

The Marshall Lectures

Cambridge University

January, 2004.

Models of Markets for Applied Work

by Ariel Pakes

(Harvard University).

Lecture 1: Static Analysis.

Lecture 2: Dynamics.

Outline.

The goal is to build up a set of tools that enable us to analyze market outcomes in oligopolistic situations. I have worked primarily on symmetric information models, and so will focus these lectures on that case¹. We begin with the “textbook” situation in which one can analyze price and quantity choices in a static framework².

Static Analysis.

Recall that static analysis conditions on

- the goods marketed (or their characteristics) and their cost functions,
- consumer’s preferences over those goods (or over characteristics tuples),
- “institutional” features like; the type of equilibrium, structure of ownership, regulatory rules

and then analyzes how prices, quantities, and the distribution of profits and consumer surplus, are determined (or change in response to an environmental or policy change).

So to do the static analysis we need the following primitives;

1. the demand system
2. the cost system
3. an equilibrium assumption

I begin with a short introduction to recent work on each of these.

¹Many of the primitives needed for this case will also be needed for asymmetric information models, so it also makes sense to start here.

²This is only a reasonable assumption when the price (or quantity) choice does not have an independent effect on: (i) future costs, (ii) future demand, or (iii) future equilibrium choices. So it clearly rules out many important empirical applications; in particular those in which either learning by doing, experience or durable aspects of the goods marketed, or collusion, are likely to play a major role. However it is a base case which, once understood, points directly to the generalizations needed for the more complex alternatives.

Demand Systems.

There has a lot of recent work on demand systems. The move has been

- away from “representative agent” models to models with heterogeneous agents, a move greatly facilitated by modern computers and the advent of simulation estimators (Pakes, 1986)
- and from models set in “product space” to those set in characteristic space (this dates back to the theoretical work of Lancaster, 1971, the early econometric work of McFadden, 1974, and more immediately the work of Berry, Levinsohn, and Pakes, 1995, henceforth BLP).

Most of the first lecture will be on the benefits of these advances, so I come back to them below.

Production and Cost Functions.

The recent work on production and cost functions has largely been motivated by access to plant (sometimes firm) level data sets (usually these data sets are panels). This in turn has focused attention on particular technical and substantive issues.

Technically there has been renewed interest in the impacts of simultaneity (endogeneity of inputs) and selection (endogeneity of attrition) on parameter estimates (see Olley and Pakes, 1996, and Levinsohn and Petrin, 2003). This corresponds to the fact that there are both large serially correlated differences in measured “productivity” among plants, and large sample attrition and addition rates in these panels (see Davis and Haltwinger, 1992, for the original work on US manufacturing data). Corrections for both simultaneity and selection based on statistical assumptions on the properties of the error (e.g., use of fixed effect and related estimators, and use of the propensity score) seem less effective than corrections based on economic models of input and exit choices.

Substantively there has been a renewal of interest in measuring productivity and how policy changes affect it (e.g., deregulation, large changes in tariffs, privatization and other liberalizations in the environment). Importantly the availability of the micro data has enabled us to distinguish between the growth in aggregate productivity that

- emanates from the growth of the productive efficiency of individual establishments, and

- the growth that emanates from more efficient output allocations among establishments of differing productivities,

(see Olley and Pakes, 1996).

There is much more to be done here, particularly in terms of expanding the model to allow us to investigate other issues. Of particular interest is the analysis of the speed and timing of input adjustments, their impact on productivity, and their relationship to unionization, regulation, legal systems, etc.. Other topics include production functions that depend on the characteristics of products, and the analysis of the sunk costs of product development. Perhaps the major reason that more resources are not devoted to these issues is the relative paucity of the requisite data; relative, at least, to the data available to analyze demand systems.

Equilibrium Assumptions.

It is clear that the least progress has been made here. Empirical work relies heavily on a Nash in prices (or in quantities) assumption, and it seems to do a fairly good job accounting for the cross-sectional variance in prices (or output), at least in markets where there are many diffuse purchasers (particularly in characteristic based models, a surprisingly good job). On the other hand, at least in my experience, these simple models do not do nearly as well in accounting for changes in prices over time in those industries, and they are also often problematic in industries where purchasers are large.

More generally we have a lot of difficulty analyzing pricing and quantity-setting decisions in goods where there are a small number of both buyers and sellers (e.g., the prices hospitals charge HMO's or producers charge retail chains...). Part of the problem is that, at least given the factors that the empirical or applied person can condition on, often many different equilibria are possible. There apparently there is not much agreement on the reasonableness of alternative refinements, nor has there been much empirical work on sorting out which equilibria are more likely (for historical, or other, reasons). Multiple equilibria become even more likely when we move to dynamic models, or when we allow for changes in the "institutional" structure the static analysis conditions on (e.g. changes in ownership patterns through mergers), so a word on their impact on empirical work is in order.

The multiple equilibria issue manifests itself in at least two ways in empirical work. First it poses an obstacle to traditional estimation algorithms.

Since we cannot associate a unique outcome with a given set of parameters and vector of observable and unobservable states, we cannot form the probability of the observed events conditional on those parameters; i.e. we cannot form likelihoods. There are a number of ways of getting around this problem, and we will come back to review them in the second lecture.

There is, however, a problem that arises due to multiple equilibria that has been more difficult to overcome. Part of the reason for building a consistent theoretical framework to underlie estimation is to provide a framework which can be used to explore counterfactuals; to make predictions as to what is likely to happen given different policies or an environmental changes. Even if we were to make the assumptions required to identify the equilibria that is currently being played, we still would not know which equilibria would result were one of the counterfactuals to become reality. That is, without further developments, we cannot make the desired predictions.

There has been disappointingly little work on this problem. One possibility that has not been explored is to exploit the recent work on simple “learning” mechanisms to choose among likely equilibria (see the book by Fudenberg and Levine, 1999, and, for models closer to those discussed here, Pakes and McGuire, 2001). Thus, for example, though we may not know how the industry evolved to its current state, we might try to use one or more of the “reinforcement” learning rules to examine how the industry is likely to change in response to a change in the institutional environment. Modern computers make these types of investigations relatively easy. Of course we would have much more confidence in the result if we had investigated whether the learning rules had good predictive ability in real world situations of institutional change. That is empirical work on just how industries respond to institutional changes would be very useful.

Dynamics.

What the static analysis does is deliver prices of all products (and therefore all quantities) as functions of the state variables of all firms. This, plus our primitives, enables us to solve for each firms profits as a function of all firms states, as well as the distribution of consumer surplus. The goal of the dynamic analysis is to analyze the changes in the firms’ state variables that the static analysis conditions on.

Thus if we put on a tax, or allow a merger, we will change prices, profits and consumer surplus conditional on the goods marketed, their cost func-

tions, and preferences. The static analysis is designed to analyze the immediate impact of a policy or environmental change on these variables. The change in policy will also, however, change investment entry and exit incentives (where investment is included broadly enough to include advertising and R&D). That is, in the longer run the environmental changes will also induce changes in the the quality and variety of goods marketed as well as their ownership structure and cost functions. It is this response of “market structure”, and through market structure the future profits and consumer surplus, that the dynamics is designed to analyze.

The primitives needed for the dynamic analysis are;

- Profit and consumer surplus as a function of the state variables of the problem. [To be obtained from the static analysis.]
- The dynamic parameters which determine the cost and impacts of investment, and the discount rate. The investment variables typically include the sunk costs of entry and exit, and a family of distribution functions which specifies³

$$p(\text{state variables tomorrow} \text{ — current states, controls}).$$

- Equilibrium assumption for the dynamic controls. We begin with limiting the strategy space to the same “payoff relevant” random variables we used for the the static control, and looking for a Nash equilibrium in the “dynamic” controls [these are typically exit and investment decisions for incumbents, and entry and investment decisions for the potential entrants].

To actually analyze dynamic responses, then we need, in addition to the output of the static problem:

- An estimation algorithm capable of estimating the additional static parameters. The difficulty here is in estimating sunk costs of entry and exit, as they typically do not have observable counterparts and their relationship to observables depends on the nature of the whole equilibrium process.

³The fact that we write this as a probability distribution, rather than a deterministic equation of motion (like $k_t = \delta k_{t-1} + i_t$ with probability one), reflects the fact that we will often be dealing with processes in which the outcome of the investment is uncertain (research, advertising, and exploration processes).

- A computational algorithm which solves for Markov perfect investment, entry, and exit policies. With these policies it is easy to simulate the evolution of market structure.

These are the two topics I come back to in the second lecture.

Advances in the Analysis of Demand.

Demand systems are the starting point for most applied analysis in I.O. Essentially the nature of demand is the primary determinant of incentives in the models we work with. Thus the demand system is the only indispensable ingredient in the analysis of how prices are set in Nash equilibrium, and therefore of how prices might change after a merger, a tariff, entry of a new good,.... Of course prices determine markups, and markups make up the incentives for entry, exit, and investment, and so on.

Not long ago graduate lectures on demand systems were largely based on representative agent models in “product” space (i.e. the agent’s utility was defined on the product per se rather than on the characteristics of the product). There were a number of problems with this form of analysis that made it difficult to apply that framework to the IO problems of interest to us.

Simulation and Aggregation.

First almost all the applications used market level data: they would regress quantity purchased on (average) income and prices. There were theoretical papers which investigated the properties of market level demand systems obtained by explicitly aggregating up from micro models of consumer choices (including a seminal paper by Houthakker, 1955). However we could not use their results to structure estimation on market level data without imposing unrealistic *a priori* assumptions on the distribution of income and “preferences” (or its determinants like size, age, location, etc.).

The introduction of simulation estimators has enabled us to aggregate up from the *observed* distribution of consumer characteristics and any functional form that we might think relevant. That is we allow different consumers to have different income, age, family size, or location of residence. We then formulate a demand system which is conditional on the consumer’s characteristics and a vector of parameters which determines the relationship

between those characteristics and preferences over products (or over product characteristics). To estimate those parameters from market level data we simply

- draw vectors of consumer characteristics from the distribution of those characteristics in the market of interest (in the U.S., say from the March CPS),
- determine the choice that each of the households drawn would make for a given value of the parameter vector,
- aggregate those choices into a prediction for aggregate demand conditional on the parameter vector, and
- employ a search routine that finds the value of that parameter vector which makes these aggregate quantities as close as possible to the observed market level demands.

The ability to obtain aggregate demand from a distribution of household preferences both increases (i) the depth of our understanding of how “aggregate preferences” are formed and are likely to differ across markets, and improves (ii) the precision of our parameter estimates. Not only does this enable us to be more detailed and precise in our analysis of industrial organization issues, it also enables us to use them to analyze a host of distributional issues of distinct interest to related fields (examples include tax incidence and voting patterns).

For example we all believe (and virtually all empirical work indicates) that the impact of price depends on income. Our micro model will therefore imply that the price elasticity of a given good depends on the density of the income distribution among the income/demographic groups attracted to that good. So if the income distribution differed across regional markets, and we used an aggregate framework to analyze demand, we would require different price coefficients for each market.

Table I provides some data on the distribution of the income distribution across U.S. counties (there are about three thousand counties in the U.S.). It is clear that the income distribution differs markedly across these “markets”. The standard deviation of the fraction of households in our nine income groups varies between twenty and a hundred percent of their means (with an average across these groups of fifty three percent of the mean and larger

coefficients of variation among the higher income groups which consume a disproportionate share of the goods markets).⁴

Table I: Cross County Differences in Household Income*

Income Group (thousands)	Fraction of U.S. Population in Income Group	<u>Distribution of Fraction</u> <u>Over Counties</u>	
		Mean	Std. Dev.
0-20	0.226	0.289	0.104
20-35	0.194	0.225	0.035
35-50	0.164	0.174	0.028
50-75	0.193	0.175	0.045
75-100	0.101	0.072	0.033
100-125	0.052	0.030	0.020
125-150	0.025	0.013	0.011
150-200	0.022	0.010	0.010
200 +	0.024	0.012	0.010

* From Pakes (forthcoming, *RIO*) “Common Sense and Simplicity in Empirical Industrial Organization”.

It is pretty clear, then, that if we rely on the aggregate demand framework we are likely to require different price coefficients in different markets. If we based our estimates on an underlying micro model, on the other hand, we could get price effects that differ in a sensible way across markets without needing to resort to parameterizing each market separately. Moreover, unlike the aggregate framework, the micro model would allow us to make sensible predictions for price elasticities in locations where the good has not yet been marketed. Again getting sensible price effects is a prerequisite for getting sensible *markups* (since markups are closely tied to the inverse of the price elasticity of demand). Moreover getting sensible predictions for markups

⁴Interestingly counties with more households tend to have a larger fraction of their populations in the highest income group, and a smaller fraction in the lowest.

and demand are the prerequisites for getting sensible incentives for *product development*, and so on.

Of course one could stay with the aggregate model and add more detailed features of the income distribution as “right hand side variables”. However were we to go this route we would quickly move to a model with too many parameters to estimate, particularly once we begin interacting the percentiles of the income distribution with demographic and location characteristics of the population. Indeed there is a very real sense in which the major advantage of moving to the micro model is that it enables us to use economics to provide a sensible and empirically useful way of constraining the impacts of the joint distribution of consumer attributes on demand. Of course another advantage is that now we can use one basic framework to analyze demand systems estimated from both micro and aggregate data, allowing us to utilize the different strengths of the two data sets in one unified framework.

The Too Many Parameters and New Goods Problems.

Having “solved” the aggregation problem, there were at least two remaining problems with the analysis of demand systems.

1. Many of the markets we wanted to analyze contained a large number of goods that are substitutes for one another. As a result when we tried to estimate demand systems in product space we quickly ran into the “too many parameters problem”.
 - J goods implies on the order of J^2 parameters (even from linear systems where each consumer has the same coefficients), and often J was one hundred or more.
 - Gorman’s polar forms for multi-level budgeting were an ingenious attempt to mitigate this problem, but they required assumptions which were often unrealistic for the problem at hand, and the reduction in parameters was not sharp enough to enable the kind of flexibility we needed for I.O. Indeed typically the grouping procedures used empirically paid little attention to Gorman’s conditions and more to the policy issue of interest with correspondingly disappointing results.
2. Demand systems in product space do not enable the researcher to analyze demand for new goods prior to product introduction.

Characteristic Space.

A “solution” to these problems which is suitable for many (but not all) markets is to switch from product to characteristic space. In characteristics space models

- Products are bundles of characteristics.
- Preferences are defined on those characteristics.
- Each consumer chooses a bundle that maximizes its utility. Consumers have different relative preferences (usually just marginal preferences) for different characteristics, and hence make different choices.
- Simulation used to obtain aggregate demand.

In these models the number of parameters required to determine aggregate demand is *independent* of the number of products per se; all we require is the joint distribution of preferences over the characteristics. Thus if there were five important characteristics, and preferences over them distributed joint normally, twenty parameters would determine demand for all products (no matter their number). Further, once we estimate those parameters we can predict outcomes after adding a new product by simply recomputing equilibrium after introducing the new product’s characteristics to the choice set.

The theoretical and econometric work that laid the ground for characteristic based models dates back at least to Lancaster (1971) and McFadden (1974)⁵ Applications of the Lancaster/McFadden framework however, increased significantly after BLP (1995) showed how to circumvent two problems that had vexed the early generation of characteristic based models.

The problems were that

1. the early generation of models used functional forms which restricted cross and own price elasticities in ways which brought into question the usefulness of the whole exercise

⁵Actually characteristics based models have a much longer history in I.O. dating back at least to Hotelling’s classic article, but the I.O. work on characteristic based models focused more on their implications for product placement rather than on estimating demand systems per se.

2. the early generation of models did not allow for unobserved product characteristics

The second problem is particularly important when studying demand for consumer goods. Typically these goods are differentiated in many ways. As a result even if we measured all the relevant characteristics we could not expect to obtain precise estimates of their impacts. What BLP suggests is to put in the “important” differentiating characteristics *and* an unobservable, say ξ , which picks up the aggregate effect of the multitude of characteristics that are being omitted. Of course, to the extent that producers know ξ when they set prices (and recall ξ represents the effect of characteristics that are known to consumers), goods that have high values for ξ will be priced higher in any reasonable notion of equilibrium. This produces an analogue to the standard simultaneous equation problem in estimating demand systems in the older demand literature, and BLP shows how one can use instruments to overcome this “simultaneity problem” in the discrete choice framework.

There were two early characteristic based models used in empirical work; the logit model used intensively by McFadden himself, and the vertical model which Shaked and Sutton (1982) had used theoretically and Bresnahan (1981) first applied empirically. For simplicity I will focus on the logit model (though analogous problems arise in the vertical model). The simple logit model has consumer utility determined as

$$U_{i,j} = x_j\beta + \epsilon_{i,j}$$

where the x_j are the characteristics of product j (including the unobserved characteristic and price) and the $\{\epsilon_{i,j}\}$ are independent (across both j for a given i and across i for a given j) identically distributed random variables (in the pure logit, they have a double exponential distribution). Thus $x_j\beta$ is the mean utility of product j and $\epsilon_{i,j}$ is the individual specific deviation from that mean.

As noted by McFadden himself this specification suffers from a rather extreme form of the “IIA” problem. The distribution of a consumer preferences over products *other than* the product it bought, does not depend on the product it bought. One can show that this implies that

- since two agents who buy different products are equally likely to switch to a third product should the price of their product rise, two goods with the same shares have the same cross price elasticities with any other

good (cross price elasticities are $s_j s_k$, where s_j is the share of good j).

- since there is no systematic difference in the price sensitivities of consumers attracted to the different goods, own price derivatives ($\partial s / \partial p$) = $-s(1 - s)$ only depends on shares. This implies that two goods with same share must have the same markup in a single product firm “Nash in prices” equilibrium.

No data will ever change these implications of the two models. If your estimates do not satisfy them, there is a programming error, and if your estimates do satisfy them, we are unlikely to believe the results.

What BLP do to circumvent this problem is to let preferences for characteristics depend on consumer attributes, that is they rewrite the utility function as;

$$U_{i,j} = x_j \beta(z_i, \nu_i) + \epsilon_{i,j}$$

where (z_i, ν_i) are observed and unobserved individual characteristics. When micro data is in fact available we condition on the z_i of the individual and simulate the ν_i . When we only have aggregate data the relevant distinction is when there is data available on the distribution of the individual characteristic (the z_i) or not (the ν_i).

For example we could let preferences for price depend on income. This would allow low income consumers to be both more price sensitive and be attracted to low priced goods. If this is what the data preferred lower priced goods would be predicted to have higher elasticities and lower equilibrium markups (regardless of their shares). Similarly now when the price of a good goes up, very particular consumers leave that good; consumer’s who liked the characteristics of that good sufficiently to induce them to choose the good at the initial prices. Consequently those people will tend to substitute to other goods with similar characteristics; just the type of substitution pattern we want to allow for.

There is a question of when we require the ν_i . The answer depends on whether the major determinants of preferences over characteristics are included in the z_i or not. The extent of preference heterogeneity conditional on unobservable characteristics can be surprising. Thus for example, in MicroBLP we conditioned on the income, the number of adults, the number of children, the age (of the head) of household, and whether their residence was rural, urban, or suburban. That study had a particularly rich data set, as

each household reported its second as well as its first best choice. Table 2 provides the best price substitutes for selected models from demand systems for automobiles that were estimated in four different ways; (i) the full model allows for both the z_i and the ν_i , the logit models in (ii) and (iii) allow for only the z_i and the σ 's only model allows for only the ν_i . It is clear that we need the ν_i . The prevalence of the Caravan and the FS pickups when we use the logit estimates is a result of them being the vehicles with the largest market shares and the absence of observed factors which cause households to prefer characteristics differentially.

Table 3 illustrates the usefulness of the characteristic based demand systems in examining the impacts of adding or deleting goods from the choice set. MicroBLP (forthcoming) examines the impact of adding new high end sport utility vehicles, and of discontinuing the Oldsmobile division of General Motors (GM recently announced a decision to do so). Here we reproduce the table estimating the impacts of closing down Oldsmobile. It shows why GM is doing this; many of the vehicles that would gain significantly from the shut down of Oldsmobile are owned by GM itself.

Characteristic Space and the CPI.

Measures of welfare increases from product introductions obtained from estimated demand systems are more problematic than the measures of substitution patterns obtained from those systems. What we would like to do to measure the welfare gains from a new good is back out the distribution of reservation prices for that good. There is data that bounds the reservation prices of individuals who did not purchase the new good at one of the observed prices but did at another. However we never have an upper bound on the reservation prices of the “inframarginal” consumers who purchased the good at all observed prices.

Characteristic based demand systems are a bit more helpful in this context than demand systems in product space, as they allow us to compare the demand for the new product to the demand for previously existing products with similar characteristics. However the inference problem on the “extremes” of the observed characteristic tuples is the same as in product space, and there remains several issues associated with the evaluation of unmeasured characteristics. What this section hopes to show is that, despite this fact, moving to characteristic space provides a natural solution to at least

**Table 2: Price Substitutes for Selected Vehicles,
A Comparison Among Models*.**

Vehicle	Full Model	Logit 1 st	Logit 1 st & 2 nd	Sigma Only
Metro	Tercel	Caravan	Ford FS PU	Civic
Cavalier	Escort	Caravan	Ford FS PU	Escort
Escort	Tempo	Caravan	Ford FS PU	Ranger
Corolla	Escort	Caravan	Ford FS PU	Civic
Sentra	Civic	Caravan	Ford FS PU	Civic
Accord	Camry	Caravan	Ford FS PU	Camry
Taurus	Accord	Caravan	Ford FS PU	Accord
Legend	Town Car	Caravan	Ford FS PU	LinTnc
Seville	Deville	Caravan	Ford FS PU	Deville
Lex LS400	MB 300	Econovan	Ford FS PU	Seville
Caravan	Voyager	Voyager	Voyager	Voyager
Quest	Aerostar	Caravan	Caravan	Aerostar
G Cherokee	Explorer	Caravan	Chv FS PU	Explorer
Trooper	Explorer	Caravan	Chv FS PU	Rodeo
GMC FS PU	Chv FS PU	Caravan	Chv FS PU	Chv FS PU
Toyota PU	Ranger	Caravan	Chv FS PU	Ranger
Econovan	Dodge Van	Caravan	Ford FS PU	Dodge Van

From Berry, Levinsohn, and Pakes (forthcoming, *JPE*) “Estimating Differentiated Product Models from a Combination of Micro and Macro Data”.

Table 3: Discontinuing the Oldsmobile Division*

	Old Share	New Share	New-Old Share
All Oldsmobiles	.237	0	-.237
All GM	3.126	3.016	-.110
All Cars	9.711	9.695	-.016
Non-Olds Share Changes.			
Chevy Lumina	0.1354	0.1548	0.0194
Buick LeSabre	0.1216	0.1336	0.0120
Pontiac Grand Am	0.1322	0.1441	0.0119
Honda Accord	0.2955	0.3039	0.0084
Ford Taurus	0.2040	0.2115	0.0075
Saturn SL	0.1465	0.1539	.0074
Toyota Camry	0.2343	0.2415	0.0072
Buick Century	0.0614	0.0683	0.0069
Pontiac Grand Prix	0.0517	0.0584	0.0067
Chevy Cavalier	0.1700	0.1767	0.0067
Pontiac Bonneville	0.0658	0.0721	0.0064

From Berry, Levinsohn, and Pakes (forthcoming, *JPE*) “Estimating Differentiated Product Models from a Combination of Micro and Macro Data”.

one related problem.

Perhaps the most immediate of the problems that arise as a result of lack of measures of the gains from new goods is the problem of constructing a consumer price index in a world of changing choice sets. An ideal price index would measure the change in base period income that would leave the consumer just indifferent between the base and reference period choice sets (the “compensating variation”). Since the reference period contains goods that were not available in the base period, and not all the base period goods survive to the reference period, an ideal price index cannot be constructed without attaching values to the inframarginal consumers of both the new and the obsoleted good. The conceptual issues cited above, combined with the technical difficulties that would have to be solved (and the questionable assumptions that would have to be made) in order to produce an estimate of the compensating variation in a timely fashion, have induced the statistical agencies to “give up” on obtaining an ideal measure. Instead they have gone with procedures designed to provide an lower bound to the (average)

compensating variation.

Bounds in Product Space; Laspeyres Indices.

As noted by Konus (1935) long ago, in a world where the same products are in the choice set in every year the Laspeyres price index provides a lower bound to the true compensating variation. This is the logic that lies behind the matched model price indices which are still used for most components of the CPI in all countries. That is the components of the CPI are obtained from the price relatives obtained by data collectors who make repeated visits to the same outlet and form the ratio of the outlet's prices for the same good over adjacent periods.

In a world where the goods marketed in a given commodity group did not change over time, these "price relatives" would be averaged to obtain the group's "matched model" index. Since the goods marketed do change over time and there is a desire to maintain representativeness of the index, the data collector is instructed to rotate a certain per cent of the goods in the index out in every period, and when a good that is not scheduled to be rotated out is no longer sold at the outlet the data collector is instructed to make a "forced substitution" of another good. The sample rotation and forced substitution processes enable new goods to enter and old goods to exit the index, and the new goods biases are a result of the mechanics of how these processes work. Since they have no way of comparing the utility of the new goods to the old, the procedures used by the statistical agencies essentially amount to taking the averages of the price changes of the continuing goods as the required component indices (for more detail see Triplett, 2003, and Pakes, forthcoming, the article which forms the basis of the following discussion).

Bounds in Characteristic Space; Hedonic Indices.

The simplest hedonic index is just the Laspeyres index in characteristic space. To construct it we regress prices on characteristics and obtain the hedonic function in each period. This is a "reduced form" construct. In terms of our "static model" it is obtained by summing the marginal cost and the markup, and regressing that sum on product characteristics (x). As a result the coefficients of the hedonic function, say $h^t(x)$, have no interpretation in terms of our underlying primitives, and are likely to change over time as either the products in the market or the distribution of consumer attributes

change. Still the values of the function itself does summarize what a consumer has to pay for a given bundle of characteristics, and in our model those characteristics determine utility.

If we let C_t be the period's choice set (the list of the x tuples that can be purchased in t), the base period hedonic adjustment to the reference period's income of a consumer who purchased x_t in period t is

$$h^{t+1}(x_t) - h^t(x_t).$$

This is the change in the base period's income that would allow the consumer to buy the same bundle of characteristics in period two that it bought in period one and still have the same amount of income left over to buy other goods. I.e. it compensates the consumer for price changes by insuring the consumer the possibility of buying the same goods in the comparison period as were bought in the base period. Since the new choice set is generally different from the old we know that *provided x_1 is available in period two* (i.e. $x_1 \subset C_2$), a consumer with income $y + h^2(x_1) - h^1(x_1)$ facing the second period's choice set and hedonic surface (h^2, C_2) will generally choose a different x than x_1 , and any such choice will result in utility greater than the utility from the period one choice (since this is still a feasible choice).

Hedonic and Matched Model Indices; Conceptual Issues.

For goods that do not exit the only difference between the hedonic and the matched model's price changes is their estimates of $h^t(x)$. One uses the base period price of good x while the other uses an unbiased estimate of it. With enough data, they will give the same result.

The major conceptual difference in the two indexes arises as a result of their treatment of cases where $x_1 \notin C_2$. The matched model index is simply *not defined* in these cases, so what practitioners do is drop the good in question and link in another. This generates a selection problem. The goods that are sold in the first period and are not in the next tend to be goods whose characteristics were desirable relative to those of the other products sold in the first period, but were not in the next. Consequently they tend to be the goods which are intensive in characteristics whose values have fallen more than the average (often due to the entrance of products with superior performance). As a result

the matched model index construction procedure tends to throw out the left

tail of the distribution of price changes and produce a price index which is biased up-wards.

In stark contrast to the matched model procedure, the hedonic procedure is the *same* whether or not $x_1 \in C_2$. If $x_1 \notin C_2$, the hedonic estimate of the price in the second period, or $h^2(x_1)$, is obtained as a weighted average of the prices of products whose characteristics *are marketed* in the second period with the weights being larger for those products that have characteristic vectors similar to x_1 (the precise form of the weights depends on how one estimates the hedonic function).

For $h^2(x_1) - h^1(x_1)$ to be an upper bound to the compensating variation when $x_1 \notin C_2$ we require conditions which insure that the consumer will be better off with $h^2(x_1) - h^1(x_1)$ added to its income and C_2 then with its original income and C_1 . Pakes (forthcoming) provide two different conditions which insure that this is true, both of which can be verified empirically. One of these conditions is based on a gradient argument and hence requires there to be goods available in C_1 that are close to x_1 in characteristic space. The other requires the major characteristics of goods to be vertical (so we can substitute an unambiguously better good that is available in C_2 for the good that dropped out).

More generally, however, for a good which exits to cause a problem for our bound it must have gone from profitable to unprofitable despite still being preferred by a significant share of consumers (at price $h^t(x)$). This is an unlikely event and when it does occur we should be able to correct for it unless; (i) the good has major “horizontal” characteristics (characteristics which some consuming units prefers more of while others prefer less) and large fixed costs⁶, *and* (ii) the good does not have close substitutes in one of the two periods. The BLS’s analyst should check for such situations. A simple check for substitutes is to see whether the exiting good’s characteristics are inside the range of characteristics of the goods that are available in C_2 . I.e. if $x_i = [x_{i,j}]$ where j indexes the various attributes of the good, $\bar{x}_j = \max_{i \in C_2} x_{i,j}$,

⁶It is only then that a non-negligible fraction of the population could prefer the old good to the new goods without making the old good profitable to market. An example might be an old automobile (say a model “T” Ford) that antique car buffs might be willing to purchase at its old price (adjusted for inflation), but that would not sell sufficient quantity at that price to cover the sunk costs of producing it. Note that neither the conditions for the bound, nor the conditions which are likely to make it suspect, have anything to do with how the hedonic function would have shifted had the good that exited not exited.

and $\underline{x}_j = \min_{i \in C_2} x_{i,j}$, check if

$$x_i \in \{x : \underline{x}_j \leq x_{i,j} \leq \bar{x}_j, \text{ for } j = 1, \dots, J\} \equiv \bar{C} \in \mathcal{R}^J.$$

There are also a number of technical problems which have to be solved before we can implement hedonic estimators in a timely fashion. These are reviewed in Pakes (forthcoming) “Hedonics and the Consumer Price Index”). For now suffice it to say that the fact that the BLS’s data gatherers now carry hand held computers whose contents are downloaded nightly onto a central processing machine makes large scale use of hedonics feasible.

A Personal Computer Example.

A personal computer example will serve to show just how large the differences between the matched model and hedonic indices can be.

Table 4 summarizes characteristics of the data. The characteristics listed in the table (which are only a small fraction of the characteristics actually used to estimate the hedonic function in this study) are the major “vertical” characteristics of PC’s. The table illustrates two empirical facts which are often observed in markets with large amounts of technological change

- all the vertical characteristics are increasing (some rather dramatically) over time,
- only a small fraction of the goods marketed in one year survive until the next (compare the number of goods “matched to t+1” to those marketed in t).

Table 5 presents four hedonic and two matched model price indices. The first point to note is that details of how we construct the index (in this case how we estimate the hedonic function) do not seem to effect the estimates. They all show a large negative average price change (from 16% to 17% per annum), which varied significantly over the years, and was especially large in 97/98 the year the pentium II took over the market.

The row labelled “% matched” indicates that only 15% of the base-period observations are matched to a comparison period product. Though this is indicative of the possibility of a large selection bias in the matched model indexes, it is startling just how large a role selection seems to play. Even the Laspeyre’s index, the matched model index with the largest price decline,

Table 4: Characteristics of Data*.

year	95	96	97	98	99
nobservations	264	237	199	252	154
matched to $t + 1$	44	54	16	29	n.r.
characteristics					
speed (Mhz)					
min	25	25	33	140	180
mean	65	102	153	245	370
max	133	200	240	450	550
ram (MB)					
min	2	4	4	8	16
mean	7	12	18	42	73
max	32	64	64	128	128
hard disk (GB)					
min	.1	.1	.2	.9	2
mean	.5	1	1.8	4.5	8.5
max	1.6	4.3	4.3	16.8	25.5

* From Pakes (2003, *AER*) “ A Reconsideration of Hedonic Price Indices with an Application to PC’s”.

had a rate of decline less than a fifth of that of the hedonic indexes⁷. In most years the positive effect of selection just about offset the negative effects of technological change on the matched model indexes, and both indices were close to zero. However 1998 was different. In 1998 the pentium II obsoleted 92% of the older machines and the 1997 products that did continue were often early pentium II models; models which increased both their prices and sales. The Tornquist index, which weighs these products more heavily, was positive in that year while the Laspeyres was negative (interestingly this produces a correlation between the two frequently used matched model indexes of minus one).

⁷We did not provide estimated variances for the matched model index. This is because the number of matches was so small that we thought those variance estimates were unreliable. We note, however, that the estimated variances for two of the years were larger than the estimated variances of the proper hedonic indexes, and for two of the years they were smaller.

Table 5: PC Price Indexes*

		95/96	96/97	97/98	98/99	av.
Hedonic Indices						
$x \in C_{t-1}$	log-log	-.102 (.037)	-.111 (.059)	-.292 (.041)	-.172 (.092)	-.169 (.09)
	NLLS	-.097 (.04)	-.108 (.063)	-.295 (.045)	-.155 (.099)	-.164 (.09)
$x \in C_{t-1} \cap \bar{C}_t$	log-log	-.100 (.032)	-.115 (.054)	-.267 (.038)	-.161 (.062)	-.161 (.08)
	NLLS	-.094 (.039)	-.111 (.052)	-.270 (.044)	-.150 (.054)	-.156 (.08)
Matched Model Indices						
$x \in C_{t-1}$	Tornquist	.012	.002	.09	.011	.028
$x \in C_{t-1}$	Laspeyres	-.013	-.002	-.08	-.011	-.027
	% matched	16.6	22.8	8.0	11.5	14.7

*From Pakes(forthcoming *AER*) " A Reconsideration of Hedonic Price Indices with an Application to PC's". Standard errors appear in brackets below estimate.

PC's provide a rather dramatic example of how selection can effect matched model indices, and we do not expect as large a difference between hedonic and matched model indices in other component indices. Still, the example presented here shows that the use of hedonics could move us quite a way to providing better measures of price movements, and through those measures, better measures of growth and structural change.

References.

- Berry, Steven; Levinsohn, James and Pakes, Ariel. " Automobile Prices in Market Equilibrium." *Econometrica*, July 1995, vol. 63(4), pp. 841-890.
- Berry, Steven; Levinsohn, James and Pakes, Ariel (forthcoming); "Differentiated Product Models from a Combination of Micro and Macro Data: The

- New Car Market”, *Journal of Political Economy*, N.B.E.R. Discussion paper 6481.
- Bresnahan, T. (1981): “Departures from Marginal-Cost Pricing in the Automobile Industry”, *Journal of Econometrics*, vol. 17, no. 2, pp. 201-227.
- Haltiwanger, J., and S. Davis (1992); “Gross Job Creation, Gross Job Destruction and Employment Reallocation”, *QJE*, vol. 107, no. 3, pp. 819-863.
- Fudenberg, D. and D. Levine (1999), *Learning in Games*, Cambridge: MIT Press.
- Houthakker, H. S. (1955) ‘The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis’, *Review of Economic Studies*, vol. 23, pp. 27-31.
- McFadden, D. (1974) “Conditional Logit Analysis of Qualitative Choice Behavior” in *Frontiers of Econometrics*, edited by P. Zarembka, New York, Academic Press.
- Lancaster, K (1971) *Consumer Demand, A New Approach*, New York, Columbia University Press.
- Levinsohn J. and A. Petrin (2003); “Estimating Production Functions Using Inputs to Control for Unobservables”, *Review of Economic Studies*, April 2003, vol. 70, no. 2, pp. 317-341.
- Olley, S., and A. Pakes (1996) ‘The Dynamics of Productivity in the Telecommunications Equipment Industry’, *Econometrica*, vol. 64, pp. 1263-1297.
- Pakes, A. (1986) ‘Patents as Options; Some Estimates of the Value of Holding European Patent Stocks’, *Econometrica* vol. 54, pp. 755-784.
- Pakes, A. and Paul McGuire (2001), “Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the “Curse” of Dimensionality”, *Econometrica*, vol.69(5), pp. 1261-81.
- Pakes, Ariel (2003), “A Reconsideration of Hedonic Price Indexes with an Application to PC’s.” the *AER*, vol. 93(5), pp. 1578-1596. Working Paper, NBER # 8715, 2002.
- Pakes, Ariel (forthcoming), “Commonsense and Simplicity in Empirical Industrial Organization” the *Review of Industrial Organization*, NBER Working Paper # 10154, 2003.

Pakes, Ariel (forthcoming), "Hedonics and The Consumer Price Index", the *Annales de l'Insee* (volume in honor of Zvi Griliches), *mimeo* Harvard University.

Shaked, Avner and John Sutton (1982), "RElaxing Price Competition through Product Differentiation" *The Review of Economics Studies*, vol. 49(1), pp. 3-13.

Triplett, Jack (forthcoming), "The OECD Hedonic Handbook." *mimeo*, The Brookings Institution.