may suffer from selection problems as firms' degree of vertical integration is chosen, not randomly assigned. With these facts in mind, we turn to the structural analysis.

Our explanation for why multinationals exhibit higher pass-through than do outsourced firms is fairly intuitive. When multinational firms are hit with an exchange-rate shock, they have only one layer of pricing decisions as the OEM and the TOPS are within the same firm, so that only one markup adjustment occurs after the shock. In contrast, in the outsourced case when vertically separated firms experience an exchange-rate shock two layers of strategic markup adjustment may occur. Not only can the OEM adjust its markup but also the TOPS. Pass-through cannot be higher than for the case of the multinational firm, and we posit that pass-through is lower the more firm boundaries a product crosses to reach the final consumer.

⁵Double marginalization arises when two independent markups are set along a supply chain and when the only contractual instrument used is the wholesale price. As a result, the sum of profits for a supplier and manufacturer may be less than their profits had they coordinated their pricing.

⁶For example, a vertically integrated supplier that functions as a division of a final-goods producer has no control rights. It does not decide the price or quantity of output it produces, while an outsourced or vertically separated supplier does possess these control rights.

Several key assumptions underlie our approach. First, our structural model assumes that each OEM sources from both domestic- and foreign-parts suppliers and has the same vertical relationship with both suppliers. A multinational is thus assumed to be vertically integrated with its domestic- as well as its foreign-parts suppliers, while a final-goods producer that outsources its foreign-parts production is assumed to do the same for its domestic-parts production. It is possible, however, that final-goods producers outsource the production of more foreign parts than domestic parts, or vice versa. The aggregate evidence shows that intrafirm transactions are more common in cross-border than in domestic trade, implying that our estimates of multinational pass-through are overstated. If a multinational outsources its domestic- but not its foreign-parts production, its composite supplier markup will be higher than if it owned both parts of the production chain, when it would be zero. In that case, some double marginalization would be introduced into the supply chain and so decrease the pass-through of the markup to the retail price. Second, we assume linear pricing by firms which creates two sources of mark-up adjustment and double marginalization in the outsourced scenario. We follow other papers in the structural IO literature in assuming linear pricing for automobiles (Berry, Levinsohn, and Pakes 1995; Goldberg and Verboven 2001; Brenkers and Verboven 2006). The second markup can be eliminated in many ways including through vertical integration or nonlinear pricing.⁷ Third, we assume for simplicity that retail automobile dealers have a constant markup and thus complete pass-through. This assumption lets us focus on the upstream strategic interactions between TOPSs and OEMs and, given the empirical patterns in the data, appears to be a reasonable way to model the retailer's pass-through behavior. In a regression analysis, we cannot reject full pass-through of changes in wholesale prices to retail prices for almost every model in our sample. Fourth, our approach in the structural model is static. While estimating a fully dynamic model would be desirable, the complexity of our structural model does not allow us to do so without simplifying aspects that we consider essential to our analysis, such as the strategic interactions between TOPSs and OEMs and the flexibility of the demand system. A primary source of the auto industry's cost variation and a focus of our analysis is the exchange rate, which is highly persistent, a stylized fact that implies that our premise of static expectations is reasonable. As Goldberg and Hellerstein

⁷We acknowledge that considering only linear pricing in the outsourced case gives us a downwardly biased estimate of pass-through. In our structural model foreign parts are combined with domestic parts in fixed proportions to form the composite supplier good for which a wholesale price is set. An alternative model might consider linear pricing for foreign suppliers but nonlinear pricing for domestic suppliers (or vice versa). This model would reduce the domestic supplier's markup and so the scope for the composite supplier's markup to adjust following a foreign cost shock. Introducing nonlinear pricing into the outsourced model would thus increase pass-through, though not to as much as in the multinational case, as some parts would still be priced linearly. Pass-through would fall between the two extreme cases of fully in-house or outsourced production of all domestic and foreign parts. In future research, we plan to examine the impact of nonlinear pricing on pass-through in vertically separated supply chains, as these would allow multiple layers of mark-ups to be contracted away and so result in higher pass-through.

(2008) set out in more detail, if the exchange rate is a primary source of cost variation and it is highly persistent, our model provides a reasonable approximation to a dynamic setup.⁸ Our approach is also standard in the IO literature and has been used in many previous studies of the auto industry (Berry, Levinsohn, and Pakes 1995; Goldberg and Verboven 2001; Brenkers and Verboven 2006).

The final step in setting out our conceptual framework is to explore possible alternative explanations for our finding of higher pass-through by multinationals. We do so in the next section after setting out the results from the industry regressions.

3 Vertical Integration and Pass-Through: Evidence of a Positive Relationship

Exchange-rate pass-through is defined as the percentage change in an industry's prices for a given percentage change in its marginal costs caused by movements in the dollar. Estimating pass-through is not a simple exercise. The sensitivity of import prices to dollar movements may differ from a simple correlation between the two variables because of independent activity in the production or demand sectors. To estimate pass-through, one must control for other forces that affect firms' choices of import prices, such as demand conditions in the importing country.

We use a standard workhorse model to estimate exchange-rate pass-through elasticities across industries and for individual automakers and models. This section describes the specification used to estimate the industry pass-through coefficients.⁹ Our pricing equation is

$$p_t = \alpha + \sum_{i=0}^4 \beta_i e_{t-i} w_{t-i} + \gamma Y_t + \varepsilon_t$$

where p_t is an industry-specific index of U.S. import prices at time t , α is a constant, e_{t-i} is the import-weighted nominal exchange rate at time t minus i , w_{t-j} is a measure of import-weighted marginal costs at time t minus j , Y_t is a control for demand shifts that may affect import prices independently of the exchange rate at time t , and ε_t is a stochastic error term that is normally distributed. Each of the variables is in percentage-change terms. Exchange-rate pass-through is defined as the sum of the coefficients on the nominal import-weighted exchange rate e interacted with the marginal-cost term w at time t plus four lagged periods. Changes in the demand for imports that reflect variation in consumer tastes or income rather than in

⁸Dynamic considerations may affect the analysis in several ways: first, to the extent that consumers or firms hold inventories of autos, the demand- and supply-side parameter estimates obtained by the static approach may be biased (Hendel and Nevo 2006); and, second, dynamics may enter through intertemporal effects on the demand side, stemming, for example, from habit formation (Ravn, Schmitt-Grohe, and Uribe 2009).

⁹Similar specifications are used in Feenstra (1989), Goldberg and Knetter (1997), and Campa and Goldberg (2005).

the dollar’s value are controlled for by including U.S. domestic demand in each regression.

The dependent variable in the regression is the quarterly industry import-price index for each NAICS’s three-digit manufacturing industry from 1985 to 2004. The industry exchange rate is the import-weighted exchange rate for each industry with weights derived from the average share of each country’s imports in the industry’s total imports from 1987 to 2004. Foreign production costs proxy for marginal-cost shocks other than exchange-rate fluctuations that may affect import prices. An import-weighted foreign consumer price index (CPI) proxies for foreign production costs w , with the same weights as those for the exchange rate. We use this pass-through specification with the marginal-cost and exchange-rate variables interacted to be consistent with the set-up of the structural model, which focuses on marginal-cost shocks. In robustness checks, we generally cannot reject that the coefficients on the marginal-cost and exchange-rates variables are identical.¹⁰

Table 1 reports the regression results by industry. Pass-through varies significantly across industries, from statistically indistinguishable from zero for Food and Beverages and Leather Goods to over 50 percent for Primary and Fabricated Metals, Industrial Machinery, and Electrical Equipment. It averages 35 percent across industries with a standard deviation of 20 percent. Figure 1 illustrates the positive relationship between an industry’s share of intrafirm imports on the x-axis and its pass-through elasticity on the y-axis. The industry shares of intrafirm imports are for the year 2000 and come from table 2 of Bernard et al (2008).¹¹ The relationship between the two variables is clearly positive as illustrated by the upward-sloping linear trend line. Table 2 reports the results from regressing the industry pass-through coefficients on the industry shares of intrafirm imports. The coefficients have a natural interpretation: a 10-percent increase in an industry’s share of intrafirm imports is associated with a roughly 6-percent increase in its pass-through to import prices.¹² The R^2 for the regression is 0.33, which implies a reasonable fit, particularly given our limited observations.

¹⁰U.S. domestic demand is defined as U.S. total domestic output (GDP) minus exports (demand from outside the U.S.) plus imports (U.S. demand not satisfied by domestic output) and comes from the U.S. Bureau of Economic Analysis. The exchange rate and CPI data come from IMF’s International Financial Statistics. The trade data come from the U.S. International Trade Commission and the U.S. Bureau of Economic Analysis and the industry import-price indexes from the U.S. Bureau of Labor Statistics. Following Antras (2003), we exclude from our analysis industries in which one country has a predominant share of U.S. imports, as in those cases it is difficult to distinguish between the effects of national origin and those of contractual form for any observed differences in pass-through. We exclude two 3-digit NAICS manufacturing industries, Wood Products and Paper Products (NAICS 321 and 322, respectively) because they each source 70-75 percent of their imports from one country, Canada.

¹¹Roughly 40 percent of U.S. imports occur via intrafirm transactions, a fact first established by Zeile (1997).

¹²We consider industry regressions with different exchange rates in appendix A.

3.1 Alternative Explanations from Prior Theoretical and Empirical Work

Is this pattern explained by omitted factors? Factors that cause or are correlated with intrafirm trade may independently cause higher pass-through by multinationals and so provide possible alternative explanations for our results. We derive several testable hypotheses for alternative explanations from the growing theoretical and empirical literature on intrafirm trade.

Capital intensity Antras (2003), Bernard et al (2008), and Nunn and Treffer (2008) find that an industry’s share of intrafirm imports has a positive relationship to its physical capital intensity. If greater capital intensity is associated with greater market power, industries with greater capital intensity may have more ability to pass-through cost shocks. To test this hypothesis, we include in our regression the same measure of industry-level capital intensity as in Antras (2003). This is the industry’s capital stock and comes from the NBER-CES Manufacturing Industry Database for 1996. Table 2 shows that the relationship between an industry’s pass-through and its average intrafirm share is robust to controlling for its capital intensity. The coefficient on the intrafirm share variable does not change significantly with the addition of this extra variable, and the coefficient on the capital-intensity variable is not significant.

Product differentiation One might also expect firms in industries with more differentiated products to exhibit higher pass-through because their more specialized products would afford them greater market power and hence a greater ability to pass through marginal-cost shocks to prices. If industries with significant intrafirm trade produce more differentiated products, this omitted variable may explain the relationship between pass-through and intrafirm trade. Along the same lines, Bernard et al. (2008) find a positive relationship between an industry’s skill intensity and its intrafirm trade shares. One would expect more skill-intensive industries to produce more differentiated products and so potentially to exhibit higher pass-through. To test this hypothesis, we define an industry’s degree of differentiation as its import-weighted share of differentiated products, where differentiated products are identified by the Rauch (1999) classification. As reported in table 2, we find that this variable is not significant in the regression and has no effect on the coefficient on the intrafirm-share variable.¹³

Market share of foreign producers A third testable alternative explanation comes from the pass-through literature. One may expect foreign producers with a larger share of the domestic market to have greater market power and therefore more ability to pass through marginal-cost shocks to prices, a relationship explored by Feenstra, Gagnon, and Knetter (1996). If industries with high shares of intrafirm trade also have substantial market share held by foreign producers, the latter may be an omitted variable that explains the relationship

¹³This result is robust to using Rauch’s (1999) conservative or liberal definition of product differentiation.

between pass-through and intrafirm trade. We examine whether our results can be explained by the average market share of foreign relative to domestic producers over the sample period. The market share of foreign producers is the mean share of imports in the domestic market (comprising imports plus domestic shipments) from 1987 to 2004. As we report in table 2, this variable is not significant, and controlling for it has little effect on the coefficient on the intrafirm-share variable.

To summarize, we identify a positive relationship between an industry’s share of intrafirm trade and its pass-through of marginal-cost shocks to import prices. As we discuss in more detail below, we hypothesize that this is a firm-level story. Multinationals are more likely to pass through a given shock because they have a lower vertical markup and so less scope for adjustment following a marginal-cost shock. Because our industry evidence is nevertheless indirect, we need to see what we can observe with data on firms. To this end, we next present evidence that supports our hypothesis from one industry, the auto industry.

3.2 Outsourcing and Pass-through in the Auto Industry

3.2.1 Data

Our price data for the auto industry come from Edmunds.com, an industry data provider. The Edmunds data contain a number of variables for individual models on a monthly basis from October 2002 to June 2006. These include the national base total market value (TMV) price for both new and used models, the manufacturer’s suggested retail price (MSRP) for new models, and data on characteristics such as horsepower, length, and sourcing. These price data build on previous pass-through studies by enabling us to estimate pass-through coefficients on the make and model, rather than just on the market segment. And unlike the prices used in some previous pass-through studies, the Edmunds TMV prices are transaction prices calculated from a sample of roughly 20–30 percent of all U.S. monthly auto sales. The TMV price is the median price paid nationally for the base model after individual transactions are weighted to represent regional trends correctly.

Busse, Silva-Risso, and Zettelmeyer (2006) note that most of the variation in retail automobile prices comes from the application of dealer and consumer promotions to the MSRP, a well established fact in industry lore and one frequently cited by practitioners, in industry publications, and the like. The most common promotions are consumer and dealer cash rebates, the latter offered by manufacturers to dealers for each auto sold during a given promotion and frequently passed through by dealers to the retail price. This stylized fact about pricing in the auto industry is precisely why we use the Edmunds TMV price in our analysis rather than the MSRP. An unusual and appealing feature of the Edmunds TMV retail prices is that they include the effects of both types of promotions. As a result, the TMV prices tend to

be several thousand dollars less than the MSRP on average and exhibit much more monthly variability than the MSRP, which generally changes at most several times over the model year in a step-like fashion. In this sense, our data are much more appropriate for a pass-through analysis than the MSRPs used in previous studies, which do not capture either the level of or variation in the true retail price over the model year. Figure 2 illustrates these patterns and compares the MSRP and the Edmunds TMV price for one model, the Ford Escape.

The Edmunds data include 24 models, each with a substantial U.S. market share. Together, they comprise slightly more than one-third of total retail sales of passenger cars and SUVs in the United States. Our sample includes top-selling models which enables us to use an identification strategy in our structural demand model that is not feasible with a broader sample.

Our monthly sales quantity data come from Ward’s Automotive. Our measure of vertical integration is each automaker’s total value added relative to its total sales and comes from Lieberman and Dhawan (2005). *Value-added* is defined as each firm’s sales less the costs of purchased materials and services and equals all payments to labor and capital.¹⁴ This value-added measure can be interpreted in a fairly straightforward fashion as a measure of vertical integration—as payments to primary factors of production directly employed by the firm relative to total payments by the firm to factors of production employed by the firm and purchased indirectly through intermediate inputs (through purchased goods and services). A key assumption underlying our use of value added as a measure of vertical integration is that firms’ coordination with their internal “suppliers” is greater than that with their external suppliers: that is, *de jure* vertical integration, measured in terms of formal ownership ties, equals *de facto* vertical integration, measured by coordination in setting markups along the distribution chain. We motivate our use of a firm-specific measure from evidence that firms exhibit very different degrees of vertical integration that persist over time, as Lieberman and Dhawan (2005) document. Some producers are historically associated with very high levels of vertical integration (e.g., GM) while others are known to outsource much more (e.g., Honda). No model-specific measures of vertical integration are available to us.¹⁵

Finally, our study includes a measure of the share of foreign content in each model’s total

¹⁴As Lieberman and Dhawan (2005) note, the measures of value added come from the Daiwa Securities Corporation’s measures of value added for Japanese firms and for U.S. firms from summing payments to capital and labor from their annual reports. Note that in the literature one finds two different definitions of value-added. One definition is the difference between the sale price and the cost of labor employed directly by the firm as well as the cost of purchased materials and services to produce a good. A second definition is the difference between the sale price and only the cost of purchased materials and services to produce a good. The latter equals the contribution of (and so income to) primary factors of production (labor and capital) to the value of the good. The measure used in the paper includes payments to labor employed directly by the firm, and so is not equivalent to a measure of accounting profitability.

¹⁵Note that our measure of vertical integration for Daimler-Chrysler comes from the Chrysler operations of the firm as our sample contains only Chrysler models.

content, enabling us to identify the role of local costs and to separate its role in any incomplete pass-through from that of inefficiencies caused by firms’ contractual arrangements. Each auto manufacturer is required by law to report the share of foreign content in each of its model’s total content.¹⁶ Table 3 reports summary statistics for prices, quantities, characteristics, and share of foreign content. Table 4 reports the sales-weighted share of foreign content by maker for the models in our sample, which ranges from 18 percent for Ford to 57 percent for Toyota.

3.3 Pass-through Analysis

Our model to estimate exchange-rate pass-through for individual auto models is:

$$p_t = \alpha + \beta^f f e_t w_t^f + \beta^d (1 - f) w_t^d + \gamma Y_t + \varepsilon_t$$

where p_t is the *Edmunds* TMV price at time t , α is a constant, e_t is the nominal exchange rate at time t , w_t^f is a measure of foreign marginal costs at time t , w_t^d is a measure of domestic marginal costs at time t , f is the model’s share of foreign content, Y_t is a control for demand shifts that may affect import prices independently of the exchange rate at time t , and ε_t is a stochastic error term that is normally distributed. The variables are expressed as logs. Exchange-rate pass-through is the coefficient on the nominal import-weighted exchange rate e interacted with the foreign marginal cost term w^f .¹⁷ We measure foreign-cost shocks by the producer price index (PPI) for the manufacturing sector in the relevant country, using the PPI for the auto industry where available. The nominal exchange rate and the foreign marginal-cost measure used in each regression depend on the sourcing information in the Edmunds data. Domestic costs are measured by the PPI for U.S. automobile manufacturing and come from the U.S. Bureau of Labor Statistics. Changes in the demand for imports that reflect variation in consumer tastes or income rather than in the dollar’s value are controlled for by including U.S. domestic demand in the regression. We examine exchange-rate pass-through over the course of one month to be consistent with the set-up of the structural model.¹⁸ Finally, we interact the domestic-cost variable with the domestic share of each model and the foreign-cost variable with the foreign share.

Table 5 reports the regression results by model. Pass-through varies significantly across

¹⁶These data come from the Automotive Trade Policy Council and the U.S. Department of Commerce and are for the year 2004. Every model in our data is assembled in the United States but has some share of its parts sourced from abroad, from 5 percent to 85 percent, with a mean of 25 percent across models and a standard deviation of 17 percent.

¹⁷In robustness checks, we generally can not reject that the coefficients on the foreign marginal-cost and exchange-rates variables are identical.

¹⁸In the industry regressions, we follow the norm in the empirical pass-through literature and estimate pass-through coefficients over the course of one year. We do not have sufficient degrees of freedom in the auto data set for an annual study.

models, from statistically indistinguishable from zero for the Ford Explorer, the Honda Accord, and the Honda Odyssey to 90 percent for the Chrysler PT Cruiser. It averages 38 percent across models, with a standard deviation of 25 percent. Figure 3 illustrates the positive relationship between an automaker’s degree of vertical integration on the x-axis and its average pass-through across models on the y-axis. The relationship between the two variables is clearly positive as illustrated by the upward-sloping linear trend line. Note that by using a measure of a firm’s vertical integration, we implicitly assume that the firm’s average degree of vertical integration affects its pricing decisions across model lines.¹⁹ Overall, this section illustrates the fact that cross-border pass-through is positively related to firms’ degree of vertical integration. We use this stylized fact to motivate the exercises conducted with the structural model.

4 The Structural Model

This section describes the supply model used in the structural analysis and derives simple expressions to compute pass-through coefficients. It then sets out the random-coefficients model used to estimate demand.

4.1 Supply

Consider two vertical supply models, an outsourced model in which all upstream production is outsourced, resulting in double marginalization, as in Brenkers and Verboven (2006), Villas-Boas and Hellerstein (2006), Villas-Boas (2007), and Hellerstein (2008) and a multinational model in which all upstream production is done in-house, resulting in a supply model in which only the manufacturer has a markup.

Consider a standard linear-pricing model in which tier-one parts suppliers, acting as Bertrand oligopolists with differentiated products, set their prices followed by original equipment manufacturers who set their prices taking the price they observe as given. Thus, a double markup is added to the marginal cost to produce the composite parts that make up the product. For simplicity, the retailer is assumed to have constant markup, and thus to have complete pass-through. Strategic interactions between TOPSs and OEMs with respect to prices follow a sequential Nash-Bertrand model. To solve the model, one uses backward induction.

¹⁹One possible alternative explanation is that our results are explained by the behavior of the exchange-rate over the sample period, as the firms with the lowest degree of vertical integration in our sample are Japanese. However, the direction and magnitude of changes in the yen-dollar relationship over the sample period do not differ significantly from those of other relevant exchange rates for the auto industry. Figure B.1 in appendix B shows that over the sample period the Japanese yen tended to move together with other major exchange-rate indexes against the dollar.

4.1.1 Retailers

Consider R retail firms that each sell some share κ^r of the market's J differentiated products. For simplicity, assume they have the same ownership structure as manufacturers. We cannot reject full pass-through of changes in manufacturer prices to retail prices for almost every model in our sample. We therefore assume a constant markup by retailers, implying that any changes to manufacturer prices are completely passed through to retail prices. This assumption lets us focus on the upstream strategic interactions between intermediate goods suppliers and manufacturers and given the empirical patterns in the data, appears to be a reasonable way to model the retailer's pass-through behavior.

4.1.2 Original Equipment Manufacturers

Consider A original equipment manufacturers that each sell a subset κ^a of the market's J differentiated products. Let all firms use linear pricing and face constant marginal costs. The profits of an auto manufacturer firm in market t are given by

$$\Pi_{jt}^a = \sum_{j \in \kappa^a} (p_{jt}^a - p_{jt}^s - ntc_{jt}^a) s_{jt}(p_t^r(p_t^a)) \quad (1)$$

where p_{jt}^a is the price the manufacturer sets for make j , p_{jt}^s is the price paid by the manufacturer to the tier-one suppliers for the composite of parts necessary to make model j , ntc_{jt}^a are local nontraded costs paid by the manufacturer to sell model j , and $s_{jt}(p_t^r(p_t^a))$ is the quantity demanded (or market share) of product j which is a function of the retail prices of all J products, p_t^r that in turn are functions of p_t^a . Assuming that each manufacturer acts as a Nash-Bertrand profit maximizer, the model price p_{jt}^a must satisfy the first-order profit-maximizing conditions

$$s_{jt} + \sum_{k \in \kappa^a} (p_{kt}^a - p_{kt}^s - ntc_{kt}^a) \frac{\partial s_{kt}}{\partial p_{jt}^a} = 0, \text{ for } j = 1, 2, \dots, J_t. \quad (2)$$

This gives us a set of J equations, one for each product. One can solve for the manufacturer markups by defining $S_{jk}^a = \frac{\partial s_{kt}(p_t^r(p_t^a))}{\partial p_{jt}^a}$ $j, k = 1, \dots, J$, as the matrix of demand substitution patterns, the marginal change in the k th product's market share given a change in the j th product's manufacturer price, and a $J \times J$ matrix T^a with the (j th, k th) element equal to 1 if both products j and k are produced by the same manufacturer, and equal to zero otherwise. If we define “ $*$ ” to denote the element by element multiplication of two matrices then the stacked first-order conditions can be rewritten in vector notation and inverted together in each market to get the manufacturer's pricing equation

$$p_t^a = p_t^s + ntc_t^s - (S_t^a * T^a)^{-1} s_t \quad (3)$$

where the manufacturer price for product j in market t will be the sum of its composite (fixed-proportions) supplier price p_t^s , composite destination-market nontraded costs ntc_t^s incurred by the supplier to produce and sell product j , and a markup.

4.1.3 Tier-One Parts Suppliers

Let there be S tier-one parts suppliers that each produce some subset Γ^s of the market's J_t differentiated products. We assume each supplier's product is combined into a composite supply product in fixed proportions over time. Each supplier chooses its price p_{jt}^s while assuming each manufacturer behaves according to its first-order condition (2). Supplier s 's profit function is:

$$\Pi_t^s = \sum_{j \in \Gamma^s} (p_{jt}^s - tc_{jt}^s - ntc_{jt}^s) s_{jt} (p_t^r(p_t^a(p_t^s))) \quad (4)$$

where tc_{jt}^s are traded costs and ntc_{jt}^s are destination-market nontraded costs incurred by the supplier to produce and sell product j . Multiproduct suppliers are represented by a ownership matrix, T^s , with elements $T^s(j, k) = 1$ if both products j and k are produced by the same supplier and zero otherwise. Consistent with industry lore, these ownership matrices are assumed to be identical to those of the auto manufacturers, as we discuss in more detail below. Assuming that a Bertrand-Nash equilibrium in prices prevails and that all suppliers act as profit maximizers, the equipment price p_{jt}^s must satisfy the first-order profit-maximizing conditions:

$$s_{jt} + \sum_{k \in \Gamma^s} T^s(k, j) (p_{kt}^s - tc_{kt}^s - ntc_{kt}^s) \frac{\partial s_{kt}}{\partial p_{jt}^s} = 0 \text{ for } j = 1, 2, \dots, J_t. \quad (5)$$

This gives us another set of J equations, one for each product. Let S_t^s be the matrix with element $\frac{\partial s_{kt}(p_t^r(p_t^a(p_t^s)))}{\partial p_{jt}^s}$, which is the change in each product's market share with respect to a change in each product's composite supplier price.²⁰ The suppliers' marginal costs are then recovered by inverting the element by element multiplication of $S_t^s * T^s$ for each market t , in vector notation. Therefore

$$p_t^s = tc_t^s + ntc_t^s - (S_t^s * T^s)^{-1} s_t \quad (6)$$

where for product j in market t the supplier price is the sum of the supplier traded costs, nontraded costs, and markup function.²¹ The supplier of product j can use its estimate of

²⁰This matrix is defined in appendix C.

²¹Equations (3) and (6) are similar to the pricing equations that produce the downstream and upstream margins of the linear-pricing models in Brenkers and Verboven (2006), Villas-Boas and Hellerstein (2006), Villas-Boas (2007), and Hellerstein (2008).

the manufacturer's reaction function to compute how a change in the supplier price will affect the manufacturer price for its product. Suppliers can assess the impact on the vertical profit, the size of the pie, and its share of the pie by considering the manufacturer reaction function before choosing a price. Suppliers may also act strategically with respect to one another. The manufacturer mediates these interactions by its pass-through of a given supplier's price change to the product's manufacturer (and so retail) price. Suppliers set prices after considering the manufacturer's pass-through of any supplier price changes to the manufacturer price, and other suppliers', manufacturers', and consumers' reactions to any changes in manufacturer prices.

4.2 Computing Pass-through Coefficients

To recover pass-through coefficients we simulate the effect of a shock to foreign suppliers' marginal costs on all manufacturers' prices by computing a new Bertrand-Nash equilibrium. Suppose a shock hits the traded component of the j th supplier's marginal cost. To compute the composite supplier's pass-through of that shock, we substitute the new vector of traded marginal costs, tc_t^{s*} , into the system of J nonlinear equations that characterize supplier pricing behavior and then search for the composite supplier price vector p_t^{s*} that will solve the system in each market t

$$p_{jt}^{s*} = tc_{jt}^{s*} + ntc_{jt}^s - \sum_{k \in \Gamma^s} (S_t^s * T^s)^{-1} s_{kt} \text{ for } j = 1, 2, \dots, J_t. \quad (7)$$

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

$$p_{jt}^{a*} = p_{jt}^{s*} + ntc_{jt}^a - \sum_{k \in \kappa^a} (S_t^a * T^a)^{-1} s_{kt} \text{ for } j, k = 1, 2, \dots, J_t. \quad (8)$$

Pass-through to the retail price is assumed to be complete. The degree of vertical integration is expressed through the ownership matrix. If a supplier and manufacturer are completely vertically integrated, then the T^s 's will be equal to zero and only the manufacturer will have a markup. If they operate in an outsourced relationship, then each diagonal entry of T^s 's will be a one. Between these two extremes, each entry of T^s will be calibrated to the average degree of integration between that OEM and its suppliers.

4.3 Discussion

Several assumptions used in the supply model merit further discussion. First, the model assumes that tier-one suppliers have the same ownership structure as OEMs. This assumption is motivated by the very large tier-one suppliers of which there are roughly 25 and which tend to have strong ties to one OEM. In addition, OEMs will commonly use contractual provisions to prevent tier-one suppliers from working with other OEMs. Ben-Shahar and White (2007) describe how OEMs will often pay for investments in machinery to be used to produce a particular set of parts. When they do, they insist on contractual provisions that prevent the supplier from serving multiple clients, that is, other OEMs, with that machinery. A common complaint of tier-one suppliers is about the wasteful exclusivity this type of provision. That said, this assumption also serves as a useful simplification of the analysis as there are many suppliers to each OEM and some do supply more than one. Our model of upstream relationships in auto manufacturing is necessarily stylized, as it would be almost impossible to map exactly the contractual relationships between OEMs and their many suppliers. Changing the model's ownership structure to allow tier-one suppliers to supply multiple OEMs has only a second-order effect on the pass-through coefficients, however.

Second, the supply model assumes that firms adjust their prices monthly, and that upstream and downstream firms transact in a spot-market-like fashion each month. But is this assumption consistent with the frequency of price adjustment observed in the auto industry and with the types of contracts OEMs commonly use with their suppliers? In the microeconomic price data used by the U.S. Bureau of Labor Statistics (BLS) to construct the producer price index from 1998 to 2008, the average duration of a price change in the transportation equipment sector is roughly one month (Goldberg and Hellerstein 2009), which is consistent with the monthly frequency of price adjustment assumed in our structural model. As for contracts, in a detailed study of the contracts OEMs use with their tier-one suppliers, Ben-Shahar and White (2007) show that while an OEM can have a long-term sourcing commitment to one supplier for a particular part, actual purchase orders are issued on a short-term basis and are typically associated with renegotiation on price. When setting up a new auto model production line, OEMs tend to create an overarching contract for four to eight years but then order parts in individual purchase orders whose duration can go from several days to twelve months. As we assume a fixed-proportions model for tier-one suppliers, we think of pass-through as conditional on any sequencing of price changes that may coincide with different contract durations for individual suppliers.²²

²²Regarding quantities, the auto industry's widespread use of "fixed-quantity-range contracts" was highlighted in a recent case that appeared before the Michigan Supreme Court (*ACEMCO Inc v Ryerson Rull Coil Processing*, September 28, 2008). Purchase orders generally specify a target quantity and allow for some deviation (e.g. +/- 20 percent) from that target.

4.4 Demand

The pass-through computations done with the sequential Bertrand-Nash supply models require consistent estimates of demand. Market demand is derived from a standard discrete-choice model of consumer behavior that follows the work of Berry, Levinsohn, and Pakes (1995) and Nevo (2000) among others. We use a random-coefficients logit model to estimate the demand system, as it is a very flexible and general model.²³

The accuracy of the pass-through coefficients depends in particular on consistent estimation of the curvature of the demand curve (see Villas-Boas and Hellerstein, 2006), that is, of the second derivative of the demand equation. The random-coefficients model allows for considerable flexibility in the specification and estimation of the curvature of demand that is not available with the simple logit. The model's parameters that characterize heterogeneity in consumers' tastes for various product characteristics, particularly their price sensitivity, also determine the curvature of demand. The scope of this heterogeneity determines a firm's ability to raise its price as its costs rise. As a firm raises its product's price, more price-sensitive consumers will respond by not purchasing the product or by dropping out of the market altogether (purchasing the outside good), meaning that the firm will retain only its less price-sensitive consumers. Note that consumers' sensitivity to price changes depends on their taste for other characteristics as well. If the price of a model with a high horsepower-to-length ratio rises, for example, consumers with a stronger preference for this characteristic will be more likely to remain in the market. The random-coefficients model also imposes very few restrictions on the demand system's own- and cross-price elasticities. This flexibility makes it the most appropriate model to study pass-through in this market.

Suppose consumer i chooses to purchase one unit of good j if and only if the utility from consuming that good is as great as the utility from consuming any other good. Using the typical notation for discrete choice models of demand, the indirect latent utility of consumer i from buying product j during market t is given by

$$u_{ijt} = d_j - \alpha_i p_{jt} + x_{jt} \beta_i + \xi_{jt} + \epsilon_{ijt}, \quad i = 1, \dots, I., \quad j = 1, \dots, J., \quad t = 1, \dots, T, \quad (9)$$

where d_j represents product fixed-effects capturing time-invariant product characteristics, p_{jt} is the price of product j , x_{jt} are product characteristics such as horsepower, ξ_{jt} identifies the mean across consumers of unobserved changes in product characteristics,²⁴ and ϵ_{ijt} represents the distribution of consumer preferences about this mean. The random coefficients α_i and

²³As the notation is fairly standard, we include in appendices D and E a more detailed description of the demand model, its estimation, and further discussion of demand identification issues.

²⁴In particular, ξ_{jt} includes changes in unobserved product characteristics, such as unobserved promotions, and changes in unobserved consumer preferences.

β_i represent the marginal disutility of price and the taste parameters with respect to characteristics x_{jt} , respectively. The marginal disutility of price varies across consumers according to $\alpha_i = \alpha + \Sigma_a v_i$, and the random coefficients on the characteristics x_{jt} vary according to $\beta_i = \beta + \Sigma_b v_i$. Unobserved consumer characteristics are contained in v_i and are assumed to be normally distributed. Finally, consumers have the option not to purchase any of the inside goods and to choose an outside good (or not to purchase at all), where the price of the outside good is assumed to be set independently of the prices observed in the sample.

Assuming that consumers purchase one unit of that product among all the possible products available in a certain market t that maximizes their indirect utility, then the market share of product j during market t is given by the probability that good j is chosen, that is,

$$s_{jt} = \int_{[(v_i, \epsilon_{it}) | u_{ijt} \geq u_{iht} \ \forall \ h=0, \dots, N_t]} dF(\epsilon) \ dF(v). \quad (10)$$

If v is fixed and consumer heterogeneity enters only through the random shock where ϵ_{ijt} is distributed i.i.d. with an extreme value type I density, then equation (10) becomes the multinomial logit model. Assuming that ϵ_{ijt} is distributed i.i.d. extreme value, and allowing consumer heterogeneity to affect the taste parameters for the different product characteristics, this model corresponds to the random-coefficients logit model, which allows for more general substitution patterns than in the multinomial logit model (McFadden and Train 2000).

5 Demand Estimation and Identification

This section describes the econometric procedures used to estimate the structural model's demand parameters consistently. Two issues arise in estimating a complete demand system in an oligopolistic market with differentiated products: the high dimensionality of elasticities to estimate and the potential endogeneity of price. Following McFadden (1973), Berry, Levinsohn, and Pakes (1995), and Nevo (2000) we draw on the discrete-choice literature to address the first issue and project the products onto a characteristics space with a much smaller dimension than the number of products. The second issue is that a product's price may be correlated with changes in its unobserved characteristics. We deal with this second issue by instrumenting for the potential endogeneity of price. We use used-auto prices for model year 2001 as instruments.

5.1 Demand Estimation

In principle, demand parameters can be estimated by choosing the value that implies predicted market shares as close as possible to observed shares. However, in this case the objective function will be nonlinear in the structural error term, which is a problem since price enters

as an endogenous variable. The key step is to construct a demand-side equation that is linear in the structural error so that instrumental variables estimation can be applied directly. Following Berry, Levinsohn and Pakes (1995) this estimation is done by equating the estimated product market shares to the observed shares and solving for the mean utility across all consumers, defined as

$$\delta_{jt}(\Sigma) = d_j + x_{jt}\beta_i - \alpha p_{jt} + \xi_{jt}. \quad (11)$$

For the random coefficient logit model, solving for the mean utility must be done numerically. Let θ be the demand-side parameters to be estimated, then $\theta = (\theta_L, \Sigma)$ where Σ are the nonlinear parameters. In the random coefficient logit model, θ is obtained by feasible method of simulated moments following Nevo's (2000) estimation algorithm, where equation (11) enters in one of the steps. This algorithm uses a nonlinear generalized-method-of-moments (GMM) procedure. The main step in the estimation is to construct a moment condition that interacts instrumental variables and a structural error term to form a nonlinear GMM estimator (Nevo 2000).

5.2 Instruments and Identification of Demand

The demand estimation requires instrumental variables because retailers consider all product characteristics when setting retail prices, not only those observed. Retailers also account for any changes in their products' characteristics and valuations. A product fixed effect is included to capture observed and unobserved product characteristics and valuations that are constant over time. The econometric error that remains in ξ_{jt} will therefore include only the changes in unobserved product characteristics such as unobserved promotions or changes in unobserved consumer preferences. Therefore prices will be correlated with changes in unobserved product characteristics affecting demand.

We use changes in the prices of used autos to instrument for prices. Prices for used autos should be correlated with those for new autos, which affect consumer demand, but not with changes in unobserved characteristics that influence consumer demand. Used-car prices are unlikely to have any relationship to the promotional activity that stimulates perceived changes in the characteristics of the sample's products. The price data for used autos come from Edmunds, and are make-, model-, and model-year-specific for used-car sales for each month in the new auto-price data. We use prices for used 2001 models. One might expect prices for used cars to be weakly correlated with those for new cars, thus generating a weak instrumental-variables problem, however, the model's first-stage results, reported in table 6, indicate that prices for used cars appear to be valid instruments.

5.3 Results of Demand Estimation

We first use the logit model to illustrate the need to instrument for prices when estimating demand. The first-stage part of table 6 reveals that the first-stage R -squared and F -statistic of the instrumental-variable specification are high, and the F -test for zero coefficients associated with the used-car series as instruments is rejected at any significance level. This outcome suggests that the instruments are useful for estimating demand parameters consistently. Table 6 presents the results from regressing the mean utility, which for the logit case is given by $\ln(s_{jt}) - \ln(s_{0t})$, on prices and product dummy variables in equation (9). The second column displays the estimate of ordinary least squares for the mean price coefficient α , and column three contains estimates of α for the instrumental variables (IV) specification. Because the consumer's sensitivity to price should increase after we instrument for unobserved changes in characteristics, it is promising that the price coefficient falls from -1.43 in the OLS estimation to -10.15 in the IV estimation.

Table 7 reports results from the random-coefficients demand model. We allow consumers' unobservable characteristics to interact with their taste coefficients for price and horsepower relative to vehicle length. As we estimate the demand equation using product fixed effects, we recover the consumer taste coefficients for time-invariant product characteristics in a generalized-least-squares regression of the estimated product fixed-effects on product characteristics. This regression assumes that changes in models' unobserved characteristics $\Delta\xi$ are independent of changes in models' observed characteristics x : $E(\Delta\xi|x) = 0$. The coefficients on the characteristics appear reasonable. The mean preference in the population is positive for more horsepower for a given auto length and negative with respect to price, as one would expect. The interaction of the horsepower variable with unobservables is positive and significant for the horsepower variable, indicating some heterogeneity in the population's preferences for this characteristic. The minimum-distance R^2 is 0.69 indicating that these characteristics explain the variation in the estimated product fixed effects fairly well.

6 Simulations

For the structural model, our empirical approach has two components: estimation and simulation of counterfactuals. At the estimation stage, we estimated the demand parameters and then the marginal costs and markups of the OEMs and TOPSs for each vertical-contractual model. To assess the overall impact of each vertical contract on firms' pass-through behavior, we employ simulation. Using the full random-coefficients model and the derived measures of marginal costs, we conduct counterfactual experiments to analyze how firms with different degrees of vertical integration respond to marginal cost shocks.

We consider the effect of a 10-percent cost shock to parts sourced from abroad on retail

prices in four scenarios, each with a different assumption about firms’ vertical integration. First, we study the effects of contracts at the two extreme points of vertical integration: the first, where all upstream value added is done in-house (the multinational case), and the second, where all upstream value added is outsourced to another firm. We compute the industry equilibrium that would emerge if firms experienced a 10-percent cost shock in the outsourced scenario and compare it to the equilibrium that prevails when one firm, a multinational, controls pricing along the distribution chain. We interpret the differential response of prices across the two cases as a measure of the overall impact of firms’ organization on their pass-through of shocks. Next, we compute pass-through with a structural model calibrated to each firm’s observed degree of vertical integration using the Lieberman and Dhawan (2005) measure from section 3. We then compare the results to the outcome given an increase of one standard deviation in the share of outsourced value added.

Simulations 1 and 2: Simulate the effect of a 10-percent cost shock when all upstream value added is outsourced or in-house. The first counterfactual experiment examines how retail prices adjust following a 10-percent change in the cost of foreign parts when all upstream production is fully outsourced. These results are reported in the third column of table 8. The median pass-through across brands is 13 percent, ranging from 7 percent for the Nissan Maxima to 27 percent for the Chrysler PT Cruiser. The second counterfactual experiment considers how retail prices adjust following a 10-percent change in the cost of foreign parts if all upstream production is done in-house. These results appear in the fourth column of table 8. The median pass-through of the cost shock to the retail price is 54 percent but varies significantly across brands, ranging from 32 percent for the Chevrolet Corvette to 120 percent for the Toyota Corolla.

The median differences between the results from the “outsourced” and the “multinational” simulations are reported in the final column of table 8: 40 percentage points across models, ranging from 18 percentage points for the Chrysler PT Cruiser to 107 percentage points for the Toyota Corolla. The median pass-through elasticity is higher in the in-house than in the outsourced counterfactual for every auto model. This table drives home the point that consumer prices are most insulated from cost shocks when there are multiple optimizations along a production chain. The differences are so stark between the two models because the margin available for adjustment in the outsourced case is much larger than that in the in-house case. The inefficiencies associated with double marginalization in an outsourced production chain have a secondary effect, given a cost shock: that is, the pass-through of that shock to prices is dampened. Outsourcing thus reduces the cross-border transmission of shocks, all else equal. This result is consistent with the patterns we documented in section 3 suggesting that multinational firms pass through marginal-cost shocks at a higher rate than do outsourced production chains. However, these simulations use a stylized model with two extreme points

of vertical integration. Although it serves to make our point, it is too abstract to compare directly to the regression results for the auto industry. To that end, we turn to a more realistic structural model in the next two simulations.

Simulation 3: Simulate the effect of a 10-percent cost shock when each firm’s vertical integration is calibrated to its observed outsourced value added. The third counterfactual experiment considers how retail prices adjust following a 10-percent change in the cost of foreign parts when each firm’s vertical integration is calibrated to its observed outsourced value added. This experiment is done by calibrating the ownership matrix of tier-one suppliers to the Lieberman and Dhawan (2005) measure. Recall that in the structural model, each production chain’s degree of vertical integration is expressed through the ownership matrix for suppliers, T_s , as defined in equation (6). If a supplier and manufacturer are completely vertically integrated, then the diagonal entries of T_s will be equal to zero and only the manufacturer will have a markup. If they operate in an outsourced supply chain, then each diagonal entry of T_s will be a 1. Between these two extremes, each entry of T_s will reflect the average degree of integration between that OEM and its suppliers measured by its total in-house value added relative to its total revenue. The results from this simulation are reported in the “observed” column in table 9. Across all the auto models in the sample, the simulated pass-through coefficient is 25 percent and ranges from 2 percent for the Nissan Maxima to 63 percent for the Ford Ranger.

Simulation 4: Simulate the effect of a 10-percent cost shock when each firm’s vertical integration is calibrated to one standard deviation above its observed outsourced value added. The final counterfactual experiment considers how retail prices adjust following a 10-percent shock to the cost of foreign-sourced parts when each firm’s vertical integration is calibrated to be one standard deviation above its observed outsourced value added. These results appear in the middle column of table 9. A one-standard-deviation increase in outsourcing decreases the value of each diagonal entry of the supplier-ownership matrix, T_s , and results in a median pass-through elasticity of 14 percent across models that ranges from 1 percent for the Nissan Maxima to 31 percent for the Ford Ranger. The last column of table 9 reports the difference between the pass-through results from simulations 3 and 4. It shows that a one-standard-deviation increase in outsourcing is associated with a 12-percent decline in the median pass-through across models, from 25 to 14 percent. The results are striking, with the most substantive changes for models where the pass-through elasticity is large in absolute terms: It drops from 63 to 31 percent for the Ford Ranger and from 51 to 23 percent for the GM Impala. The counterfactual experiments thus confirm that pass-through is higher in a vertically-integrated production chain than in a vertically separated production chain.

To give a visual impression of the fit of our structural model, in figure 4 we compare the

average estimated pass-through elasticities by firm from the section 3 regressions to the average simulated pass-through elasticities by firm from each of our three structural models: the outsourced, the multinational, and the observed-degree-of-vertical-integration models. The resulting scatterplots are in figure 4, where the values on the x-axis correspond to the average estimated pass-through from the regressions and the values on the y-axis correspond to the average simulated pass-through from each of the structural models. The figure shows that the pass-through elasticities produced by the structural model calibrated to each OEM’s observed degree of vertical integration from section 5 produces the best fit to the pass-through elasticities produced by the regressions from section 3. It illustrates that the structural model calibrated to each auto model’s observed degree of vertical integration explains 38 percent of the variation in pass-through across OEMs and roughly 20 percent of that variation that is not explained by the outsourced or multinational models. Thus, it provides some visual evidence that although our structural model is somewhat stylized, it does a reasonable job of explaining the differences in estimated pass-through across OEMs.

7 Conclusion

This paper shows that the organization of firms has a clear relationship to their pass-through of cost shocks to their prices. Our regression analysis suggests that pass-through is positively related to firms’ degree of vertical integration across industries and in the auto industry in particular. Given this stylized fact, we estimate a structural model and use it to assess the extent to which pass-through is related to vertical integration by simulating firms’ pass-through behavior in illustrative what-if scenarios. Our findings suggest that a significant portion of incomplete cross-border pass-through may result from successive optimizations by firms along a production chain that spans national borders.

One implication of our work is that the overall effect of an exchange-rate shock on a domestic economy will vary in its magnitude and its distribution across domestic firms, foreign firms, and consumers depending on the vertical contract that dominates the economy’s import and export sectors. Future research might explore the implications of these findings for exchange-rate pass-through patterns across countries with different industry mixes and thus dominant vertical contracts in their import and export sectors. To give an example, most rich-to-rich country trade is multinational, while most poor-to-rich country trade is outsourced. It is worth exploring further how the dominant form of firm organization in a given country’s cross-border transactions affects the transmission of foreign exchange-rate shocks to the domestic economy. It would be particularly interesting to determine how many of the stylized facts about the differing responses of developed and developing economies to external shocks can be attributed to the prevailing type of firm organization in their cross-border transactions.

References

- [1] Alterman, W., 1997. A Comparison of the Export and Producer Price Indexes for Semiconductors. Paper presented at the NBER's Summer Institute, July.
- [2] Antras, P., 2003. Firms, Contracts, and Trade Structure. *Quarterly Journal of Economics* 118, 1375-1418.
- [3] Ben-Shahar, O., White, J. J., 2007. Boilerplate and Economic Power in Auto-Manufacturing Contracts, in: Ben-Shahar, O. (Ed.), *Boilerplate: The Foundation of Market Contracts*, Cambridge University Press, Cambridge, pp. 29-44.
- [4] Bernard, A., Jensen, J. B., Redding, S. J., Schott, P. K., 2008. Intra-Firm Trade and Product Contractability, Dartmouth University Working Paper.
- [5] Bernard, A., Jensen, J. B., Schott, P. K., 2006. Transfer Pricing by U.S.-Based Multinational Firms. NBER Working Paper No. 12493.
- [6] Berry, S., Levinsohn, J., and Pakes, A., 1995. Automobile Prices in Market Equilibrium. *Econometrica* 63, 841-90.
- [7] Bettendorf, L., Verboven, F., 2000. Incomplete Transmission of Coffee Bean Prices: Evidence from the Netherlands. *European Review of Agricultural Economics* 27, 1-16.
- [8] Bonnet, C., Dubois, P., Villas-Boas, S. B., 2009. Empirical Evidence on the Role of Non-Linear Wholesale Pricing and Vertical Restraints on Cost Pass-Through. Mimeo.
- [9] Bordo, M., Eichengreen, B., Irwin, D., 1999. Is Globalization Today Really Different from Globalization a Hundred Years Ago? in: Rodrik D., Collins, S. (Eds.), *Brookings Trade Forum*, pp. 1-50.
- [10] Brenkers, R., Verboven, F., 2006. Liberalizing a Distribution System: The European Car Market. *Journal of the European Economic Association* 4, 216-51.
- [11] Burstein, A. T., Neves, J. C., Rebelo, S., 2003. Distribution Costs and Real Exchange Rate Dynamics during Exchange-Rate-Based Stabilization. *Journal of Monetary Economics* 50, 1189-1214.
- [12] Busse, M., Silva-Risso, J., Zettelmeyer, F., 2006. \$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions. *The American Economic Review* 96, 1253-70.
- [13] Campa, J. M., Goldberg, L., 2005. Exchange-Rate Pass-Through into Import Prices: A Macro or Micro Phenomenon? *Review of Economics and Statistics* 87, 679-90.
- [14] Clausing, K., 2003. Tax-Motivated Transfer Pricing and U.S. Intrafirm Trade Prices. *Journal of Public Economics* 87, 2207-23.

- [15] Corsetti, G., Dedola, L., 2005. Macroeconomics of International Price Discrimination, *Journal of International Economics* 67, 129-55.
- [16] Corsetti, G., Dedola, L., Leduc, S., 2010. Optimal Monetary Policy and the Sources of Local-Currency Price Stability, in: Gali, J., Gertler, M., (Eds.), *International Dimensions of Monetary Policy*. University of Chicago Press, Chicago, pp. 319-67.
- [17] Engel, C., 2002. Expenditure Switching and Exchange-Rate Policy. *NBER Macroeconomics Annual* 17, 231-72.
- [18] Feenstra, R. C., 1989. Symmetric Pass-Through of Tariffs and Exchange Rates Under Imperfect Competition: An Empirical Test. *Journal of International Economics* 27, 25-45.
- [19] Feenstra, R. C., Gagnon, J., Knetter, M., 1996. Market Share and Exchange Rate Pass-through in World Automobile Trade, *Journal of International Economics* 40, 187-208.
- [20] Goldberg, P. K., Hellerstein, R., 2008. A Framework for Identifying the Sources of Local Currency Price Stability with an Empirical Application. NBER Working Paper No. 13183.
- [21] Goldberg, P. K., Hellerstein, R., 2009. How Rigid Are Producer Prices? Federal Reserve Bank of New York Staff Report No. 407.
- [22] Goldberg, P. K., Knetter, M., 1997. Goods Prices and Exchange Rates: What Have We Learned? *Journal of Economic Literature* 35, 1243-72.
- [23] Goldberg, P. K., Verboven, F., 2001. The Evolution of Price Dispersion in the European Car Market. *Review of Economic Studies* 67, 811-48.
- [24] Gopinath, G., Rigobon, R., 2008. Sticky Borders. *Quarterly Journal of Economics* 123, 531-75.
- [25] Gopinath, G., Itskhoki, O., Rigobon, R., 2010. Currency Choice and Exchange Rate Pass-through. *The American Economic Review*, forthcoming.
- [26] Gron, A., Swenson, D., 2000. Cost Pass-Through in the U.S. Automobile Market. *Review of Economics and Statistics* 82, 316-24.
- [27] Hastings, J. S., 2004. Vertical Relationships and Competition in Retail Gasoline Markets. *The American Economic Review* 94, 317-28.
- [28] 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* 76, 14-32.
- [29] Helpman, E., 2006. Trade, FDI, and the Organization of Firms. *Journal of Economic Literature* 44, 589-630.
- [30] Hendel, I., Nevo, A., 2006. Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica* 74, 1637-73.

- [31] Hummels, D., Ishii, J., Yi, K., 2001. The Nature and Growth of Vertical Specialization in World Trade. *Journal of International Economics* 54, 75-96.
- [32] Krugman, P., 1987. Pricing to Market when the Exchange Rate Changes, in: Richardson J., Arndt, S. (Eds.), *Real-Financial Linkages among Open Economies*, MIT Press, Cambridge.
- [33] Lieberman, M. B., Dhawan, R., 2005. Assessing the Resource Base of Japanese and U.S. Auto Producers. *Management Science* 51, 1060-75.
- [34] McFadden, D., 1973. Conditional Logit Analysis of Qualitative Choice Behavior, in: Zarembka, P. (Ed.), *Frontiers of Econometrics*, Academic Press, New York, pp. 105-42.
- [35] McFadden, D., Train, K., 2000. Mixed MNL Models for Discrete Response. *Journal of Applied Econometrics* 15, 447-70.
- [36] Nakamura, E., Zerom, D., 2009. Accounting for Incomplete Pass-Through. *Review of Economic Studies*, forthcoming.
- [37] Neiman, B., 2010. Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices. *Journal of Monetary Economics*, forthcoming.
- [38] Nevo, A., 2000. A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics and Management Strategy* 9, 513-48.
- [39] Nunn, N., Treffer, D., 2008. The Boundaries of the Multinational Firm, in: Helpman, E., Marin, D., Verdier, T. (Eds.), *The Organization of Firms in a Global Economy*, Harvard University Press, Cambridge, pp. 55-83.
- [40] Rangan, S., Lawrence, R. Z., Cooper, R., 1993. The Responses of U.S. Firms to Exchange Rate Fluctuations: Piercing the Corporate Veil. *Brookings Papers on Economic Activity* 1993, 341-79.
- [41] Rauch, J. E., 1999. Networks versus Markets in International Trade. *Journal of International Economics* 48, 7-35.
- [42] Ravn, M., Schmitt-Grohe, S., Uribe, M., 2009. Incomplete Cost Pass-through Under Deep Habits. *Review of Economic Dynamics*, forthcoming.
- [43] Villas-Boas, S. B., 2007. Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data. *Review of Economic Studies* 74, 625-52.
- [44] Villas-Boas, S. B., Hellerstein, R., 2006. Identification of Supply Models of Retailer and Manufacturer Oligopoly Pricing. *Economics Letters* 90, 132-40.
- [45] Zeile, W. J., 1997. U.S. Intrafirm Trade in Goods. *Survey of Current Business* 77, 23-38.

Figure 1: *Relationship between an Industry's Share of Intrafirm Imports and Its Pass-through.*
Source: U.S. Bureau of Labor Statistics and authors' calculations.

Figure 2: *Comparison of the Manufacturer's Suggested Retail Price and Edmunds' Total Market Value Price for the Ford Escape.* Source: Edmunds.

Figure 3: *Relationship of an Automaker's Degree of Vertical Integration and Its Average Pass-through across Models, October 2002-June 2006.* Source: Authors' calculations.

Figure 4: *Pass-through and Firm Organization.* Scatterplots and OLS regression lines comparing estimated pass-through and simulated pass-through for each structural model, arm's-length (solid line and squares), multinational (dashed line and triangles), and observed (calibrated to each automakers' observed internal value added relative to total revenue: dotted line and diamonds). Source: Authors' calculations.

Table 1. Industry Pass-through Regressions

Industry	NAICS	Exchange Rate* Foreign Costs	Standard Error	R-Squared
Food	311	-0.02	(.10)	0.17
Beverages and Tobacco	312	0.20	(.08)***	0.21
Textiles	313	0.39	(.09)***	0.28
Apparel	315	0.07	(.09)	0.02
Leather Goods	316	-0.02	(.04)	0.06
Chemicals	325	0.40	(.07)***	0.37
Rubber and Plastics	326	0.37	(.13)***	0.15
Nonmetallic Metals	327	0.46	(.13)***	0.21
Primary Metals	331	0.55	(.19)***	0.24
Fabricated Metals	332	0.56	(.08)***	0.44
Industrial Machinery	333	0.55	(.10)***	0.43
Computers	334	0.45	(.16)***	0.10
Electrical Equipment	335	0.54	(.15)***	0.37
Autos	336	0.43	(.07)***	0.44
Furniture	337	0.30	(.13)**	0.12
Miscellaneous	339	0.37	(.09)***	0.35

Notes: Each regression has 80 observations. The dependent variable is the quarterly industry import-price index from the *U.S. Bureau of Labor Statistics* from 1985 to 2004. The industry exchange rate is the import-weighted exchange rate for each industry with weights derived from the average share of each country's imports in the industry's total imports from 1987 to 2004. The foreign cost variable is the import-weighted foreign CPI, with the same weights as those for the exchange rate. The exchange rate and CPI data come from the *IMF's International Financial Statistics*. The trade data come from the *United States International Trade Commission* and the *Bureau of Economic Analysis*, *U.S. Department of Commerce*. The regressions also control for U.S. domestic demand with U.S. Gross Domestic Purchases from the *Bureau of Economic Analysis*. The domestic demand coefficients are reported in Appendix Table A.1. Starred coefficients are significant at the *10-, **5-, and ***1-percent level.

Table 2. Robustness Checks for Pass-through Patterns

Variable	Industry Pass-through	Industry Pass-through	Industry Pass-through	Industry Pass-through
Industry Share of Intra-Firm Imports	0.57 (.20)**	0.59 (.22)**	0.57 (.22)**	0.56 (.18)***
Industry Capital Stock		-0.01 (.03)		
Rauch Classification			0.05 (.38)	
Market Share of Foreign Producers				-0.24 (.30)
Constant	14.03 (9.51)	14.31 (10.09)	10.39 (14.62)	19.48 (12.62)
R-Squared	0.33	0.33	0.34	0.36
Observations	16	16	16	16

Notes: The dependent variable is the industry pass-through coefficient from Table 1. The industry share of intra-firm imports is for the year 2000 and comes from Table 2 of Bernard, Jensen, Redding, and Schott (2008). The industry capital stock is from the NBER-CES Manufacturing Industry Database for 1996. The market share of foreign producers is the mean share of imports in the domestic market (comprised of imports+domestic shipments) from 1987 to 2004. The Rauch classification is the import-weighted share of differentiated products in each 3-digit NAICS industry.

Table 3. Summary Statistics for the Automobile Data

Description	Mean	Standard Deviation	Min	Max
Retail TMV Price (\$)	24,137	7,177	12,949	57,256
Manufacturer TMV Price (\$)	23,308	6,670	8,450	50,423
Retail Used TMV Price (\$)	26,999	10,853	6,165	40,420
Foreign Content (%)	25.5	16.9	5.0	85.0
Horsepower/Length	29.0	3.0	24.8	34.9

Source: *Edmunds; Automobile Trade Policy Council.*

Table 4. Sales-Weighted Share of Foreign Sourcing

Description	Share (%)
Maker	
<i>DaimlerChrysler</i>	0.29
<i>Ford</i>	0.18
<i>GM</i>	0.21
<i>Honda</i>	0.47
<i>Nissan</i>	0.56
<i>Toyota</i>	0.57

Source: *U.S. Department of Commerce; Automotive Trade Policy Council.*

Table 5. Pass-through Regressions for the *Edmunds'* Automobile Data

Maker	Model	Exchange Rate* Foreign Costs	Standard Error	Domestic Costs	Standard Error	R-Squared
<i>Daimler-Chrysler</i>	<i>Caravan</i>	0.67	(.38)*	-0.46	(.67)	0.19
<i>Daimler-Chrysler</i>	<i>Grand Cherokee</i>	0.25	(.10)**	-0.26	(.25)	0.07
<i>Daimler-Chrysler</i>	<i>PT Cruiser</i>	0.90	(.27)***	1.45	(1.03)	0.33
<i>Daimler-Chrysler</i>	<i>Town & Country</i>	0.43	(.24)*	-0.34	(.45)	0.09
<i>Daimler-Chrysler</i>	<i>Wrangler</i>	0.30	(.09)***	0.32	(.44)	0.85
<i>Ford</i>	<i>Escape</i>	0.16	(.07)**	-0.04	(.20)	0.48
<i>Ford</i>	<i>Explorer</i>	0.01	(.26)	0.03	(.24)	0.0036
<i>Ford</i>	<i>MPV</i>	0.52	(.17)***	-0.54	(.68)	0.32
<i>Ford</i>	<i>Mountaineer</i>	0.41	(.08)***	0.61	(.44)	0.41
<i>Ford</i>	<i>Mustang</i>	0.60	(.23)***	-0.17	(.41)	0.27
<i>Ford</i>	<i>Ranger</i>	0.20	(.06)***	0.34	(.27)	0.68
<i>Ford</i>	<i>Taurus</i>	0.08	(.04)**	-0.10	(.20)	0.68
<i>GM</i>	<i>Corvette</i>	0.50	(.22)**	-0.49	(.38)	0.36
<i>GM</i>	<i>Impala</i>	0.61	(.15)***	-0.14	(0.49)	0.59
<i>GM</i>	<i>Malibu</i>	0.68	(.18)***	0.37	(.75)	0.52
<i>GM</i>	<i>Monte Carlo</i>	0.71	(.17)***	0.38	(.58)	0.55
<i>GM</i>	<i>Tahoe</i>	0.56	(.12)***	0.25	(.56)	0.34
<i>Honda</i>	<i>Accord</i>	0.15	(0.11)	0.05	(.45)	0.61
<i>Honda</i>	<i>Civic</i>	0.62	(.20)***	0.03	(.32)	0.76
<i>Honda</i>	<i>Odyssey</i>	-0.02	(.22)	0.16	(.91)	0.51
<i>Nissan</i>	<i>Altima</i>	0.16	(.09)*	0.32	(.14)**	0.90
<i>Nissan</i>	<i>Maxima</i>	0.44	(.10)***	0.70	(.46)	0.74
<i>Toyota</i>	<i>Camry</i>	0.13	(.068)*	-0.02	(.10)	0.43
<i>Toyota</i>	<i>Corolla</i>	0.14	(.06)**	0.19	(.29)	0.87

Notes: The dependent variable is the monthly *Total Market Value (TMV)* price from *Edmunds* from October 2002 to June 2006. Each regression has 37 observations. The regressions also control for U.S. domestic demand with U.S. gross domestic purchases from the *Bureau of Economic Analysis, U.S. Department of Commerce*.

Table 6. Diagnostic Results from the Logit Demand Model

Variable	OLS	IV
Price	-1.43 (.26)***	-10.15 (1.60)***
Adjusted R-Squared	0.86	
Observations	888	888
First-Stage Results		
F-Statistic		61.15
Instruments		used prices

Notes: 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. Starred coefficients are significant at the *** 1-percent level. Source: Authors' calculations.

Table 7. Results from the Random-Coefficients Demand Model

Variable	Population Mean	Interaction with Unobservables
Constant	12.31 (.90)***	
Price	-2.38 (1.19)**	0.02 (.03)
Horsepower/Length	0.23 (.002)***	1.56 (.34)***
M-D Weighted R-Squared	0.69	
Observations	888	

Notes: Starred coefficients are significant at the * 10-, **5-, and *** 1-percent level. Source: Authors' calculations.

Table 8. Counterfactual experiments: Simulated pass-through of a foreign cost shock given variation in manufacturers' outsourcing

manufacturer	model	outsourced		multinational		difference	
		coefficient	s.e.	coefficient	s.e.	coefficient	s.e.
<i>Daimler-Chrysler</i>	<i>Caravan</i>	0.16	(.05)*	0.77	(.28)*	0.61	(.24)*
<i>Daimler-Chrysler</i>	<i>Grand Cherokee</i>	0.11	(.03)*	0.51	(.13)*	0.40	(.12)*
<i>Daimler-Chrysler</i>	<i>PT Cruiser</i>	0.27	(.04)*	0.45	(.19)*	0.18	(.09)*
<i>Daimler-Chrysler</i>	<i>Town & Country</i>	0.14	(.03)*	0.44	(.12)*	0.30	(.10)*
<i>Daimler-Chrysler</i>	<i>Wrangler</i>	0.15	(.04)*	0.39	(.11)*	0.24	(.08)*
<i>Ford</i>	<i>Escape</i>	0.11	(.05)*	0.74	(.21)*	0.64	(.27)*
<i>Ford</i>	<i>Explorer</i>	0.12	(.03)*	0.51	(.13)*	0.39	(.12)*
<i>Ford</i>	<i>Mountaineer</i>	0.15	(.03)*	0.48	(.15)*	0.33	(.10)*
<i>Ford</i>	<i>MPV</i>	0.23	(.06)*	0.58	(.12)*	0.35	(.17)*
<i>Ford</i>	<i>Mustang</i>	0.14	(.03)*	0.39	(.11)*	0.25	(.09)*
<i>Ford</i>	<i>Ranger</i>	0.22	(.06)*	0.90	(.19)*	0.67	(.21)*
<i>Ford</i>	<i>Taurus</i>	0.13	(.04)*	0.83	(.28)*	0.70	(.26)*
<i>GM</i>	<i>Corvette</i>	0.10	(.02)*	0.32	(.06)*	0.22	(.05)*
<i>GM</i>	<i>Impala</i>	0.13	(.04)*	0.73	(.23)*	0.60	(.22)*
<i>GM</i>	<i>Malibu</i>	0.10	(.07)*	0.50	(.15)*	0.40	(.16)*
<i>GM</i>	<i>Monte Carlo</i>	0.13	(.04)*	0.62	(.17)*	0.49	(.18)*
<i>GM</i>	<i>Tahoe</i>	0.17	(.03)*	0.46	(.11)*	0.29	(.09)*
<i>Honda</i>	<i>Accord</i>	0.08	(.03)*	0.50	(.21)*	0.41	(.15)*
<i>Honda</i>	<i>Civic</i>	0.12	(.04)*	1.09	(.47)*	0.97	(.30)*
<i>Honda</i>	<i>Odyssey</i>	0.08	(.03)*	0.45	(.12)*	0.38	(.33)
<i>Nissan</i>	<i>Altima</i>	0.10	(.03)*	0.64	(.21)*	0.54	(.22)*
<i>Nissan</i>	<i>Maxima</i>	0.07	(.03)*	0.33	(.09)*	0.26	(.09)*
<i>Toyota</i>	<i>Camry</i>	0.09	(.03)*	0.59	(.22)*	0.50	(.22)*
<i>Toyota</i>	<i>Corolla</i>	0.13	(.05)*	1.20	(.71)	1.07	(.94)
<i>Total</i>		0.13	(.03)*	0.54	(.13)*	0.40	(.13)*

Notes: Simulated coefficients from structural model with degree of vertical integration calibrated to reflect outsourced or multinational transactions along the supply chain. Starred coefficients are significant at the 5-percent level. Standard errors for the simulated coefficients are computed in 200 bootstrap simulations of the structural model. Source: Authors' calculations.

Table 9. Counterfactual experiments: Simulated pass-through of a foreign cost shock given variation in manufacturers' outsourcing

model		observed		1 std-dev rise in outsourcing		difference	
		coefficient	s.e.	coefficient	s.e.	coefficient	s.e.
<i>Daimler-Chrysler</i>	<i>Caravan</i>	0.28	(.11) [*]	0.15	(.05) [*]	0.13	(.04) [*]
<i>Daimler-Chrysler</i>	<i>Grand Cherokee</i>	0.20	(.08) [*]	0.12	(.04) [*]	0.08	(.03) [*]
<i>Daimler-Chrysler</i>	<i>PT Cruiser</i>	0.28	(.05) [*]	0.23	(.02) [*]	0.05	(.02) [*]
<i>Daimler-Chrysler</i>	<i>Town and Country</i>	0.20	(.07) [*]	0.13	(.03) [*]	0.07	(.03) [*]
<i>Daimler-Chrysler</i>	<i>Wrangler</i>	0.19	(.08) [*]	0.12	(.03) [*]	0.08	(.03) [*]
<i>Ford</i>	<i>Escape</i>	0.32	(.13) [*]	0.12	(.03) [*]	0.20	(.10) [*]
<i>Ford</i>	<i>Explorer</i>	0.37	(.16) [*]	0.16	(.05) [*]	0.21	(.09) [*]
<i>Ford</i>	<i>Mountaineer</i>	0.30	(.13) [*]	0.20	(.06) [*]	0.10	(.16)
<i>Ford</i>	<i>MPV</i>	0.23	(.08) [*]	0.21	(.04) [*]	0.02	(.04)
<i>Ford</i>	<i>Mustang</i>	0.34	(.13) [*]	0.18	(.05) [*]	0.15	(.06) [*]
<i>Ford</i>	<i>Ranger</i>	0.63	(.26) [*]	0.31	(.12) [*]	0.32	(.15) [*]
<i>Ford</i>	<i>Taurus</i>	0.55	(.23) [*]	0.25	(.10) [*]	0.30	(.14) [*]
<i>GM</i>	<i>Corvette</i>	0.12	(.05) [*]	0.09	(.03) [*]	0.03	(.03)
<i>GM</i>	<i>Impala</i>	0.51	(.20) [*]	0.23	(.09) [*]	0.29	(.10) [*]
<i>GM</i>	<i>Malibu</i>	0.46	(.18) [*]	0.23	(.09) [*]	0.23	(.07) [*]
<i>GM</i>	<i>Monte Carlo</i>	0.44	(.17) [*]	0.20	(.08) [*]	0.24	(.09) [*]
<i>GM</i>	<i>Tahoe</i>	0.31	(.14) [*]	0.21	(.06) [*]	0.10	(.09)
<i>Honda</i>	<i>Accord</i>	0.03	(.01) [*]	0.03	(.01) [*]	0.01	(.03)
<i>Honda</i>	<i>Civic</i>	0.10	(.04) [*]	0.02	(.03)	0.08	(.01) [*]
<i>Honda</i>	<i>Odyssey</i>	0.07	(.01) [*]	0.04	(.03)	0.04	(.02) [*]
<i>Nissan</i>	<i>Altima</i>	0.03	(.01) [*]	0.02	(.02)	0.01	(.01)
<i>Nissan</i>	<i>Maxima</i>	0.02	(.01) [*]	0.01	(.02)	0.01	(.02)
<i>Toyota</i>	<i>Camry</i>	0.03	(.02)	0.02	(.02)	0.01	(.02)
<i>Toyota</i>	<i>Corolla</i>	0.09	(.06)	0.03	(.04)	0.05	(.01)
<i>Total</i>		0.26	(.08) [*]	0.15	(.04) [*]	0.12	(.04) [*]

Notes: Simulated coefficients from the structural model with degree of vertical integration calibrated to each automakers' share of internal value-added over total revenue. Starred coefficients are significant at the 5-percent level. Standard errors for the simulated coefficients are computed in 200 bootstrap simulations of the structural model. Source: Authors' calculations.

A Alternative Industry Pass-Through Results

This appendix contains results from industry regressions with different exchange rates. In the import data, we observe a cluster of countries below the 2-percent average import share and another cluster below the 7.5 percent average import share. We draw on these patterns to construct two new sets of industry exchange rates that exclude the exchange rates of countries with less than 2 percent or 7.5 percent of the average industry imports, respectively. We conceive of the first set as excluding countries with a negligible industry import share and the second set as including only major trading partners. Appendix Table A.2 reports the differences between the original pass-through coefficients and those from the two sets of alternative regressions with the new exchange rates. While the different exchange rates produce some variation in individual industry's pass-through coefficients, across industries these differences net out to zero. The mean difference across industries between the >2-percent regressions and the original regressions is 1 percentage point, and between the >7.5-percent regressions and the original regressions is 2 percentage points.

The differences between the coefficients from the >7.5-percent regressions and those from the original regressions are negatively correlated with the share of differentiated products in each industry – the correlation coefficient is -.66. Pass-through coefficients for industries with less differentiated products generally rise as one excludes countries with smaller market shares, the most notable case being *Primary Metals* whose pass-through coefficient doubles from the original regressions to the >7.5-percent regressions, while the coefficients for industries with more differentiated products generally fall, the most dramatic cases being *Electrical Equipment*, whose pass-through coefficient falls by 21 percentage points, and *Furniture*, whose pass-through coefficient falls by 18 percentage points (and becomes statistically insignificant). For other industries, the differences are generally small (in the single digits). Our interpretation of these differences is that in less differentiated industries, firms with low market shares generally have little market power, and so little ability to pass-through cost shocks to prices. In more differentiated industries, in contrast, low market share may reflect other sources of market power, including specialized niches that some firms occupy in the product space. These industries may exhibit oligopolistic interactions between major trading partners that limits their pass-through relative to trading partners with smaller market shares that does not occur in less differentiated industries.

Figure A.1: *Industries with high intra-firm trade have high exchange-rate pass-through.* The industry-level exchange-rates include the bilateral exchange rates of only those countries with at least 2 percent of average industry imports over the sample period from 1985 to 2004. Source: U.S. Bureau of Labor Statistics and Authors' calculations.

Figure A.2: *Industries with high intra-firm trade have high exchange-rate pass-through.* The industry-level exchange-rates include the bilateral exchange rates of only those countries with at least 7.5 percent of average industry imports over the sample period from 1985 to 2004. Source: U.S. Bureau of Labor Statistics and Authors' calculations.

Appendix Table A.1. Original Industry Pass-through Regressions

Industry	NAICS	Exchange Rate* Foreign Costs	Standard Error	Domestic Demand	Standard Error	R-Squared
Food	311	-0.02	(.10)	1.36	(.63)**	0.17
Beverages and Tobacco	312	0.20	(.08)***	-0.50	(.22)**	0.21
Textiles	313	0.39	(.09)***	-0.08	(.23)	0.28
Apparel	315	0.07	(.09)	0.09	(.20)	0.02
Leather Goods	316	-0.02	(.04)	0.44	(.32)	0.06
Chemicals	325	0.40	(.07)***	0.49	(.27)*	0.37
Rubber and Plastics	326	0.37	(.13)***	-0.25	.34	0.15
Nonmetallic Metals	327	0.46	(.13)***	0.50	(.32)	0.21
Primary Metals	331	0.55	(.19)***	2.87	(.86)***	0.24
Fabricated Metals	332	0.56	(.08)***	0.17	-0.21	0.44
Industrial Machinery	333	0.55	(.10)***	0.20	(.37)	0.43
Computers	334	0.45	(.16)***	0.00	(.45)	0.10
Electrical Equipment	335	0.54	(.15)***	0.31	(.38)	0.37
Autos	336	0.43	(.07)***	-0.19	(.19)	0.44
Furniture	337	0.30	(.13)**	0.04	(.28)	0.12
Miscellaneous	339	0.37	(.09)***	0.54	(.27)**	0.35

Notes: Each regression has 80 observations. The dependent variable is the quarterly industry import-price index from the U.S. Bureau of Labor Statistics from 1985 to 2004. The industry exchange rate is the import-weighted exchange rate for each industry with weights derived from the average share of each country's imports in the industry's total imports from 1987 to 2004. The foreign cost variable is the import-weighted foreign CPI, with the same weights as those for the exchange rate. The exchange rate and CPI data come from International Financial Statistics. The trade data come from the United States International Trade Commission and the Bureau of Economic Analysis, U.S. Department of Commerce. The regressions also control for U.S. domestic demand with U.S. Gross Domestic Purchases from the Bureau of Economic Analysis. Starred coefficients are significant at the *10-, **5-, and ***1-percent level.

Appendix Table A.2. Industry Pass-through Regressions with Different Exchange Rates

		Mean Import Shares		Import Shares > .02			Import Shares > .075		
Industry	NAICS	Exchange Rate* Foreign Costs	Standard Error	Exchange Rate* Foreign Costs	Standard Error	Difference from Original	Exchange Rate* Foreign Costs	Standard Error	Difference from Original
Food	311	-0.02	(.10)	-0.01	(.12)	0.01	0.47	(.35)	0.49
Beverages and Tobacco	312	0.20	(.08)***	0.38	(.14)***	0.18	0.29	(.08)***	0.09
Textiles	313	0.39	(.09)***	0.60	(.09)***	0.21	0.48	(.11)***	0.09
Apparel	315	0.07	(.09)	-0.01	(.08)	-0.08	-0.07	(.09)	-0.14
Leather Goods	316	-0.02	(.04)	0.00	(.04)	0.02	0.01	(.04)	0.03
Chemicals	325	0.40	(.07)***	0.43	(.08)***	0.03	0.47	(.09)***	0.07
Rubber and Plastics	326	0.37	(.13)***	0.40	(.15)***	0.03	0.37	(.15)**	0.00
Nonmetallic Metals	327	0.46	(.13)***	0.60	(.14)***	0.14	0.33	(.22)	-0.13
Primary Metals	331	0.55	(.19)***	0.44	(.21)***	-0.11	1.02	(.36)**	0.47
Fabricated Metals	332	0.56	(.08)***	0.53	(.09)***	-0.03	0.51	(.11)***	-0.05
Industrial Machinery	333	0.55	(.10)***	0.51	(.11)***	-0.04	0.52	(.12)***	-0.03
Computers	334	0.45	(.16)***	0.41	(.16)**	-0.04	0.41	(.19)**	-0.04
Electrical Equipment	335	0.54	(.15)***	0.51	(.11)***	-0.03	0.33	(.19)*	-0.21
Autos	336	0.43	(.07)***	0.39	(.06)***	-0.04	0.39	(.08)***	-0.04
Furniture	337	0.30	(.13)**	0.28	(.12)**	-0.02	0.12	(.16)	-0.18
Miscellaneous	339	0.37	(.09)***	0.27	(.11)**	-0.10	0.27	(.12)**	-0.10
Mean		0.35		0.36		0.01	0.37		0.02

Notes: Each regression has 80 observations. The dependent variable is the quarterly industry import-price index from the U.S. Bureau of Labor Statistics from 1985 to 2004. The industry exchange rate is the import-weighted exchange rate for each industry with weights derived from the average share of each country's imports in the industry's total imports from 1987 to 2004. The foreign cost variable is the import-weighted foreign CPI, with the same weights as those for the exchange rate. The exchange rate and CPI data come from International Financial Statistics. The trade data come from the United States International Trade Commission and the Bureau of Economic Analysis, U.S. Department of Commerce. The regressions also control for U.S. domestic demand with U.S. Gross Domestic Purchases from the Bureau of Economic Analysis. Starred coefficients are significant at *10-, **5-, and ***1-percent level.

B A Comparison of Exchange Rates

Figure B.1: *A comparison of three exchange-rate indexes.* The auto-industry index includes the bilateral exchange rates of major importers of autos and auto parts each weighted by its average import share in the NAICS 3-digit industry 336 “Transportation Equipment”. The broad index comes from the Federal Reserve and is a weighted average of the foreign exchange values of the U.S. dollar against the currencies of a large group of major U.S. trading partners each weighted by its trade share. The yen-dollar index is the bilateral exchange rate between Japan and the United States over the sample period. Sources: Federal Reserve Board, authors’ calculations.

C Derivation of Matrix S_t^s

Let S_t^s be a matrix with element $\frac{\partial s_{kt}(p_i^r(p_i^a(p_i^s)))}{\partial p_{jt}^s}$ which is the change in each product's market share with respect to a change in each product's composite supplier price. The S_t^s matrix is a transformation of the manufacturer's demand substitution matrix S_t^a that was previously defined: $S_t^s = S_t^{m'} S_t^a$ where S_t^m is a J -by- J matrix of the partial derivative of each OEM price with respect to each product's TOPS price. Each column of S_t^m contains the entries of a response matrix computed as the average percent change in a product's *OEM* price for a given change in its composite supplier price. To obtain expressions for this matrix, we follow Villas-Boas (2007) and use the implicit-function theorem to totally differentiate the manufacturer's first-order condition for product j with respect to all manufacturer prices (dp_k^a , $k = 1, \dots, N$) and with respect to the composite supplier's price p_f^s with variation dp_f^s :

$$\sum_{k=1}^N \underbrace{\left(\frac{\partial s_j}{\partial p_k^a} + \sum_{i=1}^N \left(T^a(i, j) \frac{\partial^2 s_i}{\partial p_j^a \partial p_k^a} (p_i^a - p_i^s - ntc_i^a - ntc_i^s - tc_i^s) \right) + T^a(k, j) \frac{\partial s_k}{\partial p_j^a} \right)}_{v(j, k)} dp_k^a - \underbrace{T^a(f, j) \frac{\partial s_f}{\partial p_j^a}}_{w(j, f)} dp_f^s \quad (C.1)$$

Let V be a matrix with general element $v(j, k)$ and W be an N -dimensional vector with general element $w(j, f)$. Then $V dp^a - W dp_f^s = 0$. One can solve for the derivatives of all manufacturer prices with respect to the tier-one supplier's price f for the f th column of S^m :

$$\frac{dp^a}{dp_f^s} = V^{-1} W_f.$$

Stacking the N columns together gives $S^m = V^{-1} W_f$. The (j th, k th) entry of S^m is then the partial derivative of the k th product's manufacturer price (and given our assumption of constant markups, of the k th product's retail price) with respect to the j th product's supplier price for that market. The retail prices in turn affect the k th product's retail market share, that is: $\sum_m \frac{\partial s_{kt}}{\partial p_{jt}^r} \frac{\partial p_{jt}^r}{\partial p_{jt}^a} \frac{\partial p_{mt}^a}{\partial p_{jt}^s}$ for $m = 1, 2, \dots, J$.

D Demand

Market demand is derived from a standard discrete-choice model of consumer behavior that follows the work of Berry (1994) and Berry, Levinsohn, and Pakes (1995). Our estimation approach follows Nevo (2001). As the notation is fairly standard the demand section refers to the above papers for details for interested readers and we include here in the appendix a more detailed description of the demand model, its estimation, and further discussion of several demand identification issues.

The demand specification starts with consumer i choosing to purchase one unit of good j if and only if the utility from consuming that good is as great as the utility from consuming any other good. Consumer utility depends on product characteristics and individual taste parameters. Product-level market shares are derived as the aggregate outcome of individual consumer decisions. All the parameters of the demand system are then estimated from product-level data, that is, from product prices, quantities, and characteristics.

Suppose we observe $t=1, \dots, T$ markets. Let the indirect utility for consumer i in consuming product j in market t take a linear form:

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

where ε_{ijt} is a mean-zero stochastic term. A consumer's utility from consuming a given product is a function 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:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Sigma v_i \quad (\text{D.2})$$

where 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. We assume that the distribution of consumers' unobserved characteristics is multivariate normal. Indirect utility can be redefined in terms of mean utility $\delta_{jt} = \beta x_{jt} - \alpha p_{jt} + \xi_{jt}$ and deviations (in vector notation) from that mean $\mu_{ijt} = [\Sigma v_i] * [p_{jt} \ x_{jt}]$:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt} \quad (\text{D.3})$$

Finally, consumers have the option of an outside good. Consumer i can choose not to purchase one of the products in the sample. The price of the outside good is assumed to

be set independently of the prices observed in the sample. The mean utility of the outside good is normalized to be zero and constant over markets. The indirect utility from choosing to consume the outside good is:

$$u_{i0t} = \xi_{0t} + \sigma_0 v_{i0} + \varepsilon_{i0t} \quad (\text{D.4})$$

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:

$$s_{jt} = \int_{\zeta \in A_j} P^*(d\zeta) \quad (\text{D.5})$$

where $P^*(d\zeta)$ is the density of consumer characteristics $\zeta = [\nu]$ 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 independently and identically distributed with a Type-I extreme-value distribution. For this model, the probability of individual i purchasing product j in market t is given by the multinomial logit expression:

$$s_{ijt} = \frac{e^{\delta_{jt}}}{1 + \sum_{k=1}^{J_t} e^{\delta_{kt}}} \quad (\text{D.6})$$

where δ_{jt} is the mean utility common to all consumers and J_t remains the total number of products in the market at time t .

In the full random-coefficients model, we assume ε_{ijt} is i.i.d with a Type-I extreme-value distribution but now allow heterogeneity in consumer preferences to enter through an additional term μ_{ijt} . This allows more general substitution patterns among products than is permitted under the restrictions of the multinomial logit model. The probability of individual i purchasing product j in market t must now be computed by simulation. This probability is given by computing the integral over the taste terms μ_{it} of the multinomial logit expression:

$$s_{jt} = \int_{\mu_{it}} \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} f(\mu_{it}) d\mu_{it} \quad (\text{D.7})$$

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

$$s_{jt} = \frac{1}{N} \sum_{i=1}^N \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} \quad (\text{D.8})$$

We take 20 draws, which is typical in this literature. Given these predicted market shares, we search for demand parameters that implicitly minimize the distance between these predicted market shares and the observed market shares using a generalized method-of-moments (GMM) procedure, as we discuss in further detail in the estimation section.

E Demand Estimation

We estimate the demand parameters by following the algorithm proposed by Berry (1994). This algorithm uses a nonlinear generalized-method-of-moments (GMM) procedure. The main step in the estimation is to construct a moment condition that interacts instrumental variables and a structural error term to form a nonlinear GMM estimator. Let θ signify the demand-side parameters to be estimated with θ_1 denoting the model's linear parameters and θ_2 its nonlinear parameters. We compute the structural error term as a function of the data and demand parameters by solving for the mean utility levels (across the individuals sampled) that solve the implicit system of equations:

$$s_t(x_t, p_t, \delta_t | \theta_2) = S_t \quad (\text{E.1})$$

where S_t are the observed market shares and $s_t(x_t, p_t, \delta_t | \theta_2)$ is the market-share function defined in equation (D.8). For the logit model, this is given by the difference between the log of a product's observed market share and the log of the outside good's observed market share: $\delta_{jt} = \log(S_{jt}) - \log(S_{ot})$. For the full random-coefficients model, it is computed by simulation.

As in Nevo (2000), to ensure a global minimum, we start by using a gradient method (providing an analytical gradient) with different starting values of the nonlinear parameters to find a minimum of the simulated GMM objective function. Then we use that minimum as a starting value for the Nelder-Mead (1965) simplex search method.

Following this inversion, one relates the recovered mean utility from consuming product j in market t to its price, p_{jt} , its constant observed and unobserved product characteristics, d_j , and the error term $\Delta\xi_{jt}$, which now contains changes in unobserved product characteristics:

$$\Delta\xi_{jt} = \delta_{jt} - \beta_j d_j - \alpha p_{jt} \quad (\text{E.2})$$

We use brand fixed-effects as product characteristics following Nevo (2001). The product fixed-effects d_j proxy for the observed characteristics term x_j in equation (D.1) and mean unobserved characteristics. The mean utility term here denotes the part of the indirect utility expression in equation (D.3) that does not vary across consumers. Equation (23) thus relates the mean utility, which is a transformation of market shares (which are themselves functions of the data and the nonlinear parameters such as the standard deviation of retail prices and horsepower/length) to the right-hand-side variables that are product fixed-effects and prices. This equation is linear in prices and so instrumental variables can be implemented directly using it. Its only endogenous right-hand-side variable

is the retail price and used prices are its instrument. More precisely, the right-hand-side variables in our linear-demand equation after the inversion are $X = [\text{price} \quad \text{product fixed-effects}]$, our first-stage right-hand-side variables (our instruments) are used prices, and all other variables are assumed to be exogenous other than the retail price. The instrument matrix is then $Z = [\text{used prices} \quad \text{product fixed-effects}]$ given that the product fixed-effects are exogenous in our linear-demand equation. Therefore we have as many columns in Z as in X , the same number of instruments as regressors.

Note that we must restrict the number of models in our sample to use this estimation approach in our structural model. A natural question is why we don't use the instrumental variables strategy of BLP (1995) which would allow us to include many more models in the sample. The characteristics of our data make the use of BLP's identification strategy infeasible. Our *Edmunds* dataset contains monthly observations for their TMV prices and model characteristics over five years while BLP's dataset contains annual observations for MSRP's and model characteristics over twenty years. The characteristics of automobile models generally change annually. This means the time-series variation of characteristics and prices is consistent in the BLP data, but not in our data, where we observe considerable variation in monthly TMV new and used prices but almost no variation in monthly characteristics. For this reason, using BLP instruments (the sum of characteristics across own and rival makers) would be akin to using product-year fixed effects in our demand estimation. If we adopted this strategy, we would need to restrict the number of models in the sample due to the limited degrees of freedom afforded by our sample length. In addition, using the BLP approach would require us to assume that unobserved changes in product characteristics that stem from monthly promotional activity are unrelated to constant product-specific demand shifters, which may not be the case for automobiles. A new design for a bumper, for example, may be tied to more monthly promotions over the course of a model year. For these reasons, we choose to use product fixed effects in our demand estimation (to control for constant product-specific demand shifters) and employ used prices to instrument for the *Edmunds* TMV prices, as these should not be correlated with monthly unobserved changes in product characteristics, that for example, stem from promotional activity. This identification strategy requires us to limit the number of models in the sample to maintain sufficient degrees of freedom in our demand estimation.