The Evolution of Price Dispersion in the European Car Market*

Pinelopi Koujiani Goldberg
Yale University and NBER

Frank Verboven
Catholic University of Leuven and CEPR

Revised July 2000

Abstract

Car prices in Europe are characterized by large and persistent differences across countries. The purpose of this paper is to document and explain this price dispersion. Using a panel data set extending from 1980 to 1993, we first demonstrate two main facts concerning car prices in Europe: (1) The existence of significant differences in quality adjusted prices across countries, with Italy and the U.K. systematically representing the most expensive markets. (2) Substantial year-to-year volatility that is to a large extent accounted for by exchange rate fluctuations and the incomplete response of local currency prices to these fluctuations. These facts are analyzed within the framework of a multiproduct oligopoly model with product differentiation. The model identifies three potential sources for the international price differences: price elasticities generating differences in markups, costs, and import quota constraints. Local currency price stability can be attributed either to the presence of a local component in marginal costs, or to markup adjustment that is correlated with exchange rate volatility; the latter requires that the perceived elasticity of demand is increasing in price. We find that the primary reason for the higher prices in Italy is the existence of a strong bias for domestic brands that generates high markups for the domestic firm (Fiat). In the U.K. higher prices are mainly attributed to better equipped cars and/or differences in the dealer discount practices. The import quota constraints are found to have a significant impact on Japanese car prices in Italy, France and the U.K.. With respect to local currency price stability, a large percentage of the documented price inertia can be attributed to local costs, and a smaller fraction to markup adjustment that is indicative of price discrimination. Based on these results we conjecture that the EMU will substantially reduce the year-to-year volatility observed in the car price data, but without further measures to increase European integration, it will not completely eliminate existing cross-country price differences.

*We gratefully acknowledge funding from a National Science Foundation Grant to the NBER (Grant SBR-9731979), an Alfred P. Sloan Foundation Research Fellowship to Goldberg, and a grant by the Belgian Science Foundation (Fonds Wetenschappelijk Onderzoek, 774810/8/F01/6260/131200) to Verboven. We also thank the Center for Economic Research (Tilburg University) for providing financial and organizational support in collecting the data. This version has benefited substantially from the comments of the editor and the anonymous referees. We are of course responsible for any remaining errors.

1. Introduction

Car prices in Europe are characterized by large and persistent differences across countries. Figures 1a and 1b plot for several car models the maximum percentage and absolute bilateral price difference observed across five distinct markets (Belgium, France, Germany, Italy and the United Kingdom) against the average across the five markets price of the model; all prices are pre-tax, and expressed in a common currency, ECUs. The price dispersion appears to be enormous; on average, the maximum price differential is around 30% of the car price, implying that the price of a 10,000 ECU (approximately $10,000) car can vary by as much as 3,000 ECUs (or $3,000) between two countries. The percentage differences are fairly constant in the price of a car, so that the absolute price differences are the largest for the most expensive models.[1] Moreover, the raw data plots suggest no tendency at all for the price differences to decline over time; in 1993, the price dispersion is at least as pronounced as in 1980.

These price differences have been the focus of intense public debate in Europe. Consumer organizations argue that the price dispersion is the result of geographical market segmentation, allowing for anti-competitive price discriminating practices; such organizations closely monitor the evolution of price differences, and prefer those to be small; see, for example, the reports of the Bureau of European Consumers Unions (BEUC) in 1981, 1986, 1989, 1992. Industry insiders, on the other hand, defend the high cross-country price differences as the result of exchange rate fluctuations and tax policies. Empirically, the price dispersion in the European car market coincides with both large exchange rate movements and substantial, evolving cross-country differences in value-added taxes. Proponents of this view consider the currency union to be the solution to the problem. "You can only have price harmony if you have a common currency", is an industry representative quoted in the New York Times ((New York Times, Jan. 17, 1996). "Should the single currency arrive", continue the New York Times, "... Europe is expected to more closely resemble the United States. Prices would vary slightly region by region, reflecting income differences, shipping costs and other variables, but there would be fewer sharp differences". This view, however, contrasts with the findings of empirical studies that have attributed the price dispersion to differences in demand elasticities across national markets (Mertens and Ginsburgh (1985)), differences in concentration (Mertens and Ginsburgh (1985), Flam and Nordstrom (1994)), lack of uniform taxation on the value added (Gual (1993)), and differences in import quota constraints (de Melo and Messerlin (1988), Gual (1993), Flam and Nordstrom (1994)). It also contrasts with the manufacturers' active efforts to keep geographical market segmentation by maintaining the selective distribution system and preventing their

[1] Graphs of the standard deviation and coefficient of variation of car prices display similar patterns.

dealers from selling to foreign customers (see the discussion in the next section).

While the source of price disparities may be a source of disagreement among auto experts, consensus seems to exist that this is an important issue in the European Community. The Commission of the European Communities also views price differentials as a potential threat to the European market integration policies. Since 1992, it has conducted price investigations at a bi-annual basis to monitor the evolution of international price differences. The Commission has repeatedly found price differences far in excess of the 12% that were allowed in return for granting manufacturers a restrictive system of selective and exclusive distribution. Against this background, it seems particularly important to understand the sources of price dispersion in the European markets.

This paper has two goals; describe the pattern of price dispersion in five major European markets in the 1980's and early 1990's to obtain the basic "facts" that demand explanation; and identify the sources of cross- country price differentials within the framework of an oligopoly model. Given the large number of existing studies on price dispersion in the European car market, it is perhaps useful to point out which features of our analysis distinguish the current work from previous research. We believe that our approach improves on earlier work in three respects. First, in order to adjust for product quality differences that would potentially explain cross-country price differences, we constructed a major database containing extremely detailed disaggregate information on prices, sales, characteristics and production location of approximately 300 models sold in the five markets under consideration in the period 1980-1993. In addition to allowing us to control for quality differences, the disaggregate data offer the advantage of enabling us to conduct the analysis at the firm level, thus avoiding controversial aggregation assumptions. Second, because our data set has a relatively long time dimension (14 years), we can explicitly address the effects of exchange rate volatility on price dispersion; previous studies often based their conclusions on a single year cross-section, thus abstracting from the role of exchange rates. The third advance is that we analyze the data systematically within the framework of an oligopoly model with product differentiation. This framework is needed to conduct counterfactual policy simulations. In addition, it enables us to ultimately address the question whether price differences are the result of cost differences or price discrimination; and if the latter is the case, further investigate what the sources of price discrimination are. One interpretation difficulty that often arises in the context of earlier studies is that while these studies unambiguously establish correlations between exchange rates, tax differences, import quotas and cross-country price differentials, it is not clear whether these correlations are indicative of cross-country cost differences or price discrimination. This problem is particularly acute in the context of exchange rates. Because price dispersion is empirically highly correlated with nominal exchange rate volatility, price differences are often "justified" as being due to the high "costs" associated with an overvalued currency. But the response of prices to exchange rates (or the lack thereof) may itself reflect markup adjustment and be therefore indicative of price discrimination.

We start our analysis by documenting the existence of price dispersion during our sample period. Through a series of reduced form regressions we investigate whether quality adjusted price differences are systematically related to factors such as exchange rate fluctuations and tax differences. The reduced form estimation combines hedonic regression techniques with the pricing-to-market literature. We next turn into a more systematic investigation of price differences. To this end, we develop and estimate a model of oligopoly with product differentiation that incorporates exchange rate fluctuations and tax differences. This framework allows us to decompose the equilibrium price of each vehicle make into two components: its marginal cost and a markup. The markup depends on the demand side of the market (own- and cross-price elasticities), the extent of collusion, and restrictions on demand, such as import quota constraints. Our approach allows us not only to estimate the markup, but also decompose it into its determinants to understand the sources of price discrimination. We are particularly interested in identifying the role of local competition in explaining the cross-country price differentials. Are higher prices in some markets generated by strong preferences for domestic brands?

To evaluate the claim that price differences are generated by nominal exchange rate volatility, particular attention is paid to the role of exchange rates in price determination. Theoretical work in the area has shown that the response of prices to currency fluctuations can be related to the curvature of the demand and cost schedules. This has often been viewed as unsettling (see e.g. Krugman, 1987), since it implies a strong dependence of the results on particular functional form assumptions. Rather than imposing such assumptions a-priori, we estimate relatively flexible demand and cost functions, and then examine whether our parameter estimates are consistent with the observed price response to exchange rate fluctuations. In addition, we investigate the hypothesis that the incomplete response of local currency prices to exchange rate fluctuations is due to the existence of a local component in marginal cost.

2. A First Look at the Market and the Data

2.1 The Data Set and some Preliminary Descriptive Results

The data set we have constructed to analyze price dispersion has three dimensions: (1)

Products: In each year there are approximately 150 models in our sample. For each model we have information on sales, list price, and physical characteristics of the base specification. These include engine attributes (horsepower, displacement), dimensions (weight, length, width, height) and performance variables (fuel consumption at 3 levels of speed, acceleration time, maximum speed); they often vary across markets. All data are from publicly available sources.[2] (2) Time: Our sample extends from 1980-1993. (3) Markets: There are five markets in our sample: Belgium, France, Germany, Italy and the United Kingdom. We focus our attention on these five countries, both because of data availability constraints, and, more importantly, because they represent the largest markets in Europe: collectively they account for over 85% of total car sales in Europe every year. In addition, these countries represent a large spectrum for several reasons: the size of the market varies from ca. 400,000 units per year in Belgium to almost 3 million cars in Germany; the degree of import penetration ranges from ca. 30% in France and Germany to almost 100% in Belgium; the Japanese penetration varies from ca. 1% in Italy to 20% in Belgium; tax rates vary from 14% in Germany to 33% in France in the early years, and 25% later; and the C1-concentration index ranges from 53% in Italy to 16% in Belgium.

In addition, our database contains information on the production location of each model (source: Pemberton Associates); brand ownership and class; average and maximum dealer discounts for selected years (source: BEUC reports, and unpublished interviews by CECRA, the European Committee for Motor Trades and Repairs in Brussels);[3] exchange rates (source: IFS statistics); tax rates (source: the retail catalogues mentioned in the last footnote); and income distribution in each market (source: Atkinson, 1997). Some summary statistics are provided in Table 1.

A natural question that arises in the context of the European car market is why the market is segmented along national lines. The perhaps most important obstacle to cross-border trade stems from the system of selective and exclusive distribution resting on Regulation 123/85 of the European Commission. Selectivity means that the manufacturer can choose his/her dealers

[2]Price and characteristics data are available from retail catalogues, i.e., De Autogids, Auto Moto Revue, Journal de l'Automobile, Katalog der Automobil Revue, Adac Auto Special, What Car?, and Quattroruote. Sales data are the number of national registrations, and come from Nieuwe tot het Autoverkeer Toegelaten Voertuigen, l' Argus de l'Automobile et des Locomotions, Automobil Revue, MVRIS, Notziario Statistica, Tatsachen und Zahlen aus der Kraftsverkehrswirtschaft, and World Motor Vehicle Data.

[3]The data on dealer discounts allow us to compute transactions prices and use those in the estimation of the demand side of the model. Alternatively, we can employ list prices in the estimation and treat dealer discounts as part of model fixed effects (note that we only have information on dealer discounts for selected years). In the estimation of the demand system we experimented with both alternatives with no impact on the results of interest – only the intepretation of the model fixed effects in the demand estimation is affected.

and restrain them from reselling to anyone but end-users or approved sellers. Many manufacturers have in the past instructed their dealers (threatening to withdraw their concessions) not to sell to independent resellers, in particular if the purchase was intended for export. Discrimination against resellers occurred in several subtle forms: excessive delivery lags, high deposit requirements, reservations to provide guarantee outside the country of purchase, and higher prices. These practices were effective in raising the consumer cost of purchasing a car abroad.

The European Commission has threatened several times to withdraw the benefits of the Regulation if prices between member states remained significantly different over time. The Commission's stand is also known as the 12% rule: the maximally allowed price differences between any two member states (excluding the high tax countries Danmark and Greece) should not exceed 12% during a period of 6 months, and 18% at any point in time. But despite the careful monitoring of price differences, the benefits from the selective distribution system had – until recently – never been withdrawn upon observing excessive price differences. In recent years, however, the Commission has become more serious about the 12% rule.[4]

Two additional impediments to consumers seeking to purchase a car abroad are the existence of national type approval rules (these were finally harmonized in 1995), and the system of national registration. In combination with restrictive trade policies the requirement of national registration has had the effect of limiting trade of foreign, mainly Japanese models. Quantitative restrictions on imports from third countries, in particular Japan, have long existed in various European countries (France, Italy, Portugal, Spain and the U.K.). These restrictions take the form of import quotas or voluntary export restraints. The problem is, of course, that parallel imports from other European countries can undo the national restrictions. The requirement of national registration resolves this problem, since it can control cross-border trade of Japanese cars.

The best evidence for the obstacles to cross-border trade within the European Community is perhaps given by the magnitude of parallel imports, the goods imported by unauthorized resellers. Table 1 summarizes the evidence collected from various BEUC surveys. It reveals that parallel imports have been quite low in all Europeans countries of our study. This is remarkable given the very large cross-country price differences as plotted in the introduction.

[4]This is evident from the 1998 Volkswagen case. Volkswagen has been accused and convicted for pressuring Italian dealers not to sell to German and Italian customers. Proof has been found that these practices had been going on for ten years. They involved threats to 50 dealers to withdraw their licenses, with 12 licenses effectively withdrawn. The conviction included a 102 million ECU fine (about 10 percent of Volkswagen's annual profits), the largest fine ever issued to a single firm in Europe, and the removal of Volkswagen's rights as set out in the Regulation. Similar investigations are under way against Mercedes and Opel.

Other strong evidence on trade restrictions within the European Community is given by the Japanese market shares. These differ drastically across countries. In principle, this could be due to differences in local tastes. The close correspondence to the allocated national quotas, however, suggests that countries have been very successful in preventing intra-European trade of Japanese models; in countries with a quota or voluntary export restraint, the actual Japanese market share only slightly exceeds the assigned quota.

Table 1 also presents the market shares of "domestic" firms, i.e., the market shares of the models that are produced domestically. Domestic firms have high and stable market shares in Germany and France, and to a lower degree also in Italy and the United Kingdom. In all countries (except Belgium), the market share of domestic firms is more than twice the average over the five countries. Especially in Italy and the United Kingdom there is a very strong presence of domestic firms relative to the European average. These national consumption patterns may stem from differences in local tastes (possibly including differences in the sizes of dealer networks), or be the result of trade barriers within Europe.

2.2 Quality Adjusted Price Differences

To investigate whether the price dispersion documented in Figure 1 is random or systematic, we start by constructing hedonic price indices for each market and each year in our sample. These indices control for quality differences as measured by the observed physical characteristics of a car. The basic hedonic price equation we estimate is:

$$ln(p_{jmt}^{ECU}) = w_{jmt}\gamma + \delta\, tax_{mt} + \theta_c + \theta_f + \theta_{st} + \theta_{mt} + \epsilon_{jmt} \qquad (1)$$

The subscripts j, m and t refer to product j, market m and year t respectively. The vector w_{jmt} consists of physical car characteristics (horsepower, size, etc.) that may vary across markets, while θ_c and θ_f are market segment and firm dummies respectively. Price differences due to cross-country tax differences are captured in the regression through tax (the log of the value added tax). In addition, we include a set of source country/time dummies (θ_{st}) to control for differences that may be due to a common cost shock facing firms located in a particular country of origin (e.g. an increase in wages facing all Japanese firms). Given this specification, the destination/time effects θ_{mt} capture the residual cross-country price differences that cannot be explained by differences in quality or taxation across markets. All differences are measured in percentage terms relative to Belgium.

The results from estimation of equation (1) are summarized in Table 2. All parameters have intuitive signs and are precisely estimated. Given the large number of included destination/time effects we do not report them individually; however, at the bottom of the table, we report the

95% confidence interval for the destination effects in 5 distinct years: 1980, 1985, 1990, 1992 and 1993. The point estimates of the destination/time effects are plotted in Figure 2.

The first pattern that emerges from Figure 2 (and Table 2) is that there are large, persistent and statistically significant cross-country price differences, even after we adjust for quality and taxation differences. During almost the entire sample period, France, Germany, Italy and the U.K. are more expensive than Belgium; the only exception to this pattern is 1993, when the destination effects for the U.K. and Italy become statistically insignificant. The U.K. is associated with the highest prices throughout the sample period, followed by Italy. German prices are 3%-5% higher than the ones in Belgium during the entire period. If we were to characterize price dispersion during this period in a few words, we would group the five countries in our sample into three categories: Belgium (the cheapest), France and Germany, and Italy and the U.K. (the most expensive).

The second pattern that emerges from Figure 2 is that the hedonic price indices are extremely volatile. Even though the ranking of the countries remains fairly stable, year to year changes in the destination effects are fairly large. The perhaps most dramatic changes are in 1993, when both the U.K. and Italy (the relatively more expensive countries in the sample) drop below Germany and France.

Given this volatility and the 1993 effects in particular, it is compelling to relate the movements in the price differences to exchange rates. Figure 3 plots the exchange rates of France, Germany, Italy and the U.K. vis a vis Belgium. It appears that the movements in the hedonic price indices plotted in Figure 2 closely track the exchange rate movements in Figure 3. The high correlation between exchange rate and price index movements is even more pronounced in Figure 4 that plots the annual changes in price differences and exchange rates; the year-to-year changes in the two variables are highly correlated. This is indicative of the short run effects of exchange rates on price differentials; because local currency prices do not fully respond to a currency appreciation (depreciation) in the short run, prices denominated in ECUs will naturally be correlated with exchange rates.

2.3 The Role of Exchange Rates

The hedonic regression in column 2 of Table 2 controls only for changes in the exchange rate of each destination market relative to Belgium. In each market, however, there are several producers of different nationalities, each of which faces a distinct exchange rate vis a vis the destination market. To take this into account we estimate an equation of the following general form:

$$ln(p_{jmt}^{exp}) = w_{jmt}\gamma + \delta \, tax_{mt} + \theta_c + \theta_f + \theta_{st} + \theta_{sm} + \beta_{sm} ln(e_{smt}) + \epsilon_{jmt} \qquad (2)$$

This specification is similar to the hedonic pricing equation estimated before. There are three main differences: First, the left hand side variable is now expressed in exporter currency units; second, the destination market/time effects are replaced by the log of the exchange rate of each source country vis a vis the destination market (i.e., units of source currency per unit of destination currency); third, we now introduce source country/destination market fixed effects (θ_{sm}). The coefficient β_{sm} is called the pricing-to-market (or exchange rate pass-through) coefficient. If β_{sm} is equal to 0, local prices fully respond to exchange rate changes; prices remain unchanged in exporter currency units, and pass-through is characterized as complete. If, at the other extreme, local prices are completely unresponsive to exchange rate changes, β_{sm} is equal to 1. In this case, exporters fully adjust their prices in order to absorb the exchange rate change. Values of β_{sm} between 0 and 1 indicate incomplete exchange rate pass-through. Given that everything is specified in levels (and not first differences), the coefficient β_{sm} can be thought of as capturing the long term response of prices to exchange rate fluctuations. The specification in (2) is very similar to the equations estimated in the pricing-to-market literature, most notably by Knetter (1989, 1993); it has, under certain more or less plausible conditions, a semi-structural interpretation.

We start by estimating a version of (2) in which the same pricing-to-market coefficient is imposed for all source countries and all destinations ($\beta_{sm} = \beta$).The average pricing-to-market coefficient (β) is estimated at 0.46 (standard error 0.02). In the absence of a local component in marginal costs, these results imply that the price dispersion documented in the five European markets in our sample is due to price discrimination, that is highly correlated with exchange rate variation. In particular, the point estimate of 0.46 for the pricing-to-market coefficient implies that, on average, auto producers absorb 46% of the exchange rate fluctuations through markup adjustment, while local prices remain relatively stable.

Next, equation (2) is estimated allowing the pricing-to-market coefficients to vary across source countries and destinations. The results from this regression (not reported here) indicate substantial variation of the coefficient estimates across destinations and source countries. Rather than listing all β_{sm} estimates below, we report results from a more parsimonious specification in which β_{sm} was decomposed additively into a destination market and a source country component ($\beta_{sm} = \beta_s + \beta_m$). The results are listed in Table 3.

The first row reports the coefficient estimates for each source country (standard errors in parenthesis). Three out of the eleven coefficients are not statistically significant, while the point estimates are relatively low. These correspond to Belgium, Korea and the Netherlands. The observations corresponding to these countries are two few (imports from these three countries represent an almost negligible fraction of total sales in Europe) to allow a precise estimation of the coefficients. Among the other seven source countries, Spain, Sweden and the U.K. exhibit the highest coefficients, while Germany and Japan the lowest. Overall there seems to be substantial variation in the coefficient estimates across source countries. This variation is contrasted with the average (across the five markets and all sample years) market shares of each source country in row 2 of the table. It appears that exchange rate pass-through is inversely related to market shares; countries with a relatively strong presence in the destination markets (e.g., Germany) are more likely to adjust the local prices in response to exchange rates, while countries with very small shares (such as Spain and Sweden) absorb the exchange rate changes by adjusting their markups, leaving local prices relatively stable.

Row 3 of Table 3 reports the market specific components of the pricing-to-market coefficients. The only coefficient that is statistically significant is the one for the U.K.; the associated point estimate is also quite large. This is consistent with the pattern revealed in Figure 2; the U.K. is the most expensive destination for almost the entire sample period, and it is also the one with the highest degree of local price stability.

Overall, the pricing-to-market regressions suggest that local price stability is an important feature of European auto markets during periods of exchange rate volatility. Note that both the frequency of our data (annual) and the nature of the regressions suggest that this is not a stability that pertains only to the short run, due to lagged adjustment; it persists in the long run, and generates significant price differences across countries. The strong correlation between market shares and exchange rate pass-through suggests that local price stability may be related to the competitive conditions in each market.

In summary, the preliminary data analysis points to two main features of price dispersion in Europe: (1) Existence of substantial and persistent differences in quality adjusted prices across countries, and (2) Substantial year-to-year volatility. A large fraction of this volatility (but not all of it) is accounted for by exchange rate fluctuations. An analysis of the effects of exchange rates on prices within the pricing-to-market framework indicates that local prices remain relatively stable during periods of exchange rate realignments. This local price stability could in turn be generated by local production and selling costs and/or third degree price discrimination. The remainder of our analysis will be devoted to explaining the facts revealed in the descriptive analysis, and quantifying the contributions of the various components responsible for price dispersion.

3. The Empirical Model

3.1 Demand

The demand side of the market is modelled in a discrete choice framework. Because of our interest in the demand elasticities, it is particularly important to adopt a reasonable and flexible specification for the utility function. Ideally, our estimation approach would exploit data at the individual consumer level; such data would allow us to introduce consumer heterogeneity and its interactions with product characteristics in a very flexible manner. Unfortunately, such micro data are not available for any European country; we only have information on the income distribution in each country. We therefore employ product level data, but adopt a specification that allows us to at least exploit the information on the income distribution.

Specifically, we experimented with two specifications based on McFadden's (1978) generalized extreme value (GEV) model: the bi-level nested logit, and a weighted average of two one-level nested logit models, as in Bresnahan et al (1997). In both cases we utilized insights from the marketing literature suggesting that the auto market is differentiated along two dimensions: market segment or "class", and country of origin (domestic vs. foreign). However, the model proposed by Bresnahan et al did not find support in the data. Our results and discussion in section 5 are accordingly based on the bi-level nested model alone. Since Bresnahan et al's specification of the GEV model provides an attractive, and computationally feasible alternative to the bi-level nested logit model, we briefly discuss – in the spirit of sensitivity analysis – its main features and its implications for the substitution patterns in the Appendix.

Consider consumer i, in market m at time t. The consumer faces $J_{mt}+1$ alternatives: J car models offered at time t in market m, plus the option not to buy a new car ($j=0$). Assuming utility maximization, McFadden's (1978) GEV model gives rise to the following familiar expression for the probability Pij of buying car j (in the following we suppress the indices m and t):

$$P_{ij} = \frac{e^{V_{ij}} G_j(e^{V_{i0}}, ..., e^{V_{iJ}})}{G(e^{V_{i0}}, ..., e^{V_{iJ}})} \tag{3}$$

where, for the bi-level nested logit, the function G is given by:

$$G(e^{V_i}) = e^{V_{i0}} + \sum_{k=1}^{K} \left\{ \left[\left(\sum_{j \in D, C_k} e^{\frac{V_{ij}}{\rho_f}} \right)^{\frac{\rho_f}{\rho_c}} + \left(\sum_{j \in F, C_k} e^{\frac{V_{ij}}{\rho_f}} \right)^{\frac{\rho_f}{\rho_c}} \right]^{\rho_c} \right\} \tag{4}$$

The notation is as follows: C_k denotes class k (subcompact, compact, etc.); F stands for foreign products; D stands for domestic products; G_j is the partial derivative of G with respect to $e^{V_{ij}}$;

and V denotes the deterministic part of the utility function ($Uij = V_{ij} + \epsilon_{ij}$). The deterministic part is in turn given by:

$$V_{ij} = \delta_j + \alpha \, ln(y_i - p_j) \tag{5}$$

The term y_i denotes income, and p_j is the price of car j. The first term, δ_j, captures the mean evaluation for good j, common to all consumers. It can be written as:

$$\delta_j = x_j\beta + \xi_j \tag{6}$$

The vector x_j consists of observable vehicle characteristics (size, horsepower, cylinders, options, etc.), while ξ_j captures unobserved quality. In the actual estimation of this demand system we allow all parameters to vary across markets. Alternatively, one could impose the same parameters, and hence the same preference structure, across countries. Since we want to impose the minimum amount of restrictions at this stage, we adopt the first specification and examine later whether the equality restrictions are supported by the data.

Of special interest for the interpretation of our results are the two distributional parameters of the nested logit, the ρ_c and ρ_f.The first one, ρ_c, captures the pattern of dependency across products in the same class; the second one, ρ_f, parameterizes the substitutability across products of either domestic or foreign origin. To be consistent with random utility maximization (see McFadden (1978)), both ρ_c and ρ_f have to lie in the unit interval; in addition, ρ_f has to be less than, or equal to ρ_c. The interpretation of the distributional parameters is as follows: As either ρ decreases, the dependency across products in the corresponding cluster becomes stronger; as either ρ goes to 0, products in the corresponding cluster become perfect substitutes. Conversely, as ρ goes to 1, the dependency becomes weaker; in the limit case of $\rho = 1$, the error terms become independent within the cluster, and the model reduces to a one-level nested logit. Similarly, if $\rho_c = \rho_f$, the model reduces to a single level nested logit.

Between the different values that the parameters ρ_c and ρ_f can obtain and the two GEV specifications we considered, we have covered, we believe, a wide range of possible substitution patterns. While further generalizations of the demand structure are possible (see, for example, Berry et al (1995)), the specified model offers a reasonable compromise between functional form flexibility and computational tractability. The latter is high in our priority list, as it allows us to better exploit the richness of our data set by experimenting with different specifications.

Because of the unavailability of consumer level information, our expression for P_{ij}, the probability that consumer i buys product j, has to be aggregated up to the product market share function before the model can be taken to the data. While aggregating, we take advantage of the information on income distribution in each country. This information is available to us as a Lorenz curve; that is, we know what percentage of the population in each country has what percentage of income. We use this information in conjunction with aggregate population and personal disposable income data to (1) define 10 equally sized income classes (deciles), and (2) compute the per capita income of consumers within each income class. We then compute the purchase probabilities for the "average" consumer in each income class and sum up these probabilities to generate the market shares in each country. The advantage of including the data on income distribution is that we can account for cross-country differences in prices and market shares that arise from differences in the distribution of income across countries.

3.2 Marginal Cost

The marginal cost of each vehicle make is treated as unobservable, and modelled parametrically as a function of the car's physical characteristics, factor prices and total production quantity. Specifically, we adopt the following Cobb Douglas specification for marginal cost of product j in market m at time t:

$$\frac{\partial C_{jmt}}{\partial q_{jmt}} = \exp(z_{jmt}\gamma_s + \omega_s + \omega_f + \omega_m + \omega_t + \omega_{jmt}) W_{st}^{\delta} F_{st}^{1-\delta} Q_{jt}^{\zeta} \tag{7}$$

Marginal cost is expressed in the currency of the production location indexed by s. It is homogeneous of degree one in wages W_{st} and other factor prices F_{st}. The quantity Q_{jt} refers to the total European sales of product j. By including it in our specification we allow for non-constant returns to scale: a positive (negative) coefficient indicates decreasing (increasing) returns to scale.[5]

The vector z_{jmt} denotes the physical characteristics of product j. Note that this formulation allows model specifications to vary across destination markets, so that we can capture cross-country cost differences arising from differences in quality; similarly, it controls for changes in quality over time. The parameter vector γ can again be source country specific (as the notation in (7) implies), or constrained to be the same across production locations ($\gamma_s = \gamma$). We experiment with both specifications. The error term ω may be interpreted as capturing unobserved characteristics influencing marginal cost. It contains fixed effects for firms and production locations (ω_f and ω_s respectively), which may arise from differences in efficiency across firms and

[5] Note that we use the total *European* and not total *world* sales as our measure of Q_{jt}. We thus extend specifications used in previous papers that use only the local sales in each country. However, our quantity variable still misses sales in other non-European markets. This is most likely relevant for Japanese car models (for European cars, European sales are close to world sales, and American cars sold in Europe are manufactured in different facilities than American cars sold in the States). Implicit in our specification is the assumption that producers maximize profits separately for each region, and hence European markets. This assumption is not unreasonable given that different models are sold in different markets, and even the same model is often sold in different markets under different names. Note that this issue would have been irrelevant if we assumed constant returns to scale. However, we did not want to impose this assumption a-priori.

locations. In addition, it contains a market effect (ω_m), which captures unmeasured factors that influence the marginal cost of all cars sold in a particular destination market. These may arise from different local regulations or other institutional factors. For example, the required use of catalytic converters in Germany would increase the marginal costs of all cars in this country; in the U.K., it is generally believed that cars are better equipped than in other European countries, because of the importance of the leasing market.

Implicit in the denomination of marginal costs in source country currency is the assumption that all production costs occur in the source country. This assumption may be inappropriate if some inputs have been purchased abroad, or a certain component of a product's costs (mainly distribution and selling costs) is local to the destination market. To assess the importance of such a local component we also consider a specification in which the log of the local wages in the destination market is included on the right hand side of (7); one would generally expect labor costs to constitute a large fraction of local distribution and selling costs. The local wages W_{mt} are expressed in the currency of the destination market; in order to express everything in a common currency, we multiply them by the exchange rate between the firm's source country s and the destination market m (the exchange rate e_{smt} is expressed in units of source country currency per unit of destination market currency), and obtain a modified version of (7):

$$\frac{\partial C_{jmt}}{\partial q_{jmt}} = \exp(z_{jmt}\gamma_s + \omega_s + \omega_f + \omega_m + \omega_t + \omega_{jmt}) W_{st}^{\delta} (e_{smt} W_{mt})^{\theta} F_{st}^{1-\delta-\theta} Q_{jt}^{\zeta} \tag{8}$$

If 100% of the production costs occur in the source country, the local cost coefficient θ should be zero.

3.3 Oligopoly Pricing

The European automobile market is modelled as an oligopoly with multiproduct firms. In each market m at time t, firm f sells a subset F_{fmt} of the J_{mt} car models sold in this market/year. The sales of each car model in market m, $q_{jmt}(p_{mt})$, are given by the product of the market share of j and the number of potential consumers L_{mt}:

$$q_{jmt}(p_{mt}) = s_{jmt}(p_{mt}) L_{mt}.$$

The determination of the market share was already discussed in the subsection on demand. Assuming prohibitive arbitrage costs to consumers, this share will be a function only of the price vector in market m at time t, p_{mt}, and not of the price vectors in other markets. As the evidence presented in section 2 demonstrates, the assumption of prohibitive arbitrage costs is not unreasonable. The number of potential consumers is assumed to be the number of households in each market.

The equilibrium concept is Nash. We start by considering the case of Bertrand competition. The assumption of price setting is common, and consistent with industry wisdom. The framework we adopt is however flexible enough to also accommodate the case of collusion. Collusion may be present in the European markets because of the existence of quantitative import restrictions; as the trade literature has shown, these can often serve as collusive devices among producers.

The relevant price for the firm's profit maximization problem is the wholesale price p^w_{jmt}. It is assumed this is proportional to the observed list price according to:

$$p^w_{jmt} = p_{jmt}/[(1 + t_{jmt})(1 + \tau_{jmt})] \tag{9}$$

where t_{jmt} denotes value added taxes in market m, and τ_{jmt} refers to the dealer markup. This is treated as exogenous in our framework. In France, Germany and the U.K. value added taxes are the same percentage for all cars. In Belgium and Italy, this percentage is model-specific: cars with a powerful engine have a higher value added tax.
Each firm maximizes its profits over the M markets at period t, as given by:

$$\Pi_{ft} = \sum_m \sum_{j \in F_{fmt}} e_{fmt} p^w_{jmt} q_{jmt}(p_{mt}) - \sum_{j \in F_{fmt}} e_{fst} C_{jt}(q_{j1t}(p_{1t}), ..., q_{jMt}(p_{Mt}))$$

The wholesale prices p^w_{jmt} are expressed in local destination currency. $C_{jt}(.)$ refers to the production cost of model j, expressed in the currency of the production location. The term e_{fmt} denotes the exchange rate (at time t) between the firm f's country of registration and the destination market (i.e. units of f's currency per unit of m's currency), while e_{fst} is the exchange rate between the firm's registration country and model j's production location. This formulation allows us to account for the diversified activities of multinational enterprises, and capture the effects of exchange rate changes on prices.

Cars imported from Japan face quota constraints. In some countries (e.g. Italy), the constraint is an absolute upper limit on imports, i.e. $\sum_{j \in S_{smt}} q_{jmt}(p_{mt}) \square D_{smt}$, where S_{smt} is the subset of models produced in location s for import into market m. In other countries (France, Germany, and the U.K.) the constraint is specified as a percentage of total sales, i.e. $\sum_{j \in S_{smt}} q_{jmt}(p_{mt}) / \sum_{j=1}^{J_{mt}} q_{jmt}(p_{mt}) \square d_{smt}$. In both cases, the affected firms solve a constrained profit maximization problem. Solving this problem with respect to each car model j, we obtain the following set of J_{mt} first-order conditions in market m at time t; see Goldberg (1995) and Verboven (1996) for a detailed derivation:

Markets with an absolute quota:

$$\sum_{k\in F_{fmt}}\left(e_{fmt}p^{w}_{kmt}-e_{fst}\frac{\partial C_{kt}}{\partial q_{kmt}}-e_{fst}\lambda^{a}_{smt}\right)\frac{\partial q_{kmt}}{\partial p^{w}_{jmt}}+e_{fmt}q_{jmt}=0$$

Markets with a relative quota:

$$\sum_{k\in F_{fmt}}\left(e_{fmt}p^{w}_{kmt}-e_{fst}\frac{\partial C_{kt}}{\partial q_{kmt}}-e_{fst}\frac{\lambda^{r}_{smt}}{Q_{mt}}\right)\frac{\partial q_{kmt}}{\partial p^{w}_{jmt}}+e_{fst}\frac{\lambda^{r}_{smt}}{Q_{mt}}\frac{Q_{smt}}{Q_{mt}}\sum_{k=1}^{J_{mt}}\frac{\partial q_{kmt}}{\partial p^{w}_{jmt}}+e_{fmt}q_{jmt}=0$$

where $Q_{mt}=\sum_{j=1}^{J_{mt}}q_{jmt}$ and $Q_{smt}=\sum_{j\in S_{smt}}q_{jmt}$. The Lagrange multipliers λ_{smt} are expressed in the currency of the source country (as is marginal cost). In the case of an absolute quota constraint, one may interpret λ^{a}_{smt} is as the shadow marginal cost of selling an additional car from source country s into destination market m. The multipliers are identified as source- and time- specific fixed effects; the larger the multiplier, the more binding the constraint is. If the firm does not face an import constraint, its first-order condition is given by equations similar to the ones above, but with the Lagrange multipliers set to zero.

This system of J_{mt} first-order conditions can be transformed into J_{mt} pricing equations, which decompose price into a marginal cost and a markup. This transformation is useful for both econometric and interpretation purposes, as it helps us understand the sources of price discrimination. Define, for each market/year, a J_{mt} by J_{mt} matrix, $\mathbf{\Delta}_{mt}$, whose (j,k) element is $\Delta_{jk,mt}=-\partial q_{kmt}/\partial p^{w}_{jmt}$ if j and k are produced by the same firm, and $\Delta_{jk,mt}=0$ otherwise. The matrix $\mathbf{\Delta}_{mt}$ is determined by the demand side of the model; it can be estimated simultaneously with the supply side, or estimated separately in a first step (as we do in this paper). Define $\mathbf{q}_{mt}$ as a J_{mt} by 1 vector with a j-th element equal to q_{jmt}, and $\mathbf{r}_{mt}$ as a J_{mt} by 1 vector with a j-th element equal to$(Q_{smt}/Q_{mt})(1/e_{fmt})\sum_{k=1}^{J_{mt}}\partial q_{kmt}/\partial p^{w}_{jmt}$. After dividing the first-order conditions by e_{fmt}, write the system in vector notation and premultiply by $\mathbf{\Delta}^{-1}_{mt}$, the inverse of $\mathbf{\Delta}_{mt}$. This yields the following pricing equations (converted here into source currency units) for each model j in market m at time t:
Markets with an absolute quota:

$$e_{smt}p^{w}_{jmt}=\frac{\partial C_{jt}}{\partial q_{kmt}}+e_{smt}\mathbf{\Delta}^{-1}_{jmt}\mathbf{q}_{mt}+\lambda^{a}_{smt} \tag{10}$$

Markets with a relative quota:

$$e_{smt}p^{w}_{jmt}=\frac{\partial C_{jt}}{\partial q_{kmt}}+e_{smt}\mathbf{\Delta}^{-1}_{jmt}\mathbf{q}_{mt}+\frac{\lambda^{r}_{smt}}{Q_{mt}}(1+\mathbf{\Delta}^{-1}_{jmt}\mathbf{r}_{mt}). \tag{11}$$

These pricing equations, together with the expressions for the market shares and the marginal costs discussed in the previous subsections, form the empirical model we take to the data.

Before we discuss the estimation procedure, it is useful to explain how this model can be used to identify the sources of international price discrimination, and what its main limitations are. The pricing equations demonstrate that the price of each vehicle model can be decomposed in two components: its marginal cost, and a markup over marginal cost. Price differences that cannot be explained by differences in marginal costs across destinations, imply price discrimination. The markups are in turn determined by two factors: differences in the firms' perceived price elasticities of demand, and trade policies. The first factor is captured in the second term of the pricing equations by the own- and cross-price derivatives. These are in turn determined by the existence and intensity of competition. The subsection on the demand model provided more detailed intuition on the own- and cross-price derivatives. One testable hypothesis is that domestic firms face less competition than foreign firms. The second factor of price discrimination is captured by the Lagrange multipliers in the third term of the pricing equations. Binding quota in a particular country or year imply higher markups. Note that, since the quotas are based on imported units (and not on values), the markups of the inexpensive models will be relatively more affected than the markups of the luxury models if the quotas are binding. A third source of potential differences in the markups can be built into the model easily: differences in the degree of collusion. To this end, we can solve the model under alternative assumptions concerning strategic behavior, e.g. firms building different sets of coalitions in each market and engaging in joint profit maximization. Collusive behavior affects the way $\mathbf{\Delta}_{jmt}^{-1}$ is defined in (10) and (11).

The *evolution* of price discrimination can similarly be explained within the above framework. Local currency price stability can arise from a variety of factors. First, if estimation of the marginal cost function (8) indicates the presence of a significant local component in production costs, local prices will not fully respond to exchange rate shocks. Second, even without appealing to local production costs, the pricing equations (10) and (11) suggest the possibility of markup adjustment in response to cost shocks. This is indicated by several factors. First, the curvature of the perceived demand schedule is important. A firm's perceived demand elasticity for a particular model j may be increasing or decreasing in its own price and the price of its competitors. The precise shape of these elasticities is of course an empirical question; it depends on the model parameters to be estimated, the prices and the market shares. Second, the restrictiveness of import constraints may play a role. For example, quotas against Japanese firms may become less binding in a particular market as the Yen appreciates against the destination currency. Finally, to the extent that conduct affects the curvature of the perceived demand schedule, the degree of collusion may also have an impact on the degree of markup adjustment in response to cost shocks.

All these interpretations are given within the context of our static model. There is no doubt that given the durable nature of automobiles, a dynamic model of this market would be more appropriate. Unfortunately, the dimensionality of the problem makes estimation of such a model infeasible. Nevertheless, before we proceed, it is worth considering how introducing dynamics might affect our results and their interpretation.

Consider the demand side first. The effect of dynamics on car demand is two-sided. On one hand, past purchases may affect current purchases. In the aggregate data this implies that models with large market shares in the past are more likely to have large market shares in the present, once other observable factors affecting current car demand are controlled for. This effect is captured in our framework indirectly through ξ_{jmt}. The ξ_{jmt} term, that following the convention in the literature we call – perhaps misleadingly – unobserved quality, captures all unobserved factors that affect a year's demand for a model; these include past sales.[6]

The second way in which dynamics enter the picture concerns the car purchase timing, and it is substantially harder to model. Forward looking consumers base their decisions concerning car purchase, replacement and scrappage on expectations about future economic conditions, future income, the maintenance and repair cost of existing vehicles, and future car prices. At any point in time, consumers compare the cost of acquiring and operating a new vehicle, to the cost of holding an existing one. Our static framework does a poor job of explicitly capturing this behavior. At best we capture it through time dummies that proxy for macroeconomic events affecting expectations, and model dummies that proxy for different repair costs across brand names. Fortunately, this forward looking behavior is only of secondary importance for the questions we are asking in this paper; we only want to control for it in such a way, that we can obtain consistent estimates of the entities we are interested in, namely the price elasticities of demand. We believe that, while the reduced form approach we take in dealing with dynamics on the demand side often prohibits us from giving an economic explanation to some of the estimated coefficients, it still allows us to consistently estimate the price parameters.

[6]To take the "past sales" effect into account, we also experimented with a specification on the demand side that included the past year's sales of each car model in the term δ_j.This specification is rather ad-hoc and a-theoretical, given that aggregation over consumer preferences would probably give us something different; consumers buy new cars every 5-7 years, with the holding period for each vehicle depending on both vehicle and consumer characteristics. In any event, this specification did not produce results that we could interpret within our utility maximization model. One of the problems with estimating this version is that the assumption that past sales are exogenous seems implausible, yet it is hard to come up with valid instruments.

To illustrate this point consider the issue of local currency price stability.

One could plausibly argue that the reason producers keep prices in each country relatively stable is the presence of dynamics on the demand side. Producers namely anticipate that forward looking consumers would postpone a car purchase if the car price increased substantially in response to an exchange rate movement; therefore, they refrain from fully passing the exchange rate change on the price. Note, however, that this effect is accounted for in our estimation through the presence of the "outside" good. In other words, when computing the price elasticities of demand for each model, we do take into account the fact that consumers each year have the option of staying out of the new car market, by either buying a used car or postponing the purchase. In equilibrium, producers decide how much to raise the price of a car in response to an exchange rate movement taking these elasticities into account. In the absence of a dynamic model we are unable to explicitly link our estimated price elasticities to expectations about exchange rate movements. Nevertheless, we are able to obtain consistent estimates of these price elasticities.

Implicit in the above arguments is the assumption that producers maximize profits myopically, period by period. But this assumption is unlikely to hold if dynamics are important on the demand side. Dynamics on the demand side imply dynamics on the supply side. Producers are likely to maximize profits over a longer time horizon, taking into account both demand spillovers from year to year, and consumer expectations about price and exchange rate movements. It is this kind of dynamics, we believe, that would have the most severe impact on our conclusions. While modelling supply dynamics is beyond the scope of our paper, we can get a sense of how our results would be affected by utilizing the insights of some earlier work in this area. Froot and Klemperer's (1988) model is very instructive in this regard. Froot and Klemperer build a theoretical model in which demand today depends on past sales, and producers maximize the sum of current and future expected profits taking the demand side dynamics into account. It is shown that such a model generates incomplete pass-through. Moreover, the response of prices to currency fluctuations depends here on whether exchange rate movements are perceived as temporary or permanent. Temporary movements have little effect on local currency prices; producers do not want to lose current and - because of the existence of demand dynamics - also future demand, by raising prices in response to a shock that is short-lived. But a permanent exchange rate appreciation has a larger effect, and gets eventually passed through on prices. Applied to the European car market this work seems to offer an alternative explanation for the documented local currency price stability. In practice, however, we have several reasons to believe that the factors highlighted in the above model are not the primary driving force behind the observed price inertia. First, given that most of the currencies in our sample were in the EMS, it is reasonable to expect that currency realignments were viewed as permanent and not temporary by consumers. Second, previous empirical work has provided little support for Froot and Klemperer's model in which the distinction between permanent and temporary currency realignments is crucial. Finally, the persistence and magnitude of the price differentials documented in the second section makes it unlikely that these differentials were generated by the lack of adjustment to short-lived currency fluctuations. Based on these arguments, we are optimistic that our main conclusions would stand in a dynamic setting. But since there is no other definitive proof than actual estimation of a dynamic model, and since such estimation is infeasible, we interpret our estimates with caution. Our approach is simply to see how far we can go with the static model; if the static model yields counterintuitive results, or fails to explain certain features of the data, this is an indication that a more explicit consideration of dynamics is necessary.

4 Estimation

4.1 The Estimation Method

The error term ξ_j on the demand side represents unobserved product characteristics that influence the consumers' mean evaluation. The error term ω_j[7] in the marginal cost function (7) has the economic interpretation of unobserved factors that affect the producer's marginal cost for product j. Following Berry (1994) and Berry et al (1995), we proceed as follows: (1) we solve for the error terms ξ_j and ω_j as functions of the parameters and the data. This solution is derived numerically; due to the presence of the consumer specific income term y_i in the utility function, an analytical solution is not feasible. (2) we interact these error terms with a set of instruments to form a generalized method of moments estimator (GMM).

Consider the market share equations first. Let $\hat{\xi} = \delta(s, \hat{\alpha}, \hat{\rho}) - X\hat{\beta}$ be the sample analog of ξ, and Z a matrix of instruments, assumed to be orthogonal to the error term ξ. Then the GMM estimator minimizes the objective function:

$$\min_{\alpha,\rho,\beta} \hat{\xi}' Z\text{-}\ ^{-1} Z' \hat{\xi}, \tag{12}$$

where - is a weighting matrix.[8] This minimization problem involves a potentially large set of parameters. Fortunately, the computational burden can be reduced by observing that the

[7] We again supress the subscripts m and t here.

[8] In the case of homoskedastic error terms, this matrix is equal to $Z'Z$.

parameters β enter linearly. This calls for the following two-stage procedure in estimating the demand parameters. First, derive the first-order conditions with respect to β. These yield:

$$\beta = (X'Z\text{-}\ ^{-1}Z'X)^{-1}X'Z\text{-}\ ^{-1}Z'\delta(s, \alpha, \rho), \tag{13}$$

Then substitute this expression for β into the objective function (12). This 2-step procedure offers the advantage that minimization of (12) is performed only with respect to the parameters that enter non-linearly: α and ρ.

Now consider the pricing equations (10–11). Since marginal cost enters additively, one can simply substitute (7) and slightly rearrange to obtain:

$$\ln((e_s p_j^w - e_s \mathbf{\Delta}_j^{-1}\mathbf{q} - \lambda_s^*)/F_s) = z_j\gamma_s + \delta \ln(W_s/F_s) + \zeta \ln Q_j + \omega_s + \omega_f + \omega_j, \tag{14}$$

where λ_s^* refers to either the absolute or the relative quota term. Since the error term enters linearly, estimation of this equation is straightforward.

The demand and pricing equations can be estimated either jointly or separately. We chose to estimate them separately. The demand system is estimated first. The estimated demand parameters are used to construct the matrix Δ of the own- and cross-price derivatives. This matrix is then substituted into the pricing equation, which is estimated in a second step. We correct the standard errors of the parameters in the pricing equation to account for the fact that the parameters entering the matrix Δ were estimated prior to estimating the pricing equation. This two-step approach offers several advantages. First, separate estimation of demand and supply reduces the computational burden of the estimation. Second, it gives us considerable flexibility in experimenting with different supply specifications, without having to re-estimate the demand system. This is important in our case given that our data set is rich enough to allow such experimentation, and that there are several specification issues on the supply side that need to be decided empirically, e.g., pooled estimation vs. country-by-country estimation, presence of local component and hence exchange rate effects in marginal costs, Bertrand competition vs. collusion, etc... Finally, to the extent that the supply side may be misspecified, this would not affect the demand side results. The potential drawback of this procedure is of course efficiency loss; the standard errors of our parameters were, however, small enough to justify abstracting from efficiency issues.

4.2 Identification

Identification issues arise in the estimation of both the demand and supply sides of the empirical model. Consider the demand side first. The interpretation of ξ_j as unobserved product quality implies that it will be positively correlated with the product's price p_j; moreover, by virtue of the firms' first order conditions, it will also be correlated with the prices and market shares of the other products. Identification of the demand side parameters requires thus an instrumental variables matrix Z, with rank at least as large as the number of parameters to be estimated. Just as in the estimation of homogeneous product demand functions, variables that shift the producers' supply relations, but are excluded from the demand equations, make natural instruments. We therefore turn to equations (10) and (11) to identify potential instruments.

An obvious instrumental variable candidate indicated by the supply equations is the vector of exchange rates between a destination market and the source countries. As evident from the first order conditions (10) and (11), the function of exchange rates is to shift the supply relations of producers relative to each other; it is this shift that allows one to identify the demand curve suppliers face. The use of exchange rates as instruments presents two main advantages. First, they can plausibly be considered exogenous to events in the auto industry. Second, they exhibit substantial variation from year to year. On the negative side, this variation is only helpful when we estimate the demand functions facing firms located in different countries; because two firms producing in the same country face the same exchange rates vis a vis any destination market, exchange rate shocks do not affect their supply curves relative to each other. Hence, to the extent that we want to estimate the demand curves facing *individual* firms, exchange rates are helpful, but not sufficient for identification. An additional limitation stems from the fact, that when market specific time dummies are included in the demand estimation, they absorb most of the variation associated with exchange rates; we found it necessary to include such dummies in our specification to account for macroeconomic effects that may affect the purchase of the outside good (that is the option not to enter the market).

A second set of instruments can be derived by exploiting the econometric exogeneity of the products' observed characteristics (the matrix X), along the lines suggested by the recent literature. The first option that comes to mind is to include a product's own characteristics x_j in the matrix Z. Second, the supply relations (10) and (11) suggest that prices and market shares depend on the degree and closeness of competition firms face. Accordingly, we can construct functions of the exogenous characteristics of competing products, that can be thought of as proxies for the intensity of competition firms face, and use those as instruments. These functions include the number and sum over characteristics of products sharing the same cluster (that is market segment and foreign/domestic status) and/or firm ownership with product j.[9]The idea

[9]In employing these functional forms when constructing the instruments we follow a common practice in the literature that was started by Berry et al (1995). Berry et al actually prove that these instruments can be derived from a series approximation to the optimal instruments.

is that if competition in a particular cluster (as proxied by the number of products and sum of characteristics in this cluster) increases, the demand curve associated with a product in this cluster will become flatter; similarly, the greater the number of products a particular firm sells, the higher the price it will charge on each product.

In summary, the full set of instruments included in matrix Z contains: (1) A product's observed physical attributes x_{jmt}; (2) Exchange rates between source and destination countries, or market specific time dummies; (3) the sums of characteristics of other products belonging to the same cluster, and the number of products belonging to the same cluster; and (4) the sums of characteristics, and number of other products made by the same firm and belonging to the same cluster.

On the supply side, both the quantity Q_j and the quota terms (that are functions of the total car sales in each market) are likely to be correlated with the error term ω_j. In particular, a relatively low unobserved marginal cost ω_j would lead the firm to produce a higher quantity Q_j, implying that the quantity coefficient would be biased downward; similar arguments apply to the quota terms. Of course, to the extent that Q_j refers to product j's *total* sales in Europe, and not the individual sales in each market, the simultaneity bias is accentuated. Nevertheless, given that for some products, sales in individual countries of our sample constitute a large proportion of total sales,[10], it is important to instrument for the quantity and the quota terms. The quota terms were instrumented using interactions of Japanese product and destination market dummies. Finding valid instruments for quantity is, however, less straightforward.

In the search for appropriate supply side instruments we are looking for variables that shift the demand function of each model along the supply curve. But while it is relatively easy to think of variables that shift the aggregate demand for cars, it is more difficult to come up with variables that are both destination and model specific demand shifters. In instrumenting for the quantity Q_j, we experimented with the following approaches: First we interacted aggregate demand growth in each destination market with lagged market share for each model and used the resulting variable, in addition to all predetermined variables in (14), to instrument for quantity. This approach gives us destination and car make specific instruments, but may be problematic if one suspects serial correlation in the data. Hence, we experimented with a second approach: subtract from the total European sales of each product the sales to each individual market in our sample, and use the difference (again in addition to all predetermined variables in the model) as an instrument for quantity. This identification assumption is however

[10] This is especially the case with products produced by firms registered in one of the sample countries; for example, a large fraction of products produced by Italian firms is sold in Italy.

problematic if the unobserved cost of selling to one market is correlated with the unobserved cost of selling to other markets; this would for example be the case if all cars of a particular make are manufactured in the same plant.

Since neither of the instruments discussed above is perfect (they are both valid only under specific assumptions), we also considered three alternative ways of dealing with the quantity coefficient. The first one is to follow a substantial part of the empirical literature in assuming constant returns to scale, that is constant marginal cost. The second approach is to impose increasing returns to scale, that is impose a negative quantity coefficient, and estimate the remaining coefficients of the model. Finally, we also experimented with imposing decreasing returns to scale (a positive quantity coefficient).

Each of the above approaches had different implications for the presence of increasing returns to scale. Hence, as we discuss in the results section, we refrain from making strong statements about the existence or non-existence of returns to scale in car production. The other coefficients of the cost function, however, were robust to the estimation alternatives described above. As we discuss in section 5, the focus of this study is on the country fixed effects and the coefficient on the local cost component. Since these coefficients were robust to the various ways we estimated the supply side, we are fairly confident about the conclusions based on these results.

4.3 Accounting for Heteroskedasticity and Autocorrelation

Given that in our data each model appears in several years and multiple markets, it is important to account for potential heteroskedasticity and autocorrelation of the error terms in the estimation procedure. To this end, we considered two alternative procedures. The first one is to average the sample moment conditions across time, as suggested by Berry et al (1995), to obtain the moment conditions for each model in each market. This method allows for arbitrary heteroskedasticity and autocorrelation patterns in the error terms, but is not necessarily efficient. The second approach is to impose a particular covariance structure, and use that in computing the weighting matrix - and the standard errors. We considered the following structure that we believe is plausible within the context of our application:

The unobserved quality ξ_{jm} for product j in market m is assumed to follow a first order autoregressive process:

$$\xi_{jm,t} = r_m \, \xi_{jm,t-1} + u_{jmt}$$

with $E(u_{jmt}) = 0$, $E(u_{jmt}^2) = \sigma_{jm}^2$, $E(u_{jmt}, u_{jms}) = 0$ for $t \neq s$, and $\|r_m\| < 1$, while

$$E(\xi_{jmt}, \xi_{j'mt}) = 0 \; for \; j \neq j'$$

This specification, though restrictive, has the advantage of parameterizing the covariance structure in a way that allows for heteroskedasticity and serial correlation that declines geometrically over time, while limiting the number of additional parameters that need to be estimated.

A-priori, it is not clear which of the two methods is more efficient. We experimented with both, and did not find a significant difference in the results (in both cases all parameters were very precisely estimated). The results we report below are derived using the second method.

5. Results

5.1 Summary of Estimation Results

Demand: The main results from the estimation of the demand system are summarized in Table 4. The demand system was estimated separately for each country. Most of the coefficients were significantly different across countries, thus supporting the separate estimation against pooling of the data. However, inspection of the coefficients relating to price and physical car characteristics indicated the existence of collinearity. To mitigate the impact of collinearity on the results, we averaged these coefficients across countries, using the inverse of the variance of each coefficient as a weight.[11] Accordingly, Table 4 reports only one price coefficient (α) for all countries.

In general, all coefficients (including the ones not reported in the table, such as the coefficients on physical car characteristics and time dummies) had the expected signs and were very precisely estimated. Tests of overidentifying restrictions failed to reject the model (the test statistics were between 30 and 36 for the various countries, while the critical value at the 5% significance level and with 23 degrees of freedom is 35.17). The time dummies for each country are plotted in Figure 5; these capture preference for autos relative to the outside good, and can thus be thought of as proxies for macroeconomic effects that affect the likelihood of purchasing a new car. It is interesting to note that the time pattern of the dummies traces the business cycle in each country; the low time dummy values in the early 1980's, for example, coincide with the recession in European countries, the positive trend between 1985 and 1990 coincides with the rising income in these countries, and the decline of the Italian time dummies after 1990 coincides with the recession in Italy. These results are reassuring as their plausibility provides

[11] Alternatively, we could have estimated the demand system by pooling the data for all countries, imposing the same price and car attribute coefficients for all countries, and letting the ρ's and the origin dummies being country specific. This specification would have been appealing, as it implies that consumer preferences for various car attributes do not vary across markets, but there may be unobserved factors (proxied by the country dummies and captured by cross-country differences in the ρ's) that account for differences in the demand patterns. We did not pursue this alternative, because pooling of the data exceedingly complicates the covariance structure. Our procedure of averaging (using appropriate weights) the coefficients potentially affected by collinearity is, however, similar in spirit to the alternative just described.

some indirect support for the estimated specification.

Of particular interest are the country dummies reported in Table 4, that capture preferences for cars originating from a particular country. In all four countries in which there is domestic production, consumers have a distinct preference for domestic cars, as indicated by the higher coefficient for the domestic brands in each country.[12]

The other noteworthy coefficients are the ρ's, that capture the correlation pattern of the error terms within each cluster. For each country there are two ρ coefficients, one corresponding to the distinction between foreign and domestic (ρ_f), and one corresponding to the distinction between market segments (ρ_c).[13] Only in Belgium there is a single ρ coefficient as there are no Belgian domestic brands, so the model reduces to a one-level nested logit, with the nests representing different market segments.[14] Note that the ρ parameters generally satisfy the restrictions imposed by utility maximization; they lie in the unit interval, and ρ_c is greater than or equal to ρ_f. The only exception is Germany, where $\rho_c < \rho_f$; note, however, that the Wald test does not reject the hypothesis that the two coefficients are equal in this country. The other interesting case is the one of Italy; the coefficient ρ_c is equal to 1 there, implying that the model reduces to a one-level nested logit, in which each nest includes cars of the same market segment and the same domestic/foreign status. In short, we can summarize the pattern that emerges from Table 4 with regard to the demand structure as follows: In Belgium the demand model reduces to a one-level nested logit, as there are no domestic brands. In France and the U.K. the demand structure is described by a two-level nested logit, with market segments and foreign/domestic representing the two levels of nesting. In Italy and Germany the model again reduces to a one-level nested logit, but the nests look different in the two cases: in Italy products compete closely only with products of the same market segment and origin; in Germany, products compete closely with products of the same market segment irrespective of whether these products are foreign or domestic. As we will see in the next section, these differences in the nesting structures suggested by our estimation results have implications for the substitution patterns, prices and market power in each country.

Costs: The results from estimating the pricing equation (14) are summarized in Table 5. In

[12] For example, in France the largest country dummy is the one for French cars; in Italy for Italian cars, etc.. The country coefficients capture preferences relative to the U.K.; accordingly, the negative coefficients in the U.K. column indicate that cars from countries other than the U.K. are valued less than British cars.

[13] We also experimented with specifications in which the ρ coefficients were allowed to be market segment specific; this makes the model more flexible. In all five countries, however, we failed to reject the hypothesis that the ρ's were equal across different market segments, so we adopted the reported specification with only two ρ coefficients.

[14] In all other countries, between 20 and 30% of the models in each market segment are domestic.

the reported specifications we pooled the data for all countries, imposing the same coefficients on physical characteristics; we also estimated separate pricing equations by producing country, but the coefficient estimates were not statistically different across countries, while the point estimates were generally similar to the ones in Table 5. In columns 1 and 2 we report the results that were obtained by assuming Bertrand competition in all five countries. All parameters have the expected signs and are precisely estimated. The positive coefficient of the foreign firm dummy indicates that foreign firms face a cost disadvantage. The size and significance of the source country log wage coefficient indicate that marginal cost is sensitive to labor costs, i.e., there are few substitution possibilities towards other production factors.

Columns 1 and 2 differ in the choice of instruments. In column 1, we instrument for the log of the world production quantity (LQU) using the difference between total sales and sales in each market as an instrument. As noted above, this instrument is problematic if a common unobserved shock affects marginal costs in all countries (this would be for example the case if all cars of a particular make are produced in a common facility). Column 2 reports the results from an alternative instrumentation, in which the product of the aggregated demand growth in a destination market times the lagged sales for a particular product, is used as an instrument for quantity. This instrument is problematic in the presence of serial correlation. Note that in both cases we estimate increasing returns to scale (a negative coefficient on LQU), but using lagged sales interacted with demand growth (column 2) produces a larger estimate. A similar pattern emerges with respect to the wage, and car characteristics coefficients. All these coefficients have the expected signs; the signs are robust to the specification, but the size of the estimates varies depending on the choice of the instruments. The remaining coefficients (that is the foreign firm disadvantage, the quota constraint, and the local wage coefficients and the country fixed effects) exhibit even stronger robustness to the alternative ways of instrumenting; not only the signs, but also the size of the coefficients remain unchanged from specification to specification. These parameter estimates are discussed in more detail below. At this point we only want to note that since these are the parameters relevant for identifying the determinants of price dispersion, we are quite confident that our conclusions do not depend on particular instrument choices.

Note first the year quota variables for Japanese cars. As Table 5 indicates, these are jointly significant at any reasonable significance level. Figure 6 plots the estimated quota effects for each country; the estimates have the interpretation of Lagrange multipliers, thus capturing the restrictiveness of quantity constraints in each market. The plots are intuitive and consistent with our expectations. In almost every year the quota effects are highest for Italy, the country with the most stringent import constraints; Italy is followed by France, and then the U.K. and Germany, where the estimates were not statistically significant. A comparison of Figure 6 with the data on Japanese quotas reported in Table 1 (line 6) suggests that this order corresponds exactly to the shares allocated to the Japanese imports in each country. Another interesting pattern that emerges from Figure 6 is that in each country the estimated Lagrange multipliers are highest in years in which the Yen depreciates against the local currency. This is intuitive as a weaker Yen strengthens the position of Japanese producers, hence making the import constraint more binding.

Next note the coefficient on the log of local wages, LWAGELOC. Under the hypothesis that marginal costs in each country do not contain a local component, this coefficient should be zero. A positive coefficient indicates that some fraction of the marginal costs occur in local currency. This result obviously has implications for the degree of local currency price stability; when the exporting country's currency appreciates against the importing country's currency, costs denominated in exporter currency go up.[15] We explore these implications in detail in the next section. Since the results of the pricing-to-market regressions in section 2.3 indicated a higher degree of local price stability in the U.K., we also interact the variable LWAGELOC with a U.K. dummy (column 2). The estimated coefficient is statistically significant, but too small in magnitude to support the hypothesis that marginal costs in the U.K. are affected by exchange rate changes by more than in other countries.

Note finally the estimated fixed effects for each country. Their relative magnitudes are plausible as they imply that unobserved costs are highest in the U.K. and Germany, and lowest in Belgium. The obligatory use of catalytic converters in Germany is consistent with these estimates. Regarding the U.K. fixed effect, it is common wisdom that cars in the U.K. are better equipped than in the rest of Europe due to the existence of a well developed market for company cars.[16] In addition, dealer discounts are reportedly larger in the U.K. than in the other countries.[17] Since the reported specifications employ list prices, differences in discount practices across countries manifest themselves as differences in the country fixed effects in the

[15]Recall that the local wages are denominated in the producing country's currency. Hence they are defined as the product of the local wages denominated in the destination market's currency and the exchange rate between producing and destination currency. See also equation (8).

[16]The U.K. National Consumer Council (1990) reports that between 55 percent and 60 percent of all new cars sold in the U.K. are bought by a company or as a business expense. In Germany and France, company cars account for around 15 percent of sales.

[17]B.E.U.C. and the U.K. Monopolies and Mergers Commission report discounts of up to 15% in the U.K. compared to a maximum of 10% in other countries. For company cars discounts are less well documented, but the U.K. Monopolies and Mergers Commission reports that discounts to fleet customers probably vary between 16 and 22% percent.

cost equations. The right-hand drive regulation may also contribute to higher costs. Another line of interpretation for the high fixed effect in the U.K. could focus on differences in firm behavior, e.g., collusion in the U.K. Apart from the magnitude of the cost fixed effect, the U.K. stands out as the only European country in which the Monopolies and Merger Commission initiated a legal investigation for collusive behavior. This prompted us to explore the possibility of collusion in the estimation of the supply side of the model.

To this end, we reestimated the pricing equation (14) assuming collusion in the U.K., and Bertrand competition in the other four countries. As noted before, the assumption of collusion changes the matrix Δ of own and cross-price derivatives in equation (14). The results from this specification are reported in column 3 of Table 5. Ideally, we would like to employ a goodness-of-fit statistic to distinguish between the two alterative models, Bertrand vs. collusion. Both models, however, explain the data quite well and seem indistinguishable in terms of their fit. Therefore we escorted to an alternative way of assessing their relative merit: examine the estimated coefficients under the two specifications and assess which ones appear more plausible. Note that the coefficients in columns 2 and 3 of Table 5 are almost identical, except from one: the fixed effect for the U.K., which drops from 0.15 in the Bertrand specification to 0.02 in the collusive one. Intuitively, all the assumption of collusion does, is blow up the estimated markups for the U.K., hence suppressing the magnitude of marginal costs. Nevertheless, while the U.K. fixed effect estimate appeared too high under Bertrand, it seems too low under collusion, given the anecdotal evidence pertaining to better equipped cars and higher discounts in the U.K.. Given these results, it is impossible to make a definitive statement as to which behavioral assumption is more appropriate for the U.K., even though it seems that some type of behavior that is more collusive than Bertrand, but less collusive than joint profit maximization might produce the most plausible fixed effects estimates.

As noted above, the results relevant for price dispersion are robust to alternative sets of instruments used in estimating the pricing equation. Yet, none of these instruments is entirely satisfactory. Therefore, we also experimented with alternative specifications in which we imposed a particular coefficient on the quantity coefficient, and estimated the rest of the parameters. The second page of Table 5 reports results from three such specifications. In the first one, we imposed a quantity coefficient equal to zero, thus assuming constant returns to scale or constant marginal cost. This assumption has been frequently used in the literature (e.g., Bresnahan (1981), Goldberg (1995))); nevertheless, it is at odds with the estimates in the first page of Table 5 that indicated the presence of increasing returns. The second specification we considered is one in which we *impose* increasing returns to scale, that is a particular negative value for the coefficient on log quantity. Recent empirical studies of the auto market have always found increasing returns to scale (see for example Berry et al (1995) and Verboven (1996)), but there is some controversy surrounding the size of estimated coefficient (the estimated returns to scale are generally viewed as too large to be plausible). Finally, in the third specification we *imposed* decreasing returns to scale, that is a particular positive value on the log-quantity coefficient. This assumption is less plausible given the results of previous empirical studies and our own results based on instrumental variable estimation; nevertheless, we wanted to examine whether our results would remain robust to radically different assumptions concerning returns to scale. The last three columns of Table 5 offer a clear answer to this question: no matter what one believes the true returns to scale in car production are, the estimates of the country fixed effects and the coefficient on the local wages are quite robust.[18]

5.2 The Sources of Price Differentials

Having estimated the demand and supply parameters of the model, we now return to the question we posed at the beginning: Why are prices so different across European countries? Our estimation results suggest three main factors contributing to systematic price differences.

The first factor is differences in the demand patterns, and in particular, the existence of a strong bias for domestic brands. This bias manifests itself in two forms in our estimation results. First, the country dummies in the demand estimation (see Table 4 and the discussion in the previous section) unambiguously demonstrate the existence of a strong preference for domestic cars in each country. Second, the differences in the nesting structures, indicated by the differences in the relative magnitudes of the ρ parameters, also have implications about the intensity of competition between domestic and foreign brands.

Consider France and the U.K. first, the two countries in which the bi-nested logit structure is supported by the ρ estimates. The estimation results imply that consumers view products of the same domestic/foreign status as closer substitutes than products of different origin. Hence, when the price of a domestic (foreign) product goes up, they are more likely to substitute towards another domestic (foreign) product in the same class, than switching to a car of a different nationality. This pattern is to be contrasted with Germany, where domestic and foreign products are viewed as equally likely substitutes as long as they belong to the same market segment. In Italy, on the other hand, the domestic/foreign distinction is even more

[18] Columns 5 (IRS) and 6 (DRS) of Table 5 report results from specifications in which the quantity coefficients were set at -0.1 and 0.1 respectively. The results from imposing other negative or positive values were very similar. We chose to report the results corresponding to unrealistically high values for increasing (or decreasing) returns to cover the whole interval of plausible values and demonstrate the robustness of the coefficients of interest.

relevant for consumer behavior than in France and the U.K.; cars here compete closely only with products that share the same market segment and the same domestic/foreign status. From the consumer's point of view, a German subcompact lies as far from an Italian subcompact, as a German luxury car.

The implications of these estimates for the substitution patterns are summarized in Table 6, that lists the average own and cross-price elasticities of demand in each market. Note that with the exception of Germany, the average price elasticity of demand for foreign cars is always – consistent with the existence of a home bias – higher than the corresponding elasticity for domestic cars. This pattern is most pronounced in Italy which also has the lowest average price elasticity. The importance of the domestic/foreign distinction for consumer behavior in each market is also reflected in the pattern of the cross-price elasticities; in all markets but Germany the average cross-price elasticities between any two domestic (or foreign) cars are higher than the cross-price elasticities between domestic and foreign. This pattern is again much more pronounced in Italy than in the other countries. Table 7 confirms these substitution patterns for some selected models in Belgium and in Italy. Similar patterns arise for the other models.

What are the implications of these differences in demand patterns for price dispersion? Recall that the price of an automobile in each country can be decomposed into a marginal cost and a markup component. The differences in elasticities translate into differences in the cross-country markups. The average relative markups (Lerner indices) are listed in the bottom part of Table 6; again, they are broken down by foreign/domestic.

The one thing that strikes one immediately when comparing average markups across countries, is that markups in Italy are substantially higher than in the other countries. In comparison, markup differences across the other countries are relatively small; in France, for example, markups are on average only by 2% higher than in Belgium, while in the U.K. the difference goes the other way, and markups are by 2% lower. Overall, however, these differences are too small to generate economically significant price differences. In contrast, markups in Italy generate a 14% difference in prices relative to Belgium. As evident from the last two rows of Table 6, the high markups in Italy are primarily driven by the market power of domestic producers; foreign car markups are approximately of the same magnitude as in the other countries.

Why does preference for domestic cars generate such a high degree of market power in Italy? The answer seems to lie in the fact that Fiat occupies an almost monopoly position in Italy. The substitution patterns implied by our estimates suggest that when the price of a domestic car goes up, consumers are more likely to switch to another domestic car. But while in other countries the substitute car may belong to a different firm, in Italy it is very likely produced by the same firm, namely Fiat. It is this combination of home bias and a near monopoly position of the domestic firm that seems to generate market power in Italy. In short, it appears that differences in the price elasticities of demand, that translate into differences in markups, can explain a significant fraction of the price differentials in Italy relative to other countries. This explanation however does not seem to fit the other expensive country in our sample, the U.K., where the price elasticities of demand seem to be going the wrong way: they imply lower, and not higher prices, even though the difference is not statistically significant.

The second factor that generates price dispersion in the five European countries in our sample is the differential impact of the quota constraints. These appear much more restrictive in Italy and France than in the other three countries. The higher prices of Japanese cars in Italy and France could hence be attributed to the quota constraints. To assess to what extent these constraints also affect the prices of the other cars in each market, we reestimated the pricing equation (14) for each country, omitting the Japanese year/quota variables. The resulting increase in the fixed effect estimates for each country can be interpreted as the overall price (or shadow marginal cost) increase generated by the existence of import constraints. According to our results, this increase is approximately 2% for France, 1% for the U.K., and 0.5% for Italy.[19] In Germany, quota constraints do not affect costs or prices; this is of course to be expected given that the Lagrange multipliers estimated for Germany were insignificant in all years. The small impact of the quota constraints on overall costs and prices in Italy may seem surprising at first, given that the Lagrange multipliers for Italy were significantly higher than for any other country in our sample. Note, however, that due to the extremely restrictive quota allocated to Japanese imports in Italy (1%), the fraction of cars that are affected by the quota is very small (namely ca. 1% of the market). Hence, while the prices of Japanese cars are significantly higher in Italy than in the rest of Europe, their share in total sales is too small to drive the overall higher prices documented in Italy.

The third factor generating price dispersion in Europe is unobserved costs, as proxied by the country fixed effects in the pricing equation. These were discussed in detail in the previous section, and it was shown that their magnitudes were generally plausible and consistent with industry wisdom. Of particular interest is the U.K. fixed effect. Recall that our descriptive analysis at the beginning of the paper demonstrated the two countries with the systematically higher prices are Italy and the U.K.. While the demand side of the model can explain the higher prices in Italy, it fails to account for the high prices in the U.K.. The latter are explained entirely by the supply side. As discussed earlier, it is impossible to determine on the basis of our results

[19] All these figures are very precisely estimated, with t-statistics ranging from 6 to 20.

whether the high fixed effect estimate for the U.K. is due to better equipped cars and higher dealer discounts, or collusive behavior.

In summary, our estimation results suggest that approximately 20% of the cross-country price differentials can systematically be explained by structural demand side (in Italy) and supply side (in the U.K.) factors, that are unrelated to exchange rate variation.[20] Nevertheless, cross-country price differentials in each year tend to be significantly higher, approximately 30% of the average car price as shown in the introduction. Hence there remains a 10% difference to be explained. Our descriptive results in section 2.3. suggested that exchange rate volatility plays an important role in year-to-year changes in price differentials. Therefore we now turn to a more systematic investigation of the implications of our estimates for local currency price stability.

5.3 Local Currency Price Stability

To understand the implications of our results for local currency price stability, it is instructive to compare the pricing equations (10) and (11) to the pricing-to-market equation (2) we estimated at the beginning of the paper. In the pricing-to-market literature, only prices and exchange rates are observed; accordingly, the challenge for the researcher is to find a way to decompose the price variation into a component that reflects movements in the marginal cost, and a component that reflects movements in the markup, without actually observing either marginal cost or markup. As discussed in section 2.3, this is achieved by testing for correlation of export prices with exchange rates, while controlling for common across markets changes in marginal costs through time effects, and product quality differences and/or constant markups through country fixed effects.

Compared to this approach, the advantage of estimating a fully specified model of supply and demand is that it allows us to estimate the exact markups and marginal costs in each market and each year. This is evident from equations (10) and (11); prices denominated in exporter currency can now explicitly be decomposed into a marginal cost component ($\frac{\partial C_{jt}}{\partial q_{kmt}}$), the markup ($e_{smt}\mathbf{\Delta}_{jmt}^{-1}\mathbf{q}_{mt}$), and a third component that is related to the existence of import constraints. The pricing-to-market regressions established that local currency prices remain relatively stable when exchange rates fluctuate, so that a 1% change in the exchange rate between an exporting and an importing country leads to a 0.46% change in the exporter currency price. In this section

[20] In particular, the combined results of Tables 5 and 6 suggest that in the U.K., 15% of the price differential relative to Belgium can be attributed to unobserved fixed costs, while 1% is due to quota constraints; in Italy, 14% is attributed to differences in the markups, 5% to unobserved costs and 0.5% to the quotas; in Germany, 17% is attributed to cost differences, and in France 2% is attributed to markup differences, 4% to cost differences, and 2% to the impact of the quota constraints.

we investigate the determinants of this local currency price stability; is it driven by changes in the marginal cost, or does it reflect markup adjustment?

The results of the estimation of the pricing equation reported in Table 5 speak directly to this issue. The estimated coefficient on the log of the local wages in the destination country (LWAGELOC) is positive and highly significant (t-statistic: 29). This suggests the presence of a local component in marginal costs. Moreover, this component appears to be economically significant. Note that the local wages in the pricing equation estimated in Table 5 are expressed in the currency of the producer. Hence, the log of the local wages represents the sum of two variables: the log of the local wages expressed in destination currency, and the log of the exchange rate between the producing and destination country $(log(W_{mt}) + log(e_{smt}))$.[21] The coefficient on LWAGELOC can accordingly be also thought of as the "exchange rate coefficient"; it gives us the effect of an exchange rate change on the marginal cost denominated in the producing firm's currency. The point estimate of this coefficient is, depending on the specification, between 0.37 and 0.39, implying that a 1% appreciation of the exchange rate of the exporting country against the currency of the destination market changes costs denominated in exporter currency by 0.37-0.39%. This is a large effect, and naturally the question arises whether it is plausible.

To get a rough idea about the plausibility of this estimate, we went back to industry sources and reports of the European Commission to collect some information about the importance of local costs. Industry wisdom is that local costs are up to 35% of the value of a car.[22] Even though this number is vague and hard to confirm, it gives us some idea about the order of magnitude of exchange rate effects on costs. In particular, suppose that markups remained constant during periods of exchange rate volatility; then it is easy to show, using the firms' first order conditions, that the expected effect of a modest exchange rate change on costs would be around 35%.[23] Even though, in the absence of specific cost information, we are not able to make more precise statements regarding the magnitude of the local component in marginal cost, it seems safe to say that the 37-39% coefficient the pricing equation estimation yields, appears plausible.

[21] Note that the specification in (8) imposes the same coefficient on these two components, namely θ.

[22] See, for example, De Financieel Economische Tijd, January 15, 1998, for a reference to this number. This relatively high number can be partly explained by the presence of many small dealers.

[23] We put emphasis here on the word around. Note that the statement that local costs are about 35% of the value of the car, is itself inconsistent during periods of exchange rate volatility, as exchange rate changes affect the proportion of local costs in total costs. One way to interpret this statement is that exchange rate movements are not too large, so that the proportion of local costs, even though affected by the exchange rate change, remains roughly of the same order of magnitude as before.

The implications for local currency price stability are then straightforward. Out of the 46% adjustment of the export prices that the pricing-to-market regressions estimated, 37-39% can be attributed to a change in marginal costs caused by the exchange rate movement. The remaining 7-9% are due to markup adjustment. This number is significantly lower than the original 46% suggested by the pricing-to-market results. Hence it appears that local costs play an important role in generating the documented inertia of local currency prices.

6. Conclusion

In this paper we set out to investigate the sources of price dispersion in European countries paying particular attention to the role of exchange rates. Our results suggest that across the 14 years of our sample, there is an average price difference of 20% between the U.K. and Italy, and the cheapest country in our sample, Belgium. This price difference is attributed primarily to cost and discount differences in the U.K., and to price discrimination related to the existence of domestic brand bias in Italy. Around this 20% mean, there are substantial fluctuations from year to year, with cross-country price differentials becoming as large as 35-40% in individual years. These fluctuations can be attributed to the incomplete response of local currency prices to exchange rate changes. The local currency price stability reflects both the existence of a local component in marginal costs, and price discrimination (markup adjustment) that is correlated with exchange rate volatility.

If one were to judge the effects of the EMU on price dispersion based on these results, one would conclude that the EMU will very likely eliminate the year-to-year volatility observed in our data, but cross-country price differences will not completely disappear without further measures to harmonize requirements and promote integration. Of course, given that local currency prices tend to exhibit stability over time, the particular levels at which exchange rates will be fixed at the outset of the EMU, will be important. In 1993, for example, the exchange rate parities were such, that price differentials were almost eliminated across countries. By 1998, the U.K. had again become substantially more expensive than the rest of Europe. If exchange rates remained fixed at the current level in the future, the U.K. would likely remain more expensive.[24]

With the EMU in place, these predictions will of course be soon put to a test. The EMU will provide an unprecedented experiment for the purpose of assessing the role of exchange rate in generating cross-country price dispersion. But the actual data required to evaluate the effects of this experiment are at least a few years away. We hope that this exercise will provide some useful insights in the mean time.

[24]On the other hand, in 1998 the U.K. was still not part of the EMU.

REFERENCES

Atkinson, A.B., "Measurement of Trends in Poverty and the Income Distribution," University of Cambridge, *DAE Working Papers*, no. 9701, 1997.

Berry, S.T., "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal of Economics*, Vol. 25, 1994, pp. 242–262.

Berry, S.T., Levinsohn, J. and Pakes, A., "Automobile Prices in Market Equilibrium," *Econometrica*, Vol. 63, July 1995, pp. 841–890.

BEUC, "Report on Car Prices in the EEC Countries," BEUC/71/81, 1981, Brussels.

BEUC, "Car Prices in the EEC Countries," BEUC/121/86 , 1986, Brussels.

BEUC, "Car Price and Progress towards 1992," BEUC/10/89, 1989, Brussels.

BEUC, "Parallel Imports for Cars in the EC," BEUC/222/92 , 1992, Brussels.

Bresnahan, T., "Departures from Marginal-Cost Pricing in the American Automobile Industry," *Journal of Econometrics*, 17, 1981, pp. 201–221.

Bresnahan, T., "Empirical Methods for Industries with Market Power," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization*, Vol. 2, Amsterdam: North-Holland, 1989.

Bresnahan, T., Stern, S. and Trajtenberg M., "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980's," *Rand Journal of Economics*, Vol. 28, 1997, pp. 17–44.

Commission of the European Communities, "Intra-EC Car Price Differential Report," IV/4043/92, 1992, Brussels.

Davidson, R., Dewatripont, M., Ginsburgh, V. and Labbe, M., "On the Welfare

Effects of Anti-discrimination Regulations in the EC Car Market," *International Journal of Industrial Organization*, Vol. 7, 1989, pp. 205–230.

de Melo, J. and Messerlin, P.A., "Price, Quality and Welfare Effects of European VER's on Japanese Autos," *European Economic Review*, Vol. 32, 1988, pp. 1527–46.

Feenstra, R.C. and Levinsohn, J., "Estimating Markups and Market Conduct with Multidimensional Product Attributes," *Review of Economic Studies*, Vol. 62, 1995, pp. 19–52.

Flam, H. and Nordstrom, H., "The Single Market for Cars in Europe," mimeo, Institute for International Economic Studies, Stockholm University, 1994.

Froot, K. and Klemperer, P., "Exchange Rate Pass-through when Market Share Matters," *American Economic Review*, Vol. 79, 1988, pp. 637–654.

Gagnon, J.E. and Knetter, M.M., "Markup Adjustment and Exchange Rate Fluctuations: Evidence from Panel Data on Automobile Exports," *Journal of International Money and Finance*, Vol. 14/2, April 1995, pp. 289–310.

Ginsburgh, V. and Vanhamme, G., "Price Differences in the EC Car Market," *Annales d'Economie et de Statisique*, Vol. 15/16, 1989, pp. 137–149.

Goldberg, P.K., "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, Vol. 63, July 1995, pp. 891–951.

Goldberg, P.K. and Knetter, M.M., "Measuring the Intensity of Competition in Export Markets," *Journal of International Economics*, forthcoming, August 1998.

Goldberg, P.K. and Knetter, M.M., "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, September 1997, pp. 1243–1272.

Gual, J., "An Econometric Analysis of Price Differentials in the EEC Automobile Market," *Applied Economics*, Vol. 25, 1993, pp. 599–607.

Hens, T., E. Jager, A. Kirman and L. Philips., "Exchange Rates and Oligopoly," *European Economic Review*, Vol. 43(3), 1999, pp. 621–48.

Kirman, A.P. and Schueller, N., "Price Leadership and Discrimination in the European Car Market," *Journal of Industrial Economics*, Vol. 34, 1990, pp. 69–91.

Knetter, M.M., "Price Discrimination by U.S. and German Exporters," *American Economic Review*, Vol. 79, March 1989, pp. 198–210.

Knetter, M.M., "International Comparisons of Price-to-Market Behavior" *American Economic Review*, Vol. 83/3, June 1993, pp. 473–86.

Krishna, K., "Trade Restrictions as Facilitating Practices," *Journal of International Economics*, Vol. 26, 1989, pp. 251–70.

Krugman, P., "Pricing to Market When the Exchange Rate Changes," S.W. Arndt and J.D. Richardson, eds., *Real-Financial Linkages Among Open Economies*, MIT Press, 1987.

Malueg, D.A. and Schwartz, M., "Parallel Imports, Demand Dispersion, and International Price Discrimination," *Journal of International Economics*, Vol. 37, 1994, pp. 167–195.

Marston, R., "Pricing to Market in Japanese Manufacturing," *Journal of International Economics*, Vol. 29, December 1990, pp. 217–36.

McFadden, T., "Modelling the Choice of Residential Location," in A. Karlquist et al., eds., *Spatial Interaction Theory and Residential Location*, Amsterdam; New York: North-Holland, 1978.

Mertens, Y., "Modeling Price Behavior in the European Car Market," mimeo, Department of Economics, London School of Economics, 1990.

Mertens, Y. and Ginsburgh, V., "Product Differentiation and Price Discrimination in the European Community: The Case of Automobiles," *Journal of Industrial Economics*, Vol.

34, 1985, pp. 151–165.

Monopolies and Mergers Commission, "New Motor Cars – a Report on the Supply of New Motor Cars witin the United Kingdom," London: HMSO, 1992.

National Consumer Council, "Cars – the Cost of Trade Restrictions to Consumers," *International Trade and the Consumer*, Working paper 4, 1990.

Nevo, A., "Measuring Market Power in the Ready-to-Eat Cereal Industry," mimeo, 1996.

The New York Times "A Car Buyer's Heaven In a Borderless Europe,"Wed., Jan. 17, 1996, Section D.

Petrin, A., "Quantifying the Benefits of New Products: The Case of the Minivan," mimeo, 1999.

Smith, A. and A. Venables, "Completing the Internal Market In the European Community. Some Industry Simulations," *European Economic Review*, 1988, vol. 32, pp. 1501–25.

Smith, A. and A. Venables, "Economic Integration and Market Access," *European Economic Review*, 1991, vol. 35, pp. 388–95.

Smith, A. and A. Venables, "Counting the Cost of Voluntary Export Restraints in the European Car Market," in E. Helpman and A. Razin, eds., *International Trade and Trade Policy*, 1991, pp. 187–222.

Verboven, F., "International Price Discrimination in the European Car Market," *Rand Journal of Economics*, Summer 1996, vol. 27/2, pp. 240–68.

Appendix

In this Appendix we offer a more extensive discussion of the functional forms adopted in the demand estimation and their implications for the substitution patterns. For notational convenience, let us drop the subscripts m and t. Consumer i is assumed to maximize the indirect utility function:

$$U_{ij} = \delta_j + \alpha ln(y_i - p_j) + \epsilon_{ij}$$

where the notation is as explained in the main text. Let $V_{ij} = \delta_j + \alpha\, ln(y_i - p_j)$. The expression V_{ij} represents the deterministic part of the utility function. The latter can now be written as $U_{ij} = V_{ij} + \epsilon_{ij}$. In specifying the distribution of the error term ϵ, we make use of McFadden (1978), who in Theorem 1 states the conditions under which the specified distribution is consistent with random utility maximization. In particular, let $F(\epsilon_{i0}, ..., \epsilon_{iJ})$ denote the $J+1$ dimensional CDF of ϵ. According to Theorem 1:

If $G(y_0, ..., y_J)$ is a nonnegative, homogeneous-of-degree-one function satisfying certain restrictions, then

$$F(\epsilon_{i0}, ..., \epsilon_{iJ}) \;=\; exp(-G(e^{-\epsilon_{i0}}, ..., e^{-\epsilon_{iJ}}))$$

is the multivariate extreme value distribution, and

$$P_{ij} \;=\; \frac{e^{V_{ij}} G_j(e^{V_{i0}}, ..., e^{V_{iJ}})}{G(e^{V_{i0}}, ..., e^{V_{iJ}})}$$

defines the probability of choosing alternative j, where G_j is the partial derivative of G with respect to $e^{V_{ij}}$.

This model is known as the generalized extreme value (GEV) model. We considered two specifications of the G-function: the bi-level nested logit, which is given in Section 3.1, equation (4), and Bresnahan et al.'s (1997) specification:

$$\begin{aligned} G(e^{V_i}) = e^{V_{i0}} \;&+\; \alpha_c \left[\left(\sum_{j \in C_1} e^{\frac{V_{ij}}{\rho_c}} \right)^{\rho_c} + \ldots + \left(\sum_{j \in C_k} e^{\frac{V_{ij}}{\rho_c}} \right)^{\rho_c} \right] \\ &+\; \alpha_f \left[\left(\sum_{j \in D} e^{\frac{V_{ij}}{\rho_f}} \right)^{\rho_f} + \left(\sum_{j \in F} e^{\frac{V_{ij}}{\rho_f}} \right)^{\rho_f} \right] \end{aligned}$$

$$\alpha_c = \frac{1-\rho_c}{2-\rho_c-\rho_f}\ ,\ \alpha_f = \frac{1-\rho_f}{2-\rho_c-\rho_f}$$

where C_k denotes class k (subcompact, compact, etc.), F stands for foreign products, and D stands for domestic products. The G function is here a weighted average of two one-level nested logit models. The parameters ρ_c and ρ_f must again lie in the unit interval. As either ρ goes to 1, the associated weight (α_c or α_f) goes to zero, implying that the corresponding dimension of product differentiation is irrelevant.

Note that both specifications involve the same number of distributional parameters (2); hence, one cannot motivate one parameterization as being richer than the other. Neither can one say that one model is more general than the other; they are just different. Each specification has different implications about the shape of the demand function, the substitution patterns in particular.

Implications for the Substitution Patterns and Comparison to Alternative Models: Consider first the bi-level nested logit model with "class" on the top, and "foreign vs. domestic" at the bottom. This model implies that the cross price elasticities between any two products decline as we move from the bottom to the top of the nested structure. More specifically:

$$\eta_{co,co} \geq \eta_{co,co'} \geq \eta_{co,c'o} = \eta_{co,c'o'}$$

where $\eta_{co,c'o'}$ stands for the cross price elasticity between two products, one belonging to class c and origin o, and one belonging to class c' and origin o'. Note that this structure is not as restrictive as it may at first seem. The cross price elasticities between two products belonging to different nests can in principle be as large as the cross price elasticities between two products belonging to the same nest, as the above relationship indicates. More importantly, depending on the estimated values of ρ_c and ρ_f, one may reject the nested logit structure in favor of alternative specifications, e.g. a reverse ordering of the nests (if one of (or both) ρ's is outside of the unit interval), a one level nested logit (if, for example, one of the ρ's is equal to 1, or $\rho_c = \rho_f$, or even a simple multinomial logit. Tests of the "independence of irrelevant alternatives" (IIA) property within each cluster, may further serve as a specification check, potentially indicating that a further break-down of the nests is called for. In short, while the nested logit imposes some structure on the cross price elasticities, the imposed restrictions are testable.

Nevertheless, it is unsettling that within the nested logit model described above a price change in a (co) product has the same effect on demand in the partially matching category of ($c'o$) as it has on the completely unmatching category ($c'o'$). As discussed in Bresnahan et al (1997), the second specification we considered differs with respect to the above implication. In particular, this second model implies (Bresnahan et al., 1997):

$$\eta_{co,co} \geq (\eta_{co,co'}, \eta_{co,c'o}) \geq \eta_{co,c'o'}$$

The two dimensions of product differentiation are treated here in completely symmetric way.

As mentioned above, the results we obtained when we estimated the second model did not support the structure we had imposed.[25]

[25]The estimation algorithm almost always converged to ρ values well outside of the unit interval; the only exception was France where the estimates were consistent with the restrictions implied by utility maximization. Based on these results we had to conclude that the generalized extreme value model is simply not supported by the data for the automobile market. In contrast, the results we obtained by estimating the nested logit were very reasonable.

Table 1: Summary Statistics for the European Car Market[26]

	BE	FR	GE	IT	UK	ALL
1980 value-added tax (in %)	25	33	13	18	23	
1990 value-added tax (in %)	25	25	14	19	24	
Total sales (in 1,000 units)	384.4 (48.9)	1920.3 (192.1)	2508.9 (359.7)	1908.0 (293.4)	1704.1 (248.9)	8412.3 (892.4)
Parallel imports (in 1,000 units)	N/A	5–40	30–60	10–75	1–50	
Japanese market share (in %)	21.6 (1.9)	3.1 (.5)	15.5 (1.5)	1.8 (1.3)	11.3 (.6)	7.7 (1.0)
Japanese quota (in %)	—	3.0	15.0	1.0	11.0	
Domestic market share (in %)	2.5 (.4)	66.6 (5.1)	70.2 (4.0)	58.2 (6.2)	55.1 (4.0)	
European average (in %)	1.6 (.5)	24.6 (2.6)	33.4 (1.9)	16.7 (1.4)	12.1 (1.5)	
C1-ratio (in %)	16.3 (1.8) (VW)	33.5 (1.7) (PSA)	30.2 (1.2) (VW)	53.9 (5.2) (Fiat)	28.7 (3.3) (Ford)	15.7 (1.6) (Fiat)

[26] Annual averages over 1980–93. Standard deviations in parenthesis.

Table 2. Hedonic regressions

Dependent variable: log(price) in ECUs

Number of Observations: 7212

Variable	Coefficient (Std)	Coefficient (Std)
Constant	7.9562	7.6668
	(.0794)	(.0784)
Horsepower	.0076	.0076
	(.0001)	(.0001)
Weight	.0003	.0003
	(0)	(0)
Width	.0051	.0051
	(.0004)	(.0004)
Height	−.0011	−.0011
	(.0003)	(.0003)
Log value-added tax	.0526	−.1561
	(.0250)	(.0229)
Exr	—	.0207
		(.0037)
Mini	−.5119	−.5119
	(.0142)	(.0142)
Small	−.4171	−.4171
	(.0121)	(.0121)
Medium	−.3438	−.3438
	(.0111)	(.0111)
Large	−.2546	−.2546
	(.0094)	(.0094)
Executive	−.1493	−.1493
	(.0088)	(.0088)
Luxury	−.2142	−.2142
	(.0102)	(.0102)
Sports (normalized)		
Firm dummies	p-value: .0000	p-value: .0000
Source/time dummies	p-value: .0000	p-value: .0000
Market/time dummies	p-value: .0000	p-value: .0000
R-squared	.97	.97

Table 2 continued: 95% Confidence Interval for Destination Effects[27]

Country \ Year	1980	1985	1990	1992	1993
France	10–18	16–22	2–8	3–9	2–7
Germany	–	4–12	(−3)–5	3–12	5–13
Italy	11–17	19–27	10–16	12–19	(−3)–2
United Kingdom	30–36	26–32	16–22	13–19	(−3)–5

[27]Based on hedonic price regression, column 1 of Table 2. All numbers are percent deviations from the hedonic prices in Belgium that are normalized to zero every year.

Table 3: Pricing-to-market Coefficients

(by source country and destination market)

$$\beta_{sm} = \beta_s + \beta_m$$

$$\beta_{sm} = 1 \Longrightarrow \text{local currency price stability}$$

$$\beta_{sm} = 0 \Longrightarrow \text{full exchange rate pass-through}$$

	Be	Fr	Ge	It	Ja	Ko	NL	Sp	Sw	UK
β_s	.4176 (.2582)	.4445 (.0574)	.2478 (.0556)	.5684 (.0664)	.2017 (.0687)	.2096 (.3528)	.1632 (.2800)	.7395 (.1918)	.6905 (.1172)	.7910 (.0977)
S_s	.01	.25	.33	.17	.07	.001	.005	.028	.01	.12
β_m		.0281 (.0338)	−.0757 (.0470)	−.0132 (.0420)						.2024 (.0360)

Table 4: Results from the Demand Estimation[28]

(selected coefficients)

	Be	Fr	Ge	It	UK
Number of observations	1351	1196	1077	1027	1221
ρ_f	.39	.31	.42	.45	.35
	(.12)	(.06)	(.10)	(.07)	(.06)
ρ_c	—	.56	.39	1.01	.44
		(.11)	(.13)	(.23)	(.14)
α	125.9	125.9	125.9	125.9	125.9
	(9.6)	(9.6)	(9.6)	(9.6)	(9.6)
Wald-statistic for $\rho_f = \rho_c$	—	5.34	.16	5.77	.31
Preference for:					
France	.37	.95	.69	.39	$-.65$
	(.32)	(.15)	(.25)	(.15)	(.11)
Germany	.99	.38	1.68	.65	$-.15$
	(.36)	(.11)	(.45)	(.17)	(.10)
Italy	.76	$-.01$	.70	1.25	-1.06
	(.30)	(.09)	(.24)	(.25)	(.18)
Japan	.70	$-.08$	.74	$-.45$	$-.90$
	(.30)	(.08)	(.28)	(.21)	(.14)
Spain	.07	$-.12$	.47	.21	-1.58
	(.32)	(.16)	(.32)	(.24)	(.25)
Sweden	1.01	$-.01$	.89	.64	$-.61$
	(.34)	(.11)	(.27)	(.24)	(.14)

[28]Standard errors in parenthesis.

Table 5: Results from Estimation of the Pricing Equation[29]

	Bertrand Instr. I	Bertrand Instr. II	Collusion in UK Instr. II
Constant	-3.03	-5.60	-5.88
	(1.16)	(1.33)	(1.37)
Foreign firm disadvantage	.04	.04	.03
	(.59E-02)	(.62E-02)	(.63E-02)
France fixed effect	.05	.04	.04
	(.013)	(.014)	(.013)
Germany fixed effect	.17	.17	.17
	(.014)	(.015)	(.015)
Italy fixed effect	.06	.05	.05
	(.014)	(.014)	(.014)
UK fixed effect	.17	.15	.02
	(.013)	(.014)	(.014)
LHW	0.46	.43	.44
	(.02)	(.02)	(.02)
LWI	1.53	2.10	2.20
	(.17)	(.23)	(.23)
LHE	0.46	.56	.56
	(.17)	(.18)	(.18)
LQU	-.58E-02	-.03	-.03
	(.30E-02)	(.47E-02)	(.48E-02)
$LWAGE$	.25	.38	.39
	(.04)	(.04)	(.04)
$LWAGELOC$	.37	.39	.40
	(.01)	(.02)	(.02)
$LWAGEU$	—	$-.10E-02$	$-.75E-02$
		(.20E-02)	(.20E-02)
Time dummies	p-value: .00	p-value: .00	p-value: .00
Model fixed effects	p-value: .00	p-value: .00	p-value: .00
Japanese quota year dummies	p-value: .00	p-value: .00	p-value: .00

[29] *Acronym explanation:*

LHW:	log of horsepower/weight
LWI:	log of width
LHE:	log of height
LQU:	log of world production quantity
$LWAGE$:	log of the wage of the producing country
$LWAGELOC$:	log of wage in the destination country times the exchange rate between source and destination country (source currency units/destination currency units)
$LWAGEU$:	LWAGELOC * Dummy for the UK

Table 5 continued: Results from Estimation of the Pricing Equation

	Impose CRS	Impose IRS	Impose DRS
Constant	−3.13	−7.64	1.37
	(1.22)	(1.43)	(1.45)
Foreign firm disadvantage	.04	.04	.03
	(.59E-02)	(.63E-02)	(.68E-02)
France fixed effect	.06	.06	.06
	(.013)	(.014)	(.015)
Germany fixed effect	.18	.19	.17
	(.014)	(.015)	(.016)
Italy fixed effect	.07	.07	.06
	(.014)	(.014)	(.015)
UK fixed effect	.17	.17	.16
	(.013)	(.014)	(.015)
LHW	0.47	.41	.53
	(.02)	(.02)	(.02)
LWI	1.52	2.47	.57
	(.17)	(.18)	(.20)
LHE	0.47	.74	.21
	(.17)	(.18)	(.20)
$LWAGE$	.23	.55	-0.09
	(.03)	(.04)	(.04)
$LWAGELOC$	.38	.37	.39
	(.01)	(.01)	(.01)
Time dummies	p-value: .00	p-value: .00	p-value: .00
Model fixed effects	p-value: .00	p-value: .00	p-value: .00
Japanese quota year dummies	p-value: .00	p-value: .00	p-value: .00

Table 6: Substitution Patterns and Markups in the 5 European countries (in 1990)

	Be	Fr	Ge	It	UK
Average price elasticity $\left(\frac{\partial \ln s_i}{\partial \ln p_i}\right)$	-5.77	-5.60	-5.74	-4.09	-6.21
for domestic cars $\left(\frac{\partial \ln s_i^d}{\partial \ln p_i^d}\right)$	-5.77	-5.37	-6.09	-3.63	-6.03
for foreign cars $\left(\frac{\partial \ln s_i^f}{\partial \ln p_i^f}\right)$	—	-5.98	-4.96	-4.51	-6.45
Average cross-price elasticity $\left(\frac{\partial \ln s_i}{\partial \ln p_j}\right)$	.035	.047	.035	.029	.040
for domestic cars $\left(\frac{\partial \ln s_i^d}{\partial \ln p_j^d}\right)$	.035	.170	.015	.111	.076
for foreign cars $\left(\frac{\partial \ln s_i^f}{\partial \ln p_j^f}\right)$	—	.054	.033	.038	.044
between domestic and foreign $\left(\frac{\partial \ln s_i^d}{\partial \ln p_j^f}\right)$	—	.015	.043	.220E-05	.030
between foreign and domestic $\left(\frac{\partial \ln s_i^f}{\partial \ln p_j^d}\right)$	—	.024	.036	.913E-05	.025
Average markup	0.19	0.21	0.19	0.33	0.17
Average markup for domestic cars	0.19	0.22	0.18	0.40	0.17
Average markup for foreign cars	—	0.18	0.22	0.23	0.16

Table 7: Substitution Patterns for Selected Models in Belgium and Italy (in 1990) [30]

	Belgium			Italy			
	Own Elasticity	Cross-Elasticity		Own Elasticity	Cross-Elasticity		
		same group	different group		same subgroup	diff. subgr., same group	different group
Alfa 33	-5.11	.060	.53E-03	-3.16	.399	.45E-02	.22E-02
	(.72)	(.013)	(.32E-04)	(.49)	(.093)	(.046)	(.26E-0
Citroën AX	-3.07	.198	.15E-02	-2.02	.104	.13E-02	.88E-03
	(.42)	(.042)	(.90E-04)	(.32)	(.024)	(.87E-02)	(.10E-0
Fiat Uno	-2.96	.113	.85E-03	-1.55	.534	.011	.70E–02
	(.41)	(.024)	(.51E-04)	(.23)	(.124)	(.069)	(.83E-0
Fiat Tempra	-5.08	.019	.15E-03	-2.80	.656	.50E-02	.19E-02
	(.72)	(.41E-02)	(.93E-05)	(.44)	(.153)	(.061)	(.22E-0
Ford Escort	-4.26	.210	.19E-02	-2.86	.151	.15E-02	.71E-03
	(.59)	(.045)	(.11E-03)	(.45)	(.035)	(.015)	(.84E-0
Opel Corsa	-3.05	.242	.18E-02	-2.00	.061	.77E-03	.51E-03
	(.42)	(.052)	(.11E-03)	(.32)	(.014)	(.50E-02)	(.61E-0
Toyota Corolla	-3.93	.247	.22E-02	-4.55	.78E-02	.76E-04	.37E-04
	(.54)	(.052)	(.13E-03)	(.73)	(.18E-02)	(.77E-03)	(.43E-0
VW Golf	-3.98	.563	.50E-02	-2.40	.578	.56E-02	.27E-02
	(.53)	(.120)	(.30E-03)	(.37)	(.135)	(.056)	(.32E-0

[30]Standard errors of the estimated elasticities are in parentheses.

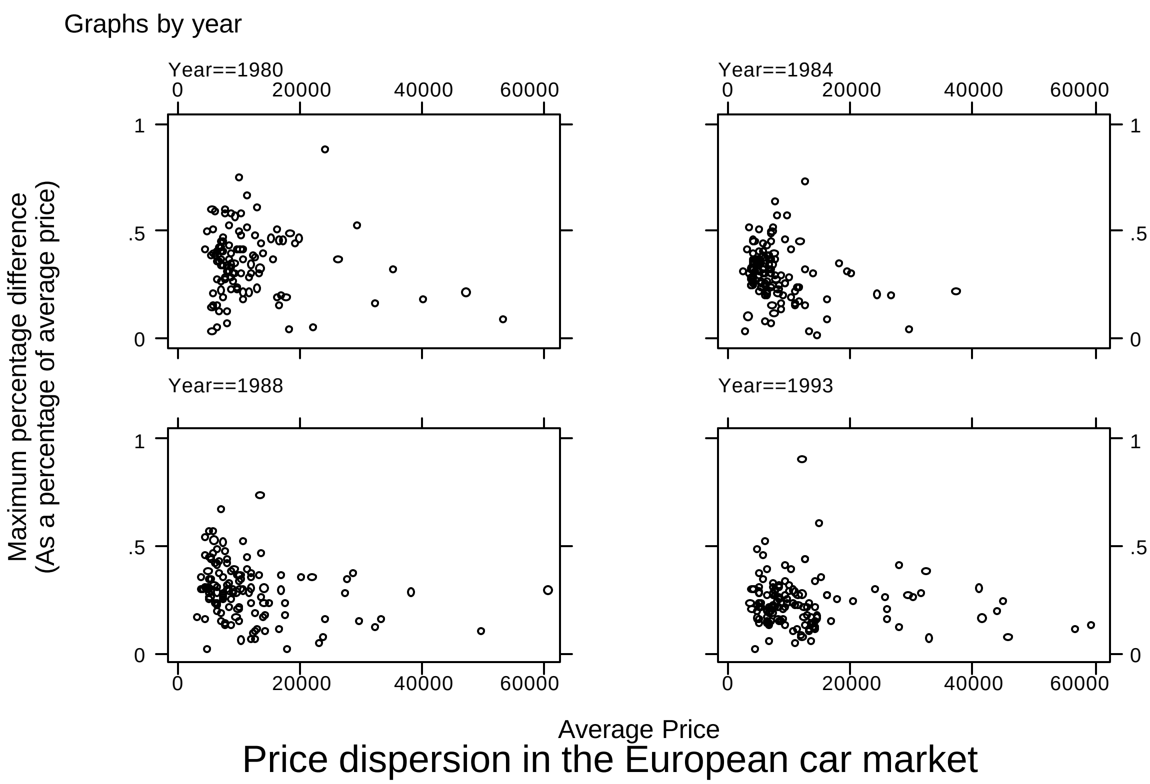
Graphs by year
Year==1980
Year==1984
Year==1988
Year==1993
0
20000
40000
60000
1
.5
0
Maximum percentage difference
(As a percentage of average price)
Average Price
Price dispersion in the European car market

Figure 1:

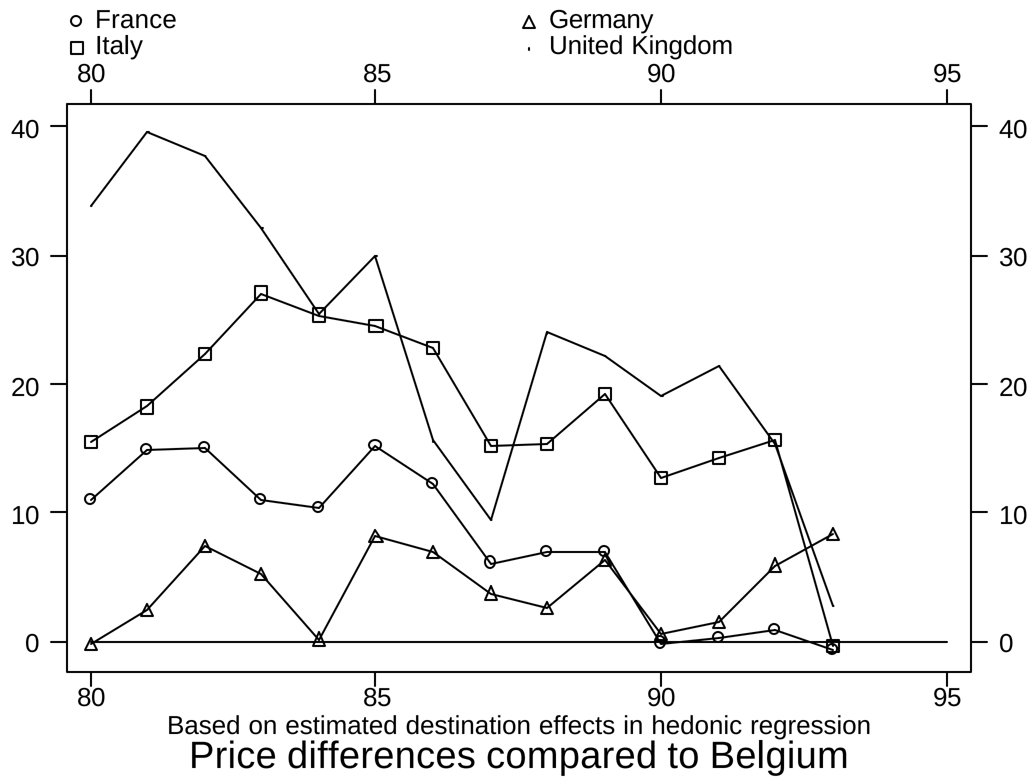
France
Germany
Italy
United Kingdom
80
85
90
95
40
30
20
10
0
Based on estimated destination effects in hedonic regression
Price differences compared to Belgium

Figure 2:

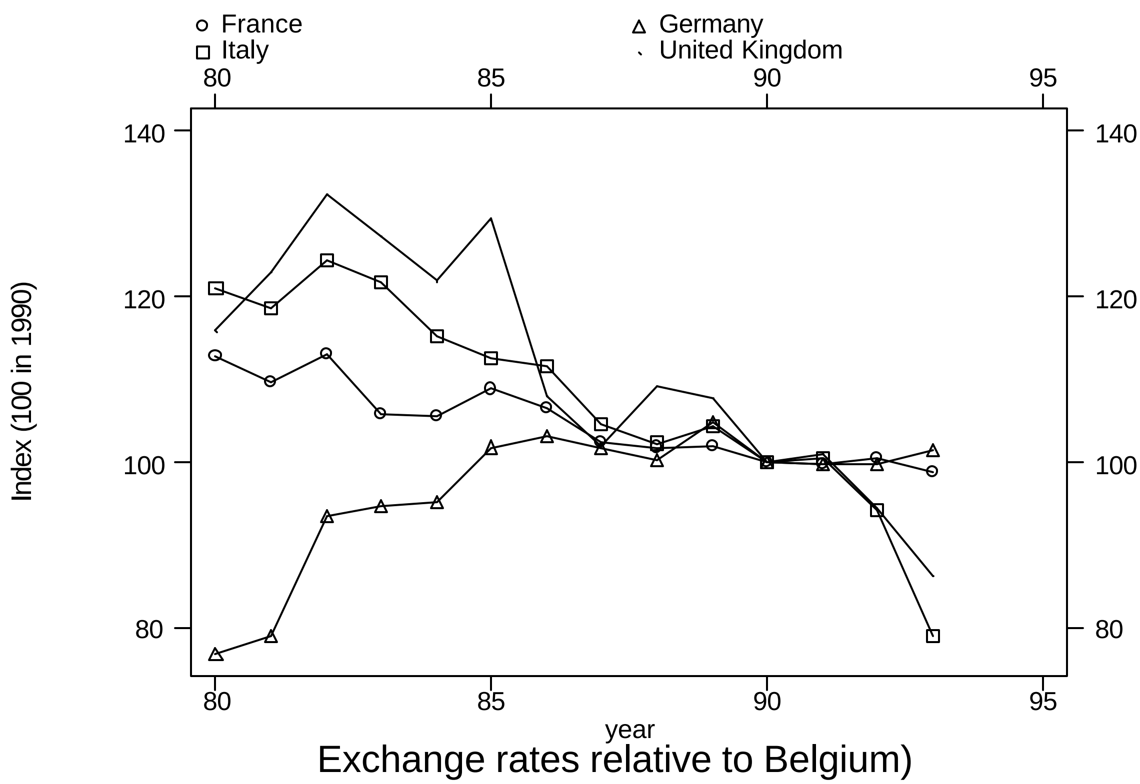
France
Germany
Italy
United Kingdom
80
85
90
95
140
120
100
80
Index (100 in 1990)
year
Exchange rates relative to Belgium)

Figure 3:

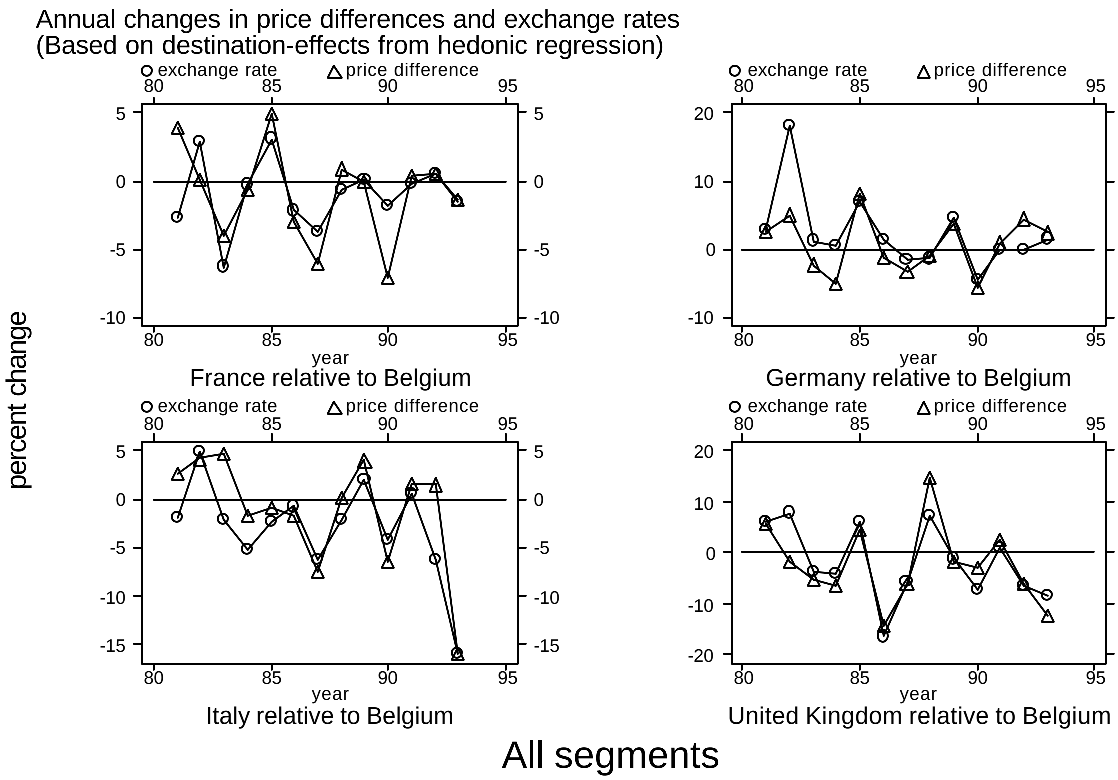

Annual changes in price differences and exchange rates
(Based on destination-effects from hedonic regression)
exchange rate
price difference
percent change
year
France relative to Belgium
Germany relative to Belgium
Italy relative to Belgium
United Kingdom relative to Belgium
All segments

Figure 4:

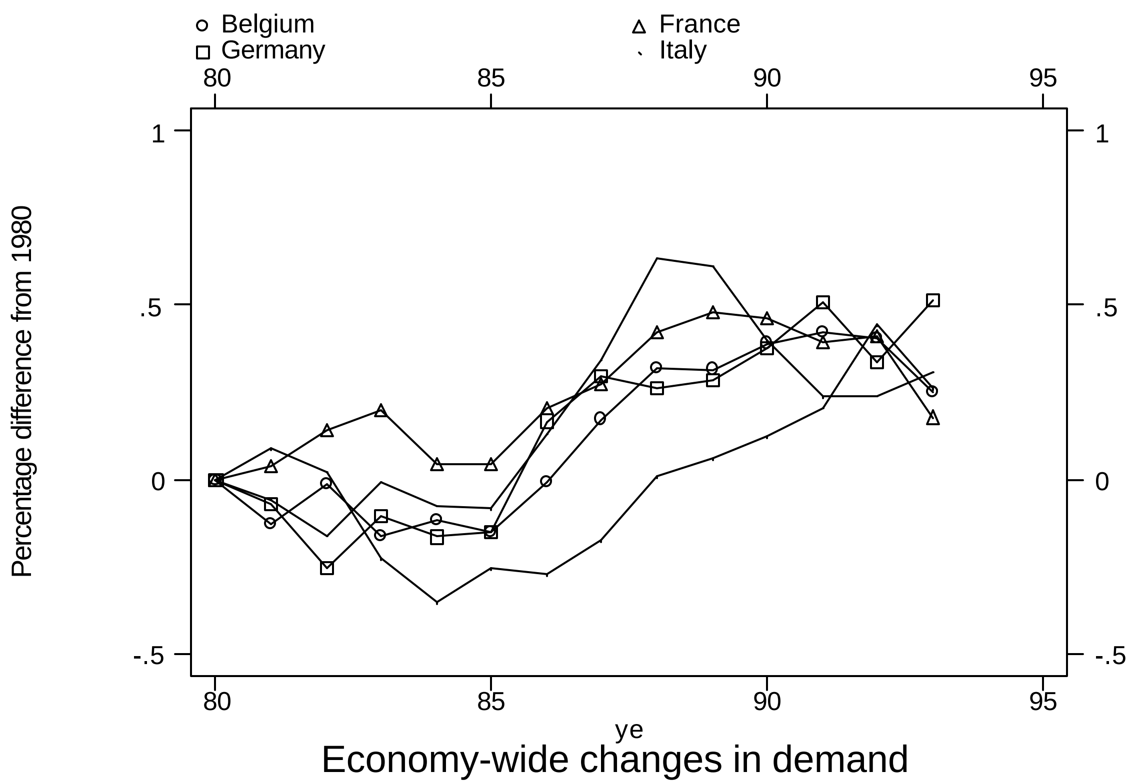

Belgium
Germany
France
Italy
80
85
90
95
1
.5
0
-.5
Percentage difference from 1980
ye
Economy-wide changes in demand

Figure 5:

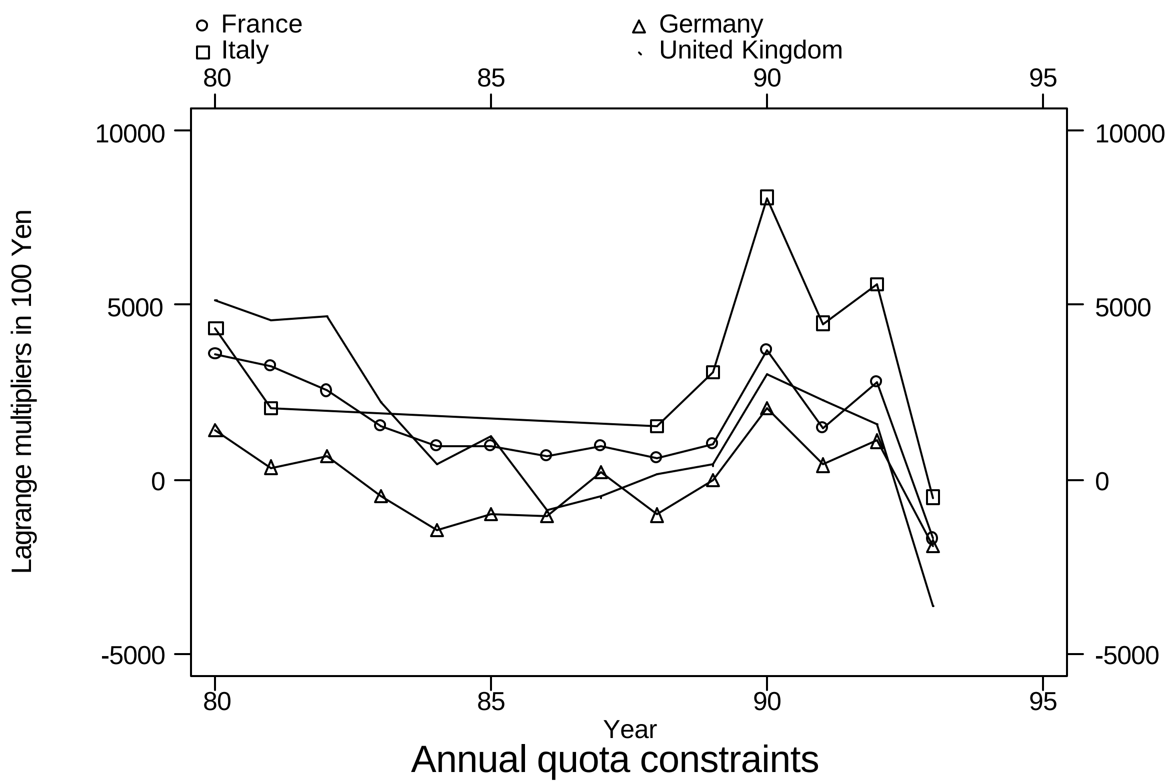
France
Germany
Italy
United Kingdom
80
85
90
95
10000
5000
0
-5000
Lagrange multipliers in 100 Yen
Year
Annual quota constraints

Figure 6: