RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

School of Public Policy
University of Michigan
Ann Arbor, Michigan 48109-1220

Discussion Paper No. 392

Productivity in the U.S. Auto Industry: Evidence from Multiple Perspectives

James Levinsohn
University of Michigan
National Bureau of Economic Research

January 28, 1997

Productivity in the U.S. Auto Industry: Evidence from multiple perspectives

by

James Levinsohn
University of Michigan
National Bureau of Economic Research

Current version: January 28, 1997

Address. Department of Economics, University of Michigan, Ann Arbor, MI 48109; Telephone:(313) 763-2319; Fax:(313) 754-2769; e-mail: JamesL@umich.edu

Productivity in the U.S. Auto Industry:
Evidence from multiple perspectives

James Levinsohn
University of Michigan
National Bureau of Economic Research

1. Introduction

Productivity in the U.S. automobile industry is sort of like religion. Lots of people have very strongly held views about a concept that few can firmly define and none can persuasively measure. In principle, comparisons of productivity are simple. If a firm can produce the same output with less inputs, it is considered to be more productive. Applying this widely-agreed-upon definition to the U.S. auto industry, though, is problematic.

There are at least three main reasons for this. First, outputs of the multi-product firms are not homogeneous. Second, input mixes differ widely and accounting for fixed costs, while enormously important, is difficult. And third, factors that are completely outside the control of the firm, such as foreign currency exchange rates, can wreak havoc on costs. In this environment, productivity comparisons get fuzzy. A few examples illustrate these points.

The problems introduced by the heterogeneity of outputs are pretty obvious. If Toyota could produce a Buick LeSabre at lower cost than can GM, then most would agree that for this particular product Toyota is more productive. The problem of course is that Toyota doesn't produce the Buick LeSabre. Toyota produces its own flavor of a mid-to-large size sedan (the Camry) and one never observes what it would have cost Toyota had Toyota produced the Buick product.

Suppose for a moment, though, that GM and Toyota really did produce an identical (or homogeneous) product so the two firms' outputs were the same. Their inputs (and not just the mix thereof) would still differ. The Camry is assembled in a fairly new and highly automated factory in Kentucky. The LeSabre is assembled in an older factory in Flint, Michigan. Suppose for the sake of argument that the Camry is assembled with many fewer worker hours but with

This paper is part of the NBER/Sloan program on productivity. I am grateful to Mustafa Mohatarem of General Motors and Martin Zimmerman of Ford for facilitating plant visits and to plant managers who generously gave their time.

much more expensive and less depreciated machinery. A simple comparison of labor hours used for assembly would not be an accurate measure of productivity. If the machinery that replaces the labor gets expensive enough, it will be more efficient to hire more labor even if this drives up the factory wage. One might imagine then trying to compute the user cost of the capital and treat it appropriately. But this too is problematic since its value is sometimes simply not known and in other cases is proprietary information. Further clouding matters, some assembly plants also include stamping facilities in which the body parts are formed while others bring such parts in from outside the factory, and some assembly plants include their own paint shops while others do not. Add to all this the fact that many factories produce multiple products so attributing the fraction of capital costs to each particular product is tricky at best. The heterogeneity of inputs, then, also contributes to the difficulty in precisely comparing productivity.

Finally, the international aspect of the U.S. auto industry introduces subtleties to discussions of productivity. Discussions of productivity typically focus on firms' technologies and costs. But if a firm is either receiving revenues and/or paying costs in a foreign currency, exchange rate fluctuations will impact the bottom line. Are these fluctuations (which are typically responses to decisions made by central bankers and legislatures and have nothing to do with how firms produce cars) part of the productivity calculus? For example, does Acura become less productive when the yen falls relative to the dollar? If so, does this in turn make Mazda more productive since they can move assembly to their existing U.S. plants?

All of these issues-- heterogeneous outputs, big fixed costs, heterogeneous inputs, and multi-product plants, and exchange rate issues-- make the discussion and measurement of productivity difficult. One of the goals of this paper is to highlight some of these difficulties. Another goal is to investigate whether the application recent econometric methods can inform the productivity discussion. There have been recent advances in estimating the demand and cost of differentiated products using product-level data. Can these methods be used to further our understanding of productivity in a differentiated product industry? And do these methods give rise to a view of productivity that is consistent with or at odds with the the related (non-econometric) literature on productivity in the auto industry?

This paper is organized as follows. Section 2 contains a critical review of the most cited productivity studies of the U.S. auto industry. Section 3 outlines two econometric frameworks for investigating productivity-- one using proprietary plant level data and one using product-level data. The former is the subject of the paper by Berry, Kortum, and Pakes for this project. Section 4 contains results of the latter approach. Section 5 concludes and summarizes what one can and cannot learn from the various approaches to measuring productivity in the auto industry.

2. A Critical Review of the Literature

In an industry as huge as the U.S. automobile industry, it is not surprising that there have been numerous attempts to measure productivity in the industry. Two of these studies have received much more attention than the rest. They are the International Motor Vehicle Program at the Massachusetts Institute of Technology and the Harbour Report produced annually by a consulting firm located near Detroit. Each of these reports is reviewed below. Together, they provide a picture of productivity in the industry which serves as a benchmark against which econometric evidence can be compared.

MIT's International Motor Vehicle Program

The International Motor Vehicle Program (IMVP) at MIT has studied issues relating to productivity in the auto industry for well over ten years. This program has spawned two widely read books on the industry, ‘‘The Future of the Automobile’’ and ‘‘The Machine that Changed the World.’’ The latter of these books was specifically directed at what economists might consider the issue of productivity. The basic idea in the book was that the design and assembly of automobiles varied drastically throughout the world and that the methods used principally in Japan and given the moniker of ‘‘lean production’’ were going to change the way cars would be made in the coming years.

The concept of lean production encompassed many aspects of the production process. ‘‘The Machine that Changed the World’’ specifically considered the design process, the supply chain, and the assembly of cars. Although not typically considered the production process, the book also examined the varying methods of retailing cars too. Many of the notions put forth in the book are by now well known.

The book emphasized the importance of being able to move quickly from the idea to design to product in an era of changing consumer tastes. The benefit of integrating the design process, as opposed to having lots of individual and separate design teams, was also stressed. The IMVP documented the large differences between Japanese and American firms in terms of the time and resources required by the design process. During the mid-1980's, they note that U.S. car manufacturers used almost twice as many engineering hours per new car developed and took about 40 percent longer to bring a new car to market. They note that U.S. firms used about twice as many employees in a project team and that only one in six Japanese products was delayed (relative to its initial schedule) compared to half of American products. Finally, they note that it took U.S. firms an average of 11 months to bring quality back to normal after introduction of a new model while it took Japanese firms only one and a half months to do so. All of these figures suggest that the Japanese were more productive in terms of product design. Still, the caveats given in the introduction of this paper apply. First, the cars produced by U.S. and Japanese firms were, in the early and mid-1980's, quite different. Second, none of these comparisons makes note of factor prices. Across countries, using fewer engineer-hours doesn't necessarily translate into being less expensive. Despite these caveats, though, the picture painted in ''The Machine that Changed the World'' is one in which the design of cars in the U.S. seems, on the surface anyway, pretty inefficient. The response of the U.S. industry has been a movement toward many of the practices previously associated with the Japanese firms, and this too suggests that the design practices of the Japanese firms really were more efficient.

The IMVP study also studied differences in the relationships between the auto firms and their suppliers. The study emphasized three ways in which Japanese firms were different from their U.S. counterparts in terms of the supply chain. First, Japanese firms relied much more heavily on multi-tier suppliers from outside the firm. Put another way, firms were less vertically integrated and more likely to instead rely on outside firms to supply parts. The productivity implication here is that less vertical integration was more efficient. Second, when dealing with a supplier, many of the decisions regarding manufacturing were left to the supplier. The IMVP study paints a picture in which the car companies would often present the supplier with the desired specifications and leave engineering decisions to the supplier. The productivity implication is that by allowing suppliers, who would frequently have long-term relationships with the car firms, more freedom, the suppliers would be free to be more efficient. The presence of long-term relationships between suppliers and producers was important, though, since in the absence of this relationship, suppliers might be hesitant to undertake the necessary capital investment. The third big difference in the supply chain was the early adoption of just-in-time inventory practices by Japanese producers. Less inventory means less carrying costs, less space devoted to storage (and the attendant costs), and, by implication, more productivity. Again, the picture painted is one in which the system of lean production imparted productivity advantages to the Japanese firms. This advantage, while very nicely documented, is never really quantified. To varying extents, all these differences in the supply chain have been recently adopted by the U.S. firms. Visits to Ford and GM plants vividly illustrated the just-in-time inventory system. Rather than warehouses full of parts, parts were delivered several times a day by truck and, in some cases, rail. This has not been a costless transition, though. The recent strike at a brake plant in Dayton, Ohio brought several GM assembly plants to a crawl and cost GM dearly. Just-in-time inventory practices can be very costly when labor relations are not smooth, and this illustrates the point that a practice that enhances productivity in one country may not work in another.

The third area the IMVP study focussed on was the assembly of cars. Here, the study highlighted several differences between Japanese practices and then-domestic practices. First, line workers in Japan could stop the line when necessary. In the U.S., the faulty cars continued through the line and would later (hopefully) be corrected. Second, (though related to the first), the rework areas in Japanese factories were much smaller or altogether absent. This is the area where problems with a car are remedied. This is much more of a craftsman-like procedure and it is costly. Third, the assembly plants in Japan put a high weight on quality and quantity while it was the IMVP's impression that U.S. manufacturers discounted the importance of the former. Visits to U.S. assembly plants suggested that all these differences have been adopted (at least by the plants visited.) Workers could now stop the line, banners throughout the factory urge workers to make ‘‘Quality Job One’’ (at a Ford plant), and rework areas even at the Cadillac plant were small. Again, the general impression is one in which the U.S. firms have adopted several of the more productive Japanese practices when it comes to assembly.

Six years after the publication of this book, the authors appear to have been pretty much on target. Many of the practices highlighted by the study have been adopted by U.S. firms. Insofar as these practices represented a productivity advantage for Japanese firms, that advantage would be expected to have narrowed. On the other hand, Japanese firms have not stood still while American ones have played catch-up.

The Harbour Report

While ‘‘The Machine that Changed the World’’ highlighted several practices that impacted productivity, productivity was not its primary focus. The Harbour Report is an annual report put out by the consulting firm of Harbour and Associates, Inc. and it is indeed focussed on productivity. The Report was first issued in 1992. The subtitle of the 1995 annual report reads ‘‘Manufacturing Productivity, Company by Company -- Plant by Plant.’’ Several of the authors of the study have worked in the auto industry for years, and their experience lends credibility to the study. The news media has also taken note of the Harbour Report, and its release is typically accompanied by lengthy articles in publications such as the New York Times and the Wall Street Journal as well as by the leading trade publication-- Automotive News. The bottom line, then, is that the Harbour Report has earned a fair amount of credibility and visibility and its focus on productivity makes it a natural focal point.

The Harbour Report makes use of detailed plant level information provided by (most of) the major automobile producers in the U.S. The report is especially useful for cross-sectional comparisons at a particular point in time, but it is not as useful for a longer time series. That is, the Report carefully highlights differences between firms (and plants) in a given year, but it is less useful when considering, for example, how productivity at a given firm has changed over the last 20 years. It simply does not address the latter. Nonetheless, the Report highlights some of the ways productivity is measured in the auto industry. It is also interesting because it highlights differences between how engineers think about productivity and how economists think about it. With a critical eye, the 1995 Harbour Report’s picture of industry productivity is reviewed.

The first section of the Report is devoted to auto assembly, and this is the section which traditionally attracts the most attention. The adopted measure of productivity is workers per vehicle produced. Using very careful measures of the number of workers and the *daily* output of vehicles, the Report computes that Ford used 3.12 workers per vehicle (excluding Mexican production), Chrysler 3.41, and GM 3.86. Their was substantial intra-firm variation in this figure. For example, in the case of GM, the Bowling Green plant required 9.70 workers per vehicle while its Spring Hill plant required only 3.09. The analogous figures for the transplants (Japanese plants located in the U.S.) were uniformly smaller. Honda required 2.58 workers per vehicle at its Marysville plant while Toyota used 2.42 at its original Georgetown plant.

Do these figures mean that Ford is more productive than GM, Honda more productive than Ford, and Spring Hill more productive than Bowling Green? Not necessarily, and this is why these worker per vehicle figures, while widely cited, can be misleading. As noted at the outset, cars are not homogeneous. The Bowling Green plant produces the Chevrolet Corvette while the Spring Hill plant produces the Saturn models. The 1994 list price for the Corvette ranged from $38,285 to $67,543, while Saturns were between $9995 and $12,595. Clearly, a measure of value-added per worker might be more informative than a simple measure of cars per worker (or the inverse thereof.)

There are reasons beyond product heterogeneity to view the worker per vehicle numbers as perhaps not adequately representing productivity. Wages might matter. In fact, this is not as big an issue as it could be due to the common bargaining framework in much of the industry. Although pension burdens vary substantially by company, the wage structure does not. Transplant firms are not subject to the same rules since many are not unionized. But there is little evidence of these plants using more labor since it is possibly cheaper. Indeed, the evidence points the opposite way. Mexican plants, though, do pay lower wages and, in some cases they also use substantially more labor. Ford's Cuautitian plant used 7.93 workers per vehicle while GM's Mexico City plant used 10.63 workers per plant. (The exception to this pattern is Ford's Hermosillo plant which required only 2.99 workers per vehicle.)

Also, labor and capital are, to some degree, substitutes in assembly. Plant visits illustrated this point vividly. Some plants had replaced entire sections of workers with huge, complex, and very expensive robotics. One reason behind the greater vehicles per worker figure for the transplants may be that the transplants are, on average, newer plants, and newer plants are often more highly automated. From an economic as well as business perspective, fewer workers per vehicle all else equal does indeed imply greater productivity. But all else is not equal. Absent such equality, a focus on costs is more informative than a focus on workers per vehicle. Costs, though, are harder to measure and frequently quite proprietary.

Finally, assembly is not a huge component of the value added to a vehicle. The manager of a plant that produces a very expensive luxury car estimated his total costs of assembly per car (labor and amortized capital costs) to be in the range of $3500. While it is not clear how he came up with this figure, it does suggest that while it's nice to be efficient at assembling cars, assembly is not the entire picture.

The assembly productivity figures in the Harbour Report are cross-sectional. They vary across assembly plants. There is also some minimal time series evidence in that figures from one year previous are also given. These figures have the same drawbacks (discussed immediately above) that the current figures have. They yield some interesting insights, though. The intra-firm variation in changes in workers per vehicle typically is greater than the between firm averages. For example, while the workers per vehicle figure for Ford fell from 1993 to 1994, almost half of the Ford plants experienced increases in this measure. Much of the intra-firm variation is due to model changeovers. When a new model is introduced, the line speed is slower and output less. As workers learn and the glitches are worked out, the line speed rises and workers per vehicle falls. In any given year, several plants will experience these model changeovers. Other plant-level annual changes in the workers per vehicle figures are attributed to the purchase of new machinery, and this relates to the previous point regarding capital and labor substitutability.

Perhaps because assembly is not the entire picture, the Harbour Report also includes detailed productivity information on stamping and power train production. Given that the actual assembly is only a small part of the total value-added, productivity in the other stages of manufacturing is a very important part of the total productivity picture. The organizational structure varies across firms and this clouds productivity comparisons. For example, the Japanese transplants purchase many parts from outside suppliers whereas the domestic firms tend to do less outsourcing. Even among the Big Three, there are substantial differences. GM is thought to produce about 70 percent of its parts in house, Ford about 50 percent, and Chrysler 30 percent. (New York Times, June 25, 1996) While having efficient sources can contribute to lower vehicle costs for, say Toyota-U.S., does this make Toyota more productive? The question suggests again that perhaps costs, not inputs per output, might provide a more reasonable basis for comparison in an industry where the degree of outsourcing varies substantially by firm.

Stamping of parts is a significant step in the production of cars and the Harbour Report provides several measures of productivity in stamping. Stamping is typically done in-house and is not often outsourced, for it is viewed as ‘‘critical to all the other elements of making a car.’’ (Harbour, 1995, p. 36.) Stamping is also very expensive and accounts for a big part of fixed costs. Newer plants almost always have the stamping done where assembly takes place. Some plants have contiguous stamping facilities, but stand-alone stamping facilities appear to be on their way out.

The economics of the move away from stand-alone facilities is basically a story of returns to scale and geography. In the 1970's, the number of truly different models produced by a firm was much smaller than it is now and the production per model was much higher. With high volume, it made more sense to build huge stand-alone stamping plants that would service several assembly plants in different locales. As the number of models produced by a firm has increased, a stand-alone plant would have to change dies more frequently, and this is costly. Instead, the newer plants have been constructed with stamping facilities on location. This also has the benefit of reducing transport costs and better accommodating the just-in-time inventory practices that are becoming the standard. Not surprisingly, all of the transplant factories have on site stamping facilities.

The Harbour Report has several measures of stamping productivity. A typical car has about 300 stamped parts, but the car companies usually focus on producing about 100 of these in house. These 100 are the most important parts and Harbour counts these parts as a generic stamped vehicle. The Report then provides four measures of productivity for stamping based on this generic stamped vehicle. These measures are: i) workers per vehicle; ii) vehicles per line; iii) hits per hour; and iv) sets of dies per line. The first attempts to measure labor productivity, the second and third equipment productivity, and the last is an indicator of operating complexity. Based on these measures, the Harbour Report finds that the transplants are the most productive. Using the transplants as the benchmark, Ford and Chrysler are reported to be less efficient than transplants while GM is way behind its domestic rivals. Due to the multiple measures of productivity, it is hard to find a single summary measure. Unfortunately, Harbour's methods are new and there are not any comparisons over time.

In addition to assembly and stamping, the Harbour Report also provides some productivity measures on power-train (engine and transmission) production. As with assembly, productivity is measured by workers per unit produced (engines in this case.) This sort of measure seems quite useful to measure how plant performance changed absent any capital investment and absent product changes. It is less useful when either of these change or when comparisons are made across plants or firms. The same caveats that apply to assembly apply to engine production. With these caveats in mind, the productivity rankings are by now familiar-- the transplants are (by a factor of 2) the most productive, followed by Ford, Chrysler, and finally GM. GM's and Ford's least productive plants (keeping in mind the metric is workers per engine) are both located in Mexico.

Again, if wages are low enough, these plants might still be quite cost-effective even with the high workers-per-engine. Unfortunately, the Harbour Report works more with engineering data than economic data.[1]

Summarizing, the Harbour Report provides detailed and very useful information about firms at a particular point in time. The general message is one in which the transplants are the most productive, followed by either Ford or Chrysler depending on what is being measured, with GM being the least productive. The same measures of productivity that inform these firm rankings suggest that variation within firms is very substantial.

The IMVP project and the Harbour Report combine to suggest that U.S. producers lag behind Japanese firms, but that the gap may be narrowing as more of the production practices previously associated with Japanese manufacturing are adopted by U.S. firms. The Harbour Report does not give much information on European firms, while the IMVP study provides some evidence that the European firms are not as productive as the Japanese firms. In order to get a better view of productivity trends, one must turn to a longer time series. This means, in this instance, turning to less detailed data. The next section of this paper applies recent econometric methods to a 25 year panel of data. The section investigates what these newer methods imply about productivity and asks whether the econometric evidence is broadly consistent with the picture painted by the IMVP and Harbour.

3. An Econometric Approach to Productivity and Costs in the Automobile Industry

Econometric estimates of costs using firm-level data have a long history. This literature, though, consistently assumes a homogeneous product industry-- an assumption that is simply inappropriate for the U.S. auto industry. A first pass at estimating cost surfaces using plant-level data for a differentiated product industry is Berry, Kortum, and Pakes (hereafter Kortum et. al. in order to differentiate this approach from the other Berry approach outlined below) in a paper prepared for the Sloan conference on productivity. The approach used in that paper is briefly reviewed here and contrasted with the approach used in this paper. The main differences between the two approaches are dictated by their respective data requirements. Kortum et. al.'s approach

[1] The Harbour Report does include some financial data. Much of this information consists of repackaging information contained in firms' annual reports coupled with their estimates of labor usage.

required plant-level data while the approach used in this paper requires only product-level data. The latter are much more easily obtained, but, for some purposes, less informative. Both approaches, though, are explicitly designed to accommodate the differentiated products nature of the auto industry. The remainder of this section is organized as follows. First, the work of Kortum et. al. is reviewed. Next, the approach I use is briefly outlined. The benefits (and drawbacks) of each are outlined.

Measuring Costs with Plant Level Data

Kortum et. al. provide a methodology and the first estimates of what they call a hedonic cost function. The interested reader is referred to their paper for a very clear discussion of their methods. The basic idea is to use plant-level factor requirements (obtained from the LRD) combined with the product attributes of the cars produced to infer the cost of various product characteristics. Once these estimates are obtained, one can then back out how the cost of any particular product attribute changes over time (as well as how overall costs changed over time.)

The first step in Kortum et. al.'s methodology is to use plant level data on material inputs, labor, and capital as well as output mix to estimate factor input requirement equations. These equations map a relationship between the minimum amount of a factor needed to produce a given amount of output and the attributes of the product mix. Total factor usage for a generic factor M is given by:

$$M_p = \mu_m + J_p\delta_m + \sum_{j \in J(p)} M_{jp},$$

where p indexes plants, j indexes models, $J(p)$ is the set of models produced by plant p, μ_m is the fixed factor usage associated with plant upkeep, δ_m is the sum of any per model factor usage associated with the change over between models produced in the plant, and M_{jp} is the factor usage of model j. The estimating equations take the form of per unit factor requirements. A typical estimating equation is given by:

$$\frac{M_p}{Q_p} = \beta_0 + \frac{1}{Q_p}\mu + \frac{J_p}{Q_p}\delta + \sum_r x_{pr}\beta_r + \epsilon_p.$$

When taken to the data, year dummies are included as are firm-level fixed effects in some specifications. The β's give the factor usage associated with the various product attributes. To illustrate the idea, consider the following hypothetical example. Suppose in the labor demand equation that one saw that the coefficient on horsepower was steadily falling over time. That is, the labor requirements associated with the production of horsepower were decreasing. This would imply that one could increase horsepower and keep costs constant over time. Kortum et. al. estimate factor demands for labor, capital, and parts/material while the included product attributes are a dummy variable for trucks, a dummy variable for the presence of air conditioning, the miles per gallon (MPG), horsepower, and weight. Kortum et. al. summarize their results as follows:

> The parts and materials requirements of auto assembly, and these are about 85 % of the costs of automobile assembly, are largely determined by the characteristics of the vehicles assembled at the plant. There are no noticeable material costs of plant upkeep or model change over. On the other hand, the labor and capital requirements of labor and capital requirements are almost entirely determined by the fixed upkeep costs of the plant and, in the case of labor, possible ‘‘change over’’ costs of assembling different models at the same plant and a per unit cost. That is the labor and capital requirements bear little relationship to the characteristics of the vehicles assembled at the plants.’’

The authors find that most of the action is in the materials factor usage equation, and this is entirely consistent with plant visits. In the plant level data, well over 80 percent of the unit value of output is due to materials usage, while labor and capital together typically account for around 15-20 percent. (Assembly plant managers indicated that the vast majority of their overall expenses were well outside control of the assembly plant.) The data set spans 1972-1982 (and 1987 is also included in some specifications.) Patterns in the estimated coefficients on the year dummies might be interpreted as giving changes in input requirements conditional on the changing product mix. These patterns, though, are hard to find and change depending on the specification adopted. Also, these year dummies are usually imprecisely estimated. Instead, Kortum et. al. focus on the decreasing material inputs required to produce MPG.

The big advantage of this entire approach is that it does not rely on typically arbitrary assumptions regarding the nature of industry equilibrium to back out costs from a pricing first order condition. Rather, it relies directly on plant level factor usage and product mix information and product attribute information. No demand elasticities are ever required nor does one have to specify the nature of competition in the industry (as would be the case in the standard oligopolistic pricing first order condition.) This truly is a big advantage.

There are some drawbacks, too, and Kortum et. al. are entirely forthright about these. Three of these drawbacks are potentially especially important in the context of analyzing what has happened to productivity in the auto industry. First, much of the focus of the Machine That Changed the World and the Harbour Report is on comparisons across firms. For example, how productive is GM relative to Ford or Chrysler or Toyota? While the Kortum et. al. methodology can address this question for firms producing in the U.S., confidentiality concerns prohibit publishing those answers. That is, while one can estimate plant and firm-effects using the LRD, those effects are secret. This is always a trade-off when using proprietary data, but the Census rules are especially strict. Also, the LRD plant-level data only allow the researcher to address productivity concerns for firms in the LRD sample, and this excludes firms exporting to the U.S. Hence, European firms, which until this year did not assemble cars in the U.S., and many Japanese models not assembled in the U.S. will not enter the productivity comparisons when using the LRD sample.

The second drawback of the Kortum et. al. approach relates to the time series implications of their set-up. In order to keep their methods reasonably tractable, Kortum et. al. assume a fixed coefficients production technology. This precludes the possibility that labor, capital, and materials are substitutable for one another. Since the purpose of this paper is to examine how productivity varied over a couple decades, this assumption is problematic. Indeed, much of the message of The Machine That Changed the World could be interpreted as suggesting that materials and capital are quite substitutable for labor. The specification adopted by Kortum et. al. precludes this possibility.

The third drawback is perhaps the most serious since there is basically nothing one can do about it. The data requirements for the cost function approach are severe, for it is difficult to obtain access to the requisite data (one must become a sworn agent of the U.S. Census) and once access is obtained, one must do the work on-site in either Boston or Suitland, Maryland.

Measuring Costs Without Plant-Level Data

The alternative to using plant-level data is to rely only on product-level data (although household-level data could be used on the demand side of the model) to back out marginal costs. This is the approach labeled the New Empirical Industrial Organization approach. The basic idea is to use demand side estimates to nail down product level demand elasticities, and then combine these elasticities with assumptions on the nature of competition to back out marginal cost. This approach was employed by Bresnahan (19xx) and was adapted to a reasonably flexible discrete choice-differentiated products setting by Berry, Levinsohn, and Pakes (BLP) (1995.) The advantages of this approach are that it does not require the plant-level data used in the direct estimation of the cost function. By not requiring plant-level data, it is possible to make use of a long time series and to infer costs for all products sold (as opposed to produced) in the U.S. auto market. On the other hand, there is certainly something potentially unsettling about estimating costs in the absence of any direct data on costs. Somewhat arbitrarily imposed structure matters a lot.

The approach adopted in this study, not surprisingly, is that of BLP. The idea is to see whether the picture of productivity that emerges from this approach supports or is at odds with that of the other approaches.

The basic set-up is taken from Berry, Levinsohn, and Pakes (1995) (hereafter BLP). For purposes of brevity, an intuitive discussion of BLP is given here and the interested reader is referred to BLP for a (much!) more in-depth discussion. The model estimated has two parts--a utility-based consumer framework on the demand side and a cost-function-based model of a multi-product oligopolistic firm on the supply side. Each are discussed in turn.

Following a strategy developed by Pakes (1986), demand in this model is computed by aggregating over simulated heterogeneous consumers. Consumers' utility functions are assumed to have the same functional form, but the parameters of the function vary across the population. This is because consumer tastes vary throughout the population. The distribution of tastes is one of the primitives that is estimated. I assume that tastes for product attributes such as horsepower, weight, and size are normally distributed in the population. The estimation procedure estimates the mean and variance of these normal distributions. Price is treated a bit differently than other product attributes. I assume sensitivity to price is inversely proportional to income, and it is income that varies throughout the population. Rather than estimating the distribution of income as, say, the distribution of tastes for horsepower is estimated, the empirical distribution of income is used. There is also a random idiosyncratic component to utility. A simulated consumer then consists of a draw from each of the distributions of tastes and income as well as a draw from the distribution of idiosyncratic terms. This simulated consumer then chooses to either buy a car or spend nothing and instead buy the ‘‘outside good.’’ The utility of the outside good is normalized to zero, and its presence allows substitution out of the auto market. Conditional on this set of draws, one can then compute which product gives this simulated consumer the greatest utility. One can imagine simulating about 90 million consumers (the number of households in the U.S.), hence effectively simulating the demand for automobiles. One would keep track of the most preferred product of each of these consumers and aggregate up to compute market shares. Loosely speaking, the objective of the estimation procedure is to find the means and variances of the underlying distribution of tastes that come as close as possible to fitting the observed market shares.

The above description ignores many important aspects of the demand side of the model. These include econometric issues such as allowing for product characteristics that are unobserved by the econometrician but observed by the consumer, the probable correlation of these unobserved characteristics with price and the econometric endogeneity thus induced, and sampling techniques (in particular, importance sampling).

On the supply side of the model, Each product is assumed produced with constant returns to scale and a (log) marginal cost function is estimated for each product. Marginal cost is assumed to depend on attributes of the product as well as cost shifters such as wages and exchange rates (when applicable.) The firms are modeled as multi-product oligopolists who set prices in a Nash fashion (i.e. Bertrand competition). That is, firms set prices to maximize firm-level profits taking as given the prices of their competitors. To compute the prices that maximize profits, firms make use of information on demand elasticities. Price is composed of marginal cost plus the markup. Since the demand system is not a constant elasticity system, markups will depend on quantities demanded. (i.e. Demand elasticities vary along the demand schedule.) The demand and pricing sides of the model are simultaneous, because demand depends on prices and the prices set by the firms depend on quantity demanded. Put another way, firm's first order conditions for optimal prices depend on demand elasticities, and the underlying (indirect) utility function itself depends on the prices firms charge. The pricing and utility sides of the model are estimated simultaneously.

In this framework, a firm is more productive if, conditional on it's product mix, it's marginal costs are lower. Product mixes vary both over time for a given firm (a 1972 Chevrolet Caprice is much different than its 1994 edition) and across firms (a 1994 Ford Escort and a 1994 BMW 735 bear little resemblance to one another.) Hence it is crucial, in any cost and productivity comparison to control for changes in these product mixes. The BLP framework, complete with an accounting of observed and unobserved cost shifters, allows one to account for and condition on changes in products. Hence, if log marginal cost (mc) is given by

$$ln(mc_{jt}) = w_{jt}\gamma + \omega_{jt}$$

where j indexes products, t indexes year, w are observed cost shifters, ω unobserved cost shifters, and γ a vector of estimated parameters, then one can account for time trends in marginal costs by including trending terms. One can also account for firm-specific shifts in (log) marginal costs by including firm fixed effects. Finally, one can allow for firm-specific trends by interacting the firm-effects with the trends. This specification is given by:

$$ln(mc_{jt}) = w_{jt}\gamma + \beta_0 year_t + \sum_F \beta_F firm_F + \sum_F \delta_F(firm_F * year_t) + \omega_{jt}$$

where $firm_F$ is a firm-specific dummy variable. This specification allows costs to differ across firms and over time, albeit in a log linear fashion. A less restrictive specification would allow year dummies instead of a linear trend, but when interacted with firm dummies, there are too many parameters to precisely estimate given the data. (A discussion of alternative specifications is given in the next section.) An intermediate step between the linear trend and year dummy variables is to allow a non-log-linear trend, and this specification is the subject of future investigation.

The decision of how to enter the trends is dictated mostly by data constraints. There is just not enough data to precisely estimate everything we would like to know. A more conceptual issue is whether wages and, especially, exchange rates ought to enter the specification estimated. On the one hand, these variables are surely cost shifters and as such they belong in the marginal cost function. Previous work included these variables and they typically showed up as quite statistically significant. On the other hand, if the experiment one has in mind here is to evaluate whether, say, Japanese firms are more productive than U.S. firms, there is an argument for excluding these cost shifters. One interpretation of being more productive is being able to produce the same vehicle at lower cost, *not* at a lower cost conditional on wages and exchange rates. Since exchange rates are typically outside the control of any given firm, I opt to include them in the base case. By conditioning on exchange rates, I am essentially controlling for the cost advantage or disadvantage imposed by changes in exchange rates. As a sensitivity analysis, the model is re-estimated without wages or exchange rates.

To summarize, I estimate a demand system and cost system simultaneously. As in any oligopolistic equilibrium, demand elasticities feed into the firm's first order conditions, while the prices charged by firms feed into the demand system. Although, much attention is paid to estimating the demand system in a flexible way that does not impose artificial patterns of substitution, this is not the focus of this study. Rather, this study employs these methods, which rely only on product-level data, to investigate changes in productivity as they are reflected in changes in marginal costs. Some changes in productivity, though, may be reflected in changes in fixed costs, and the methods used here simply will not capture those. This is a drawback, but fixed costs in the automotive industry are hard to measure and harder still to attribute to specific products. The approach adopted here is to measure what we can, and see what the comparison yields.

4. Results Using Product-Level Data

The data set is comprised on product-level information on almost all cars sold in the U.S. from 1971 to 1994. There are 2809 individual models sold over this period. This data set is augmented with annual data on exchange rates (obtained from the IMF), annual data on manufacturing (i.e.*not* automobile-specific) wages (obtained from the OECD), and annual data on the empirical distribution of income (obtained from the CPS.) The first two of these are used in the cost function, while the last is used in the demand-side estimation. (Again, details are found in Berry, Levinsohn, and Pakes, 1995.) The product-level data includes detailed information on product attributes of the base model as well as list price of the base model. Not all product attributes for which data are available are actually used in the cost function, since many are collinear with one another and there are just too many to obtain precise estimates of the parameters of interest. The data set is not proprietary and is available upon request by e-mail.

On the demand side, I estimate means and standard deviations of the (assumed normal) distribution of tastes for the following product characteristics: a constant (the standard deviation of which captures variability in substitution to the outside good), horsepower divided by weight (a good measure of acceleration), external size of the car (width times length), whether air conditioning is standard equipment (a decent measure of luxury) and whether the car is made in the U.S. The product price and consumer income also enter consumer utility and the specification allows for a distribution of ‘‘tastes’’ for how price impacts utility. Experiments with other variables in the utility specification suggested that the resulting elasticities were not that sensitive to the choice of variables included in the utility function. It is these elasticities that enter the pricing equation and hence impact the estimate of marginal cost (which is what we care about). Because adding more variables considerably increases computational time, I adopted the more parsimonious demand side specification and concentrated efforts on the cost side instead.

On the cost side, the base case specification essentially projects log marginal cost on a constant, seven product attributes, lagged (log) exchange rate, the (log) of the wage, a dummy variable for cars produced in Europe, a dummy variable for cars produced in Japan, separate dummy variables for each of the Big Three U.S. firms, a trend, and a set of interaction terms that allows European and the three U.S. firms to each have their own trends. The eight product attributes are the logs of horsepower (HP), weight, length, width, and miles per gallon (MPG), and dummy variables for air conditioning (AIR) and automatic transmission (AT) as standard equipment.

Hence, (omitting subscripts) I estimate :

$$
\begin{aligned}
ln(MC) =& \gamma_0 + \gamma_1 ln(HP) + \gamma_2 ln(WEIGHT) + \gamma_3 ln(LENGTH) + \gamma_4 ln(WIDTH) + \gamma_5 AIR \\
&+\gamma_6 ln(MPG) + \gamma_7 AT + \gamma_8 LAG(ln(E-RATE)) + \gamma_9 ln(WAGE) \\
&+\beta_1 EUROPE + \beta_2 FORD + \beta_3 GM + \beta_4 CHRYSLER \\
&+\delta_1 TREND + \delta_2 (TREND * EUROPE) + \delta_3 (TREND * FORD) \\
&+\delta_4 (TREND * GM) + \delta_5 (TREND * CHRYSLER) + \omega.
\end{aligned}
$$

Japan is the the omitted country, so the trend term captures trends in Japanese costs, while the interacted trend terms capture how these firms cost trends differ from those of the Japanese firms. Similarly, the dummy variables for Europe and the Big Three capture the difference in log marginal costs relative to the Japanese firms.

The estimated coefficients from this specification can inform the discussion on productivity. The first ‘‘step’’ is to control for differences in the cars themselves. As noted in the introduction, cars with different vectors of attributes would be expected to cost different amounts to produce. In order to be able to make reasonable comparisons across different products, we need to control for these differences, and this is the role of the included attributes. Ideally, one would also like to allow costs to differ for each firm. That is, conditional on the attributes of the car, some firms produce their product at higher or lower costs. There are too many firms in the sample (26) to precisely measure each of these effects. Instead, I measure five such effects-- one for Japanese cars, one for European cars, and one for each of the Big Three domestic firms. These are the coefficients, respectively, on trend, Europe, Ford, GM, and Chrysler. Finally, costs for a given firm or group of firms may well change over time. That is, it may simply cost less to produce a car with given specifications over time as productivity increases. These changes in costs may not be the same for every firm or group of firms, and the coefficients on the trend terms capture these movements.

Table 1 gives the results of the estimated specification. Details of estimation are found in BLP. I focus discussion on the estimated parameters of the marginal cost function. All but one of the vehicle attributes enters marginal cost in a statistically significant manner. Not surprisingly, controlling for vehicle attributes matters. One could allow the implicit cost of these attributes to themselves vary over time. This, for example, is a prime focus of the Kortum, Berry, and Pakes study. The focus of this study, though, is how firms' costs have changed over time and I am less occupied with whether it is progress in the ''production'' of horsepower or MPG or weight underlying the productivity differences. (Due to the proprietary nature of the LRD, Kortum et. al. simply could not report results along the lines in Table 1.)

The country and firm dummy variables are generally statistically significant. The sole exceptions are the coefficients on Europe and Europe interacted with a trend. The probable reason for the fuzzy estimates of the European coefficients is that the composition of European cars changes drastically over the sample-- from Volkswagen Beetles to BMW's and Mercedes. Combining these groups and asking whether European cars are more or less expensive (even conditional on the included product attributes) is probably wrong. The country and firm dummy variables may be interpreted as giving the log marginal cost differences between firms in 1971 (the start of the sample.) The estimates imply that in 1971, the U.S. firms had lower costs (conditioning on exchange rates as well as vehicle attributes). Since the dependent variable is log marginal costs, one must exponentiate the coefficients in order to interpret their magnitudes (which would then be in 1000's of 1983 dollars.) The trend term on Japanese cars ($Trend$) is negative implying that Japanese costs are falling over time. The positive coefficients on the other interacted trend terms are all positive and smaller than the absolute value of the coefficient on the Japanese trend. This implies that marginal costs for the non-Japanese firms are also falling, but not as quickly as for the Japanese firms.

A perhaps easier way to interpret the information implicit in the coefficients on the firm/country dummies and the trends is to plot the pattern implied by those coefficients. The estimated β's in the estimating equation show how costs change with vehicle attributes. By conditioning on these attributes, I control for the many of the differences in automobiles and the estimated δ's indicate how marginal costs differ across countries (and U.S. firms) and over time. In order to make comparisons easier, I normalize costs relative to the estimated Japanese productivity. This is the visual analog to the fact that Japan is the excluded country in the estimated regression. Such a plot is given in Figure 1. Figure 1 takes as a baseline the Japanese pattern of productivity increases and graphs productivity (as proxied by marginal costs) relative to the Japanese. The horizontal axis in Figure 1 measures time and each of the points in the graph represents an annual observation. The vertical axis measures the difference in log marginal costs. To put the numbers on that axis in perspective, -1 corresponds to the case in which a firm has marginal costs about $2700 less than a Japanese firm conditional on vehicle attributes, wages, and exchange rates. The pattern is broadly one of a trend towards convergence, but one in which Japanese costs are the highest. Ford's costs are the lowest, while GM's and Chrysler's costs are between Ford's and those of the European firms. By the end of the sample period, European, GM's, and Chrysler's costs are all quite quite similar. Ford's costs remain about $1400 lower than those firms, while Japanese costs remain higher than those of the converged firms.

Are the implications of Figure 1 reasonable? This pattern is only partially consistent with the picture of productivity painted by the other. The notion that productivity differences have narrowed is indeed consistent with received wisdom on this topic. The idea that Japanese costs are *higher* than others' is less credible. One reason for this result may be the role of the Voluntary Export Restraints (VERs) that were in effect throughout most of the 1980's. Within the context of the structural model used in this paper, the VERs introduce another wedge between price and marginal cost, much like that of a tax. By ignoring this wedge, and instead attributing price minus the markup to marginal cost, I may be over-estimating Japanese marginal costs.

An advantage of this structural econometric approach is that it is possible to at least try to figure out what might be explaining the results. Using the methods described in Berry, Levinsohn, and Pakes (1997), I re-estimated the model while now accounting for the VER. The parameter estimates remain quite statistically significant and the pattern of implied marginal costs is given in Figure 2. Japanese costs are given by the horizontal line at zero (since they are the baseline in this exercise.) The pattern is clearly one of convergence. European costs are higher than Japanese costs, although this mostly disappears by the end of the sample. In 1971, Japanese costs are substantially higher than domestic costs, but this gap is mostly eliminated by 1994. Put another way, relative to U.S. costs, Japanese firms have made large productivity gains over the last two decades. By 1994, all firms have reasonably similar marginal costs. The results displayed in Figure 2 are much more in line with non-econometric discussions of productivity than are the results in Figure 1. It remains the case in Figure 2, though, that Japanese costs are estimated to be *higher* than domestic costs, and this is not consistent with the non-econometric discussions.

Most of the non-econometric discussions of productivity (and implicitly costs) ignores the role of exchange rates. Over the sample period, exchange rates have undergone massive shifts. It is worthwhile to untangle the role exchange rates play in Figures 1 and 2. There are multiple ways to do this. I opt to re-estimate the cost system with the same variables but now excluding the exchange rate variable. I do this both accounting for and ignoring the VER. The results when exchange rates are excluded but the VER is ignored are given in Figure 3. To reiterate, the only difference between the models generating Figures 1 and 3 is that in Figure 3, exchange rates are not an included regressor. (The results when exchange rates are excluded and the VER is modeled are very similar to those in Figure 3.) The baseline, given by the solid horizontal line, is still Japanese marginal costs. The relative pattern of the domestic firms is of course unchanged. (Exchange rates play no role in this ranking, although this finding is not guaranteed, since when the system was re-estimated, exchange rates might have co-varied with included cost-shifters.) When exchange rates are not an included cost-shifter, the broad pattern of converging marginal costs is quite strong. European costs are the highest, followed by Japanese costs, and then Chrysler, GM, and Ford. By the end of the sample period, the large initial costs differentials have substantially narrowed.

There are are few things that do not explain these results. I have experimented with other vectors of vehicle attributes, and find the results robust to a variety of attributes. I have also experimented with how transplant production (Toyotas produced in the U.S.) are treated and find that the results are robust to whether they are considered Japanese or American cars. (The relevant wages and exchange rates differ.) The pattern is also reasonably robust to the exclusion of the wage variable in the cost function. I have also experimented with more flexible functional forms, but find that the inclusion of too many terms gives very imprecise estimates of most of what I care about. For example, I have estimated a translog cost function, and almost nothing was precisely estimated.

5. (Preliminary) Conclusions

This paper has reviewed productivity in the auto industry from three quite different approaches and then suggested and implemented a fourth approach. The approach I use is the only econometric approach to measuring firm productivity that permits reporting firm-level productivity results, and it does so by using only widely available product-level data. Unlike the other econometric approach which requires firm-level data, though, my approach requires that more structure be imposed on the nature of equilibrium in the auto industry.

When my econometric results are compared with the non-econometric studies of productivity, some interesting patterns emerge. The main messages regarding productivity of ‘‘The Machine that Changed the World’’ are that the Japanese are more productive but that the gap will narrow if U.S. firms adopt the methods of lean production. These methods have indeed been widely adopted and the gap, according to my estimates has indeed narrowed. My results show productivity converging for almost all manufacturers. The possibly puzzling feature of my results is that starting in 1971, the Japanese are estimated to have higher costs than other manufacturers, so that while every domestic firm is estimated to be getting more productive, it is the Japanese that became ‘‘more productive’’ quicker as industry productivities converged.

There are other possible explanations for why my econometric estimates might differ from the non-econometric studies. First, my estimates are based on comparisons of marginal cost, while many of the differences highlighted by ‘‘The Machine that Changed the World’’ are differences in fixed costs. In this case, there is little reason to expect the econometric and descriptive productivity results to mesh. Second, the descriptive studies often do not take into account exchange rates. As the analyses above indicate, over the last 25 years, exchange rates have played an important role and whether one conditions cost estimates on the exchange rate matters.

The Harbour report compared the domestic firms in a given year using detailed engineering and manufacturing data. The report found that there was substantial inter-firm variation in productivities, but that Ford and Chrysler tended to be more productive than GM. Using 24 years of data instead of one, but much less detailed data, a similar pattern emerges from my estimates. Ford is consistently the most productive, while Chrysler and GM follow. The picture one gets from industry documents is that Chrysler has made the most dramatic gains in productivity over the last few years, and it is quite possible that the linear productivity trend I have imposed is inappropriate for Chrysler. This is the subject of ongoing research.

Because the econometric approach of Kortum et. al. does not include foreign manufacturers and because it cannot report firm-level productivity measures, I have no idea whether those results are broadly consistent with mine. Such a comparison would be useful, since if methods that impose lots of structure but require less data (as in this paper) give similar results to those that require less structure but more data, researchers might be more confident in using publicly available product-level data to estimate productivity. Similarly, the converse holds if the results are massively different.

References

Berry, S., Kortum, S., and Pakes, A. (1995a). Hedonic cost functions for automobiles and the fuel efficiency of new cars. mimeo, Yale University.

Berry, S., Levinsohn, J., and Pakes, A. (1995b). Automobile prices in market equilibrium. *Econometrica*, *63*(4), 841--890.

Berry, S., Levinsohn, J., and Pakes, A. (1997). Voluntary exports on automobiles: Evaluating a strategic trade policy. mimeo, University of Michigan and Yale University.

Bresnahan, T. (1981). Departures from marginal cost pricing in the American automobile industry. *Journal of Econometrics, 17,* 201--227.

Griliches, Z. (1961). Hedonic price indexes for automobiles: An econometric analysis of quality change. In *The Price Statistics of the Federal Government*. National Bureau of Economic Research, New York.

Harbour, J. E. (1995). *The Harbour Report: 1995*. Harbour and Associates, Troy, Michigan.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, *82,* 34--55.

Womack, J., Jones, D., and Roos, D. (1990). *The Machine That Changed The World*. Harper Collins, New York.

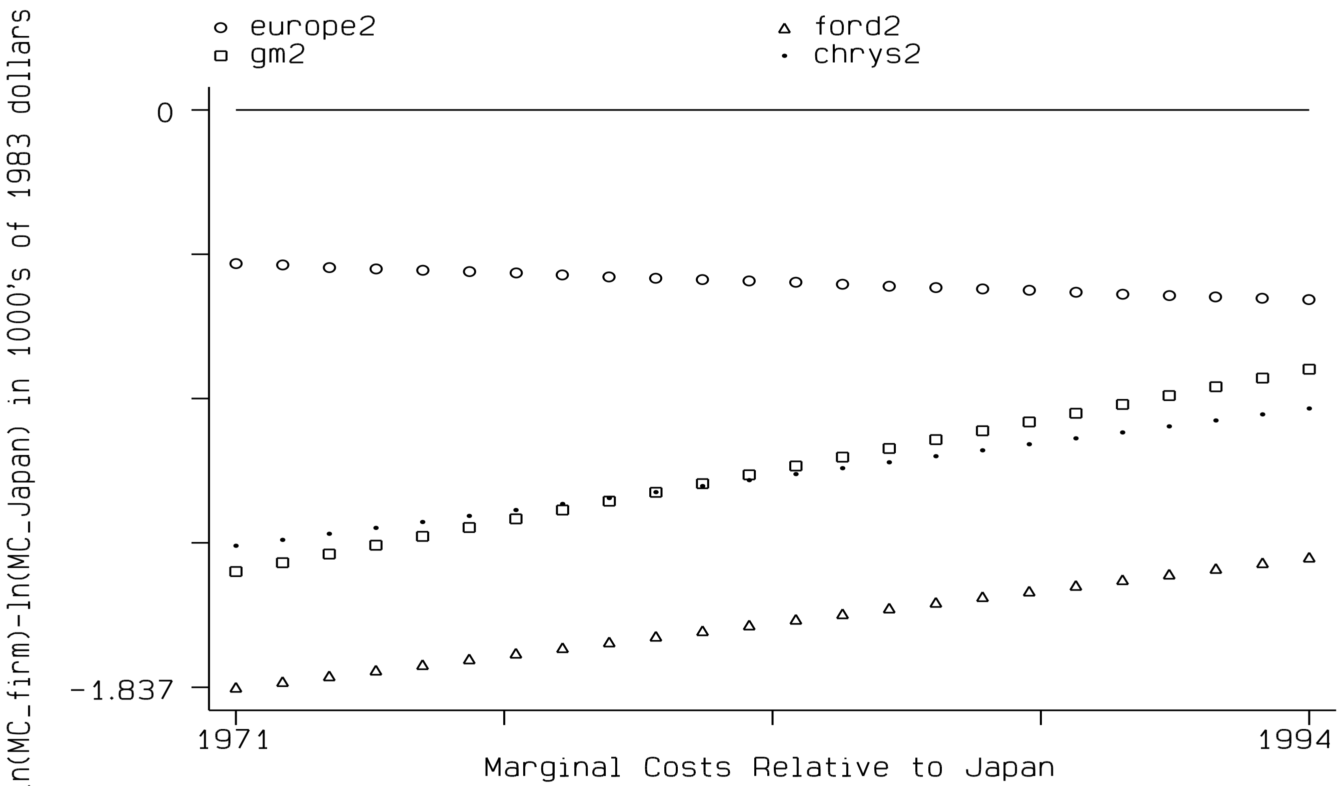
europe2
gm2
ford2
chrys2
0
-1.837
1971
1994
ln(MC_firm)-ln(MC_Japan) in 1000's of 1983 dollars
Marginal Costs Relative to Japan
Figure 1

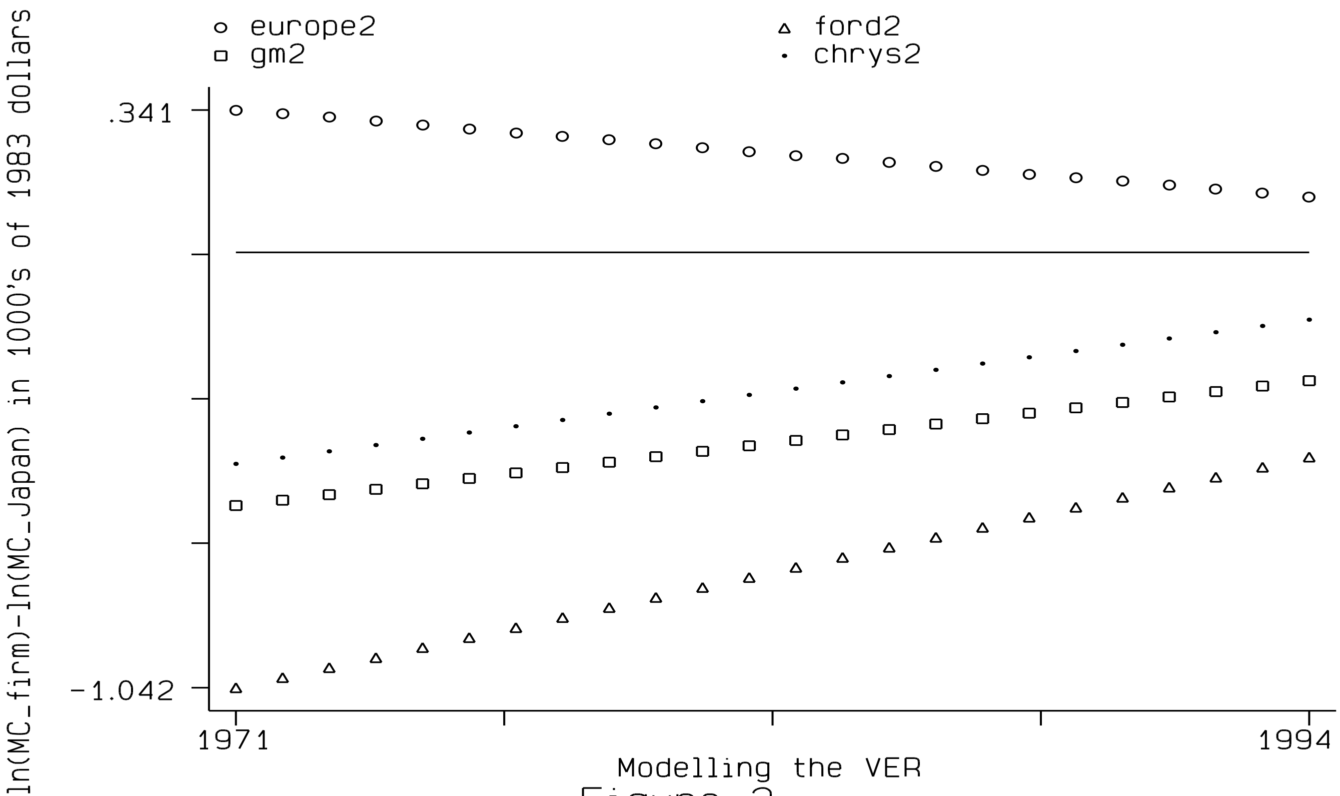
europe2
gm2
ford2
chrys2
.341
-1.042
1971
1994
ln(MC_firm)-ln(MC_Japan) in 1000's of 1983 dollars
Modelling the VER

Figure 2

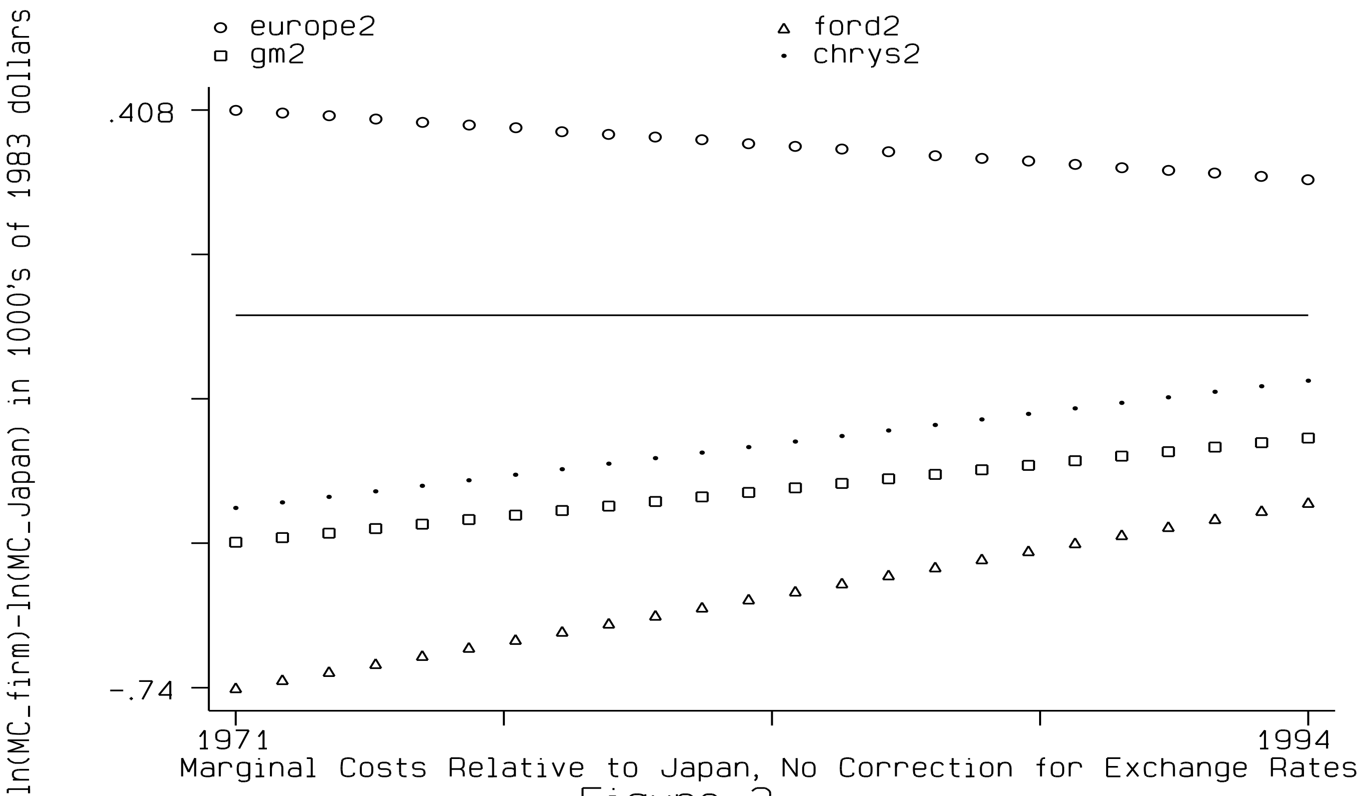
europe2
ford2
gm2
chrys2
.408
-.74
1971
1994
ln(MC_firm)-ln(MC_Japan) in 1000's of 1983 dollars
Marginal Costs Relative to Japan, No Correction for Exchange Rates

Figure 3

TABLE 1

Estimated Parameters of the Marginal Cost and Demand Equations:

Base Case Specification

1971-1994 Data

	Variable	Parameter Estimate	Standard Error
Cost Side Parameters			
	Constant	-1.629	0.700
	ln(HP)	0.803	0.077
	ln(Weight)	0.815	0.138
	ln(Length)	-0.705	0.205
	ln(Width)	-0.078	0.108
	Air	0.291	0.029
	ln(MPG)	0.122	0.054
	AT	0.095	0.020
	Europe	-0.485	0.459
	Ford	-1.855	0.553
	GM	-1.497	0.562
	Chrysler	-1.406	0.562
	Trend	-0.031	0.006
	E-Trend	-0.005	0.004
	G-Trend	0.028	0.006
	F-Trend	0.018	0.006
	C-Ttrend	0.019	0.006
	Lag(ln(ER))	0.201	0.095
	ln(wage)	0.891	0.091
Demand Side Parameters			
Means (β's)	Constant	-7.026	0.491
	HP/Weight	0.384	0.640
	Size	3.354	0.224
	Air	0.135	0.153
	Foreign	-1.040	0.289
Std. Deviations (σ_β's)	Constant	1.655	0.896
	HP/Weight	-2.359	0.903
	Size	-0.494	0.696
	Air	1.273	0.268
	Foreign	2.171	0.427
Term on Price (α)	$\ln(y - p)$	26.183	2.592