

New Economic Tools for Merger Analysis
by
Ariel Pakes, Yale University
for
”Economist’s Perspectives on Antitrust”
Charles River Associates Incorporated
Westin Hotel, April, 1997. *

*This is a selective overview of the new techniques available, concentrating primarily on the techniques that my coauthors and I have been developing. The notes themselves are copies of slides. There are many more slides in the notes than I will cover in the talk. The additional material provides detail and references for those who are interested. Most of the slides I will not cover have “Details” as the first word in their title.

Tools for Merger Analysis
Ariel Pakes, Yale University
April, 1997.

I will divide the main discussion into two parts.

- “Static Analysis” . An analysis of price changes that are likely to result from the merger if we hold the non-price characteristics of all firms (those of the merging firms, the other incumbents, and of the potential entrants) fixed at their premerger values. By “characteristics” we mean the goods they market, their capital stocks, etc.
- “Dynamic Analysis” . An analysis of the change in prices and firm characteristics that are likely to occur after the merger (this takes account of the fact that the merger will change the incentives which determine entry, investment, and exit decisions).

Static Analysis

I will try to give you a flavor of the analysis that can be done, the assumptions needed for the analysis, and what can go wrong when the assumptions are inappropriate. The goals of the analysis are to

- Provide a model which determines pre and post merger prices as a function of empirically analyzable magnitudes.
- Consider sources of information on the needed magnitudes, giving special emphasis to the role of demand systems and current techniques for estimating them.

Static Analysis:Model

I am going to focus on a "unilateral" pricing "differentiated product" model, as it underlies much of modern analysis. In this model

- Different products compete for consumers, and the extent of competition between any two of them (or their "similarity") is determined by the extent to which consumers view the goods as substitutable for one another (a concept that will be formalized below).
- Firms set prices and we assume that price choices maximize the profits of the firm given the prices and the characteristics of other firms (that is a firm could not earn higher returns from a unilateral change in its price).

Details: Unilateral Pricing Models

Note. The latter condition is the “Nash in Prices” or Bertrand, equilibrium condition. A similar analysis can be done for the other “unilateral” models which appear in the literature.

However models which allow for collusive pricing require a different (and slightly more refined) type of reasoning. Collusive models which are realistic enough to be analyzed for their detailed numerical implications are just now being developed (a bit more on this below).

Profit Maximization Conditions; Pre Merger.

We let the quantity demanded of good i be $q_i(\cdot)$, and it will be a function of the prices, as well as perhaps the characteristics, of all J competing goods (say (p_1, \dots, p_J) , and (x_1, \dots, x_J)). If we let $c_i(x_i, q_i)$ be the cost of producing q_i units of good i then the equilibrium condition is that firm i 's price, p_i , maximizes profits, or

$$p_i q_i(p_1, \dots, p_J, x_1, \dots, x_J) - c_i(q_i, x_i),$$

(conditional on the observed prices and characteristics of its competitors.)

Note this analysis is "static" because the characteristics of the products being marketed, and their cost functions, will be held fixed.

Static Pricing Equations; Pre Merger.

Letting the marginal cost of good i be $mc(q_i, x_i)$, the equilibrium condition insures that price is set equal marginal cost plus a markup equal to the inverse of the semi-elasticity of demand, i.e.

$$p_i = mc(x_i, q_i) + \frac{1}{[\partial q_i / \partial p_i] / q_i}$$

or, equivalently, that the percentage markup (markup/price) is equal to the inverse elasticity of demand (η_{ii}), or

$$\frac{p_i - mc_i(q_i, x_i)}{p_i} = \frac{1}{\eta_{ii}}.$$

Notes on the Markup Equation.

- If prices and marginal costs are known we can simply compute markups as their difference. Unfortunately reliable cost data are often proprietary and very difficult to obtain.
- An alternative way of getting some idea of the pre-merger markup is to estimate the elasticity of demand; i.e. estimate a demand system.

Details: Computation of Markups.

Note that this elasticity depends (through the demand function) on the prices and characteristics of all the goods marketed. However estimates generally imply that it is a fairly "smooth" function of these variables, so that a small change in price, say from pre to post merger prices, will not affect its value very much (the effect will be second order). When we use the term "approximate" below, we will mean that we are ignoring such second order effects.

Post Merger Prices.

In the firms that produce the first two goods merge, then the merged firm is concerned with the sum of the profits generated from the two goods, and the equilibrium condition for good one is that its price maximizes

$$[p_1 q_1(p_1, \dots, p_J) - c_1(q_1, x_1)] + \\ [p_2 q_2(p_1, \dots, p_J) - c_2(q_2, x_2)],$$

which is just profits from good one plus profits from good two.

If we index post merger prices with the superscript m , then for the post merger price to satisfy our equilibrium we require that the post merger markup be

$$p_1^m - mc(x_1, q_1) = \frac{1}{[\partial q_1 / \partial p_1] / q_1} + (p_2^m - mc_2) \frac{\partial q_2}{\partial p_1} / \frac{\partial q_1}{\partial p_1}.$$

The first term, $\frac{1}{[\partial q_1 / \partial p_1] / q_1}$, is (approximately) the premerger markup, so if

$$\theta_{21} \equiv \frac{\partial q_2}{\partial p_1} / \frac{\partial q_1}{\partial p_1}$$

i.e. if θ_{21} is the fraction of those customers who stop purchasing good one when good one's price rises, who then substitute to good two, the difference between the post and pre merger markup is

$$\theta_{21}(p_2^m - mc_2),$$

which is positive for a horizontal merger (i.e. for substitutes).

The Logic of The Post Merger Increase in Price.

Pre merger:

$p_1 \uparrow \Rightarrow \uparrow$ profits by q_1 , and
 \downarrow profits by $\Delta q_1(\Delta p_1)(p_1 - mc_1)$.

At equilibrium we increase price until the sum is zero. Post merger:

$p_1 \uparrow \Rightarrow \uparrow$ profits by q_1 , and \downarrow profits
by $\Delta q_1(\Delta p_1)[(p_1 - mc_1) - \theta_{21}(p_2 - mc_2)]$

Now increase price until the sum of these terms is zero.

Pre vs Post Merger Prices.

A similar pricing equation holds for the second good. Consequently we can solve for the post merger markups of each good in terms of the diversion ratios and the premerger markups of both goods.

Recalling that the p' s without the superscript are the premerger prices, we have

$$p_1^m - mc_1 = \frac{1}{1 - \theta_{21}\theta_{12}} [(p_1 - mc_1) + \theta_{21}(p_2 - mc_2)].$$

The increase in markup as a result of the merger will be higher

- the higher the two “diversion” ratios (θ_{21}, θ_{12}), or of those who substitute out of one good when its price rises, the higher the fraction that substitute to the second good (see Shapiro, 1995), and
- the higher the premerger markup on *each* of the merging goods.

Empirical Magnitudes

There are at least two ways of getting these.

- If we obtained pre merger markups by estimating a demand system (recall that they equal the price semi-elasticities), then generally that same demand system will generate estimates of $(\theta_{12}, \theta_{21})$.
- If we obtained premerger markups in some other way, and demand system estimates are not obtainable in the relevant horizon, we need another way of obtaining approximate estimates of the diversion ratios. Various survey techniques have been used to infer the diversion ratios (see below).

Survey Results: Caveats.

- If we use responses to a question about second choices to estimate the θ'_s , it is important to keep in mind that we need the second choices of the subsample of purchasers who would substitute out of good one were its price to rise, not the second choices of typical purchasers.
- Survey responses are not likely to be as accurate as actual purchases (the incentives underlying the choices are different).

First issue. Typically the “survey instrument” is designed to sample randomly from the group of purchasers of good one. The second choices of a random sample can give a very misleading estimate of θ . The purchasers who would substitute out if the price of good one increased,

are a select sample of the purchasers: they are particularly price sensitive. As a result their second choice will tend to be to different than those of the randomly chosen purchaser (they will tend to be to lower priced goods). One possibility is to ask first about whether the consumer would substitute out of the good when the price changes by

Details: Cost “Synergies” and Price Changes

We have not considered any cost changes that might result from the merger, and it is possible for cost synergies to overturn the result that post merger prices are higher than premerger prices. Since prices will fall with marginal cost, the net effect of mergers on prices depends on whether this effect can outweigh the “internalization” effect discussed above. Here is a suggestion for analyzing this possibility.

- Note, first that if we have estimates of the demand system, we can actually back out estimates of pre merger marginal costs. I.e.our formulas imply that $p_i(1 - 1/\eta_{ii}) = mc_i$.

- If we are willing to specify how synergies change these marginal costs and had the diversification ratios, then in most cases we can solve for the post merger prices implied by those synergies and the demand parameters (there is a system of “J” equations, one for the price of each good, in the “J” unknown prices).
- The suggestion is to do this for a reasonable range of cost synergies and trace out the prices they imply.

Caveats for Synergy Analysis.

- The only synergies which count for the static analysis are synergies which lower the *marginal* cost of production. Synergies which effect fixed costs do not matter.

- Note that if there are increasing returns, then the marginal cost we have from the pre-merger analysis would have to be changed for *both* possible synergies, and for possible scale effects, and these might tend to offset each other.
- This discussion only involves cases where the effects of the synergies are limited to the firms involved in the merger. If the merger generates new products or processes which are then (perhaps slowly) diffused to other firms in the industry, we would also want to quantify those effects.

Models for Demand Estimates.

There are two types of demand system models which are in current use.

- Demand systems where consumer preferences are defined directly on the products they purchase. E.g. Hausman (1992).
- Demand systems where consumer preferences are defined to be a function of the characteristics of the products they purchase, and products are defined to be particular bundles of those characteristics. E.g. Berry, Levinsohn and Pakes (1995).

Demand in Product Space.

- Almost all applied work in this genre uses the representative agent framework. It assumes a single agent who has a utility function defined on all goods, and derives quantities demanded by maximizing this utility function subject to a budget constraint.
- The solution gives the demand for each good as a function of the prices of all goods and total expenditure. If we estimated only one parameter for each price and expenditure effect, there are $J + 1$ parameters for each of J equations; or about J^2 parameters to estimate. For markets with a reasonable number of goods this is far too many parameters to estimate precisely from most currently available data sets.

Details: Demand in Product Space

It is possible to generalize from the representative agent framework, derive different demands for different households as a function of the household's characteristics, and then use the distribution of household characteristics in some public use data base to sum up to get market demand. This is a more computationally demanding strategy which, though not currently used, may well be used in the near future (since improvements in computers make it easier to accomplish).

The number of parameters that need to be estimated can be reduced by imposing the constraints implied by the representative agent utility maximizing framework (these include adding up, symmetry, and homogeneity restrictions). However the number of parameters we need to estimate will still grow like J^2 , and we will quite generally still need more restrictions.

Dealing With “Too Many” Parameters.

The usual way to deal with this is to place “a priori” restrictions on the system. This reduces the number of parameters at the cost of forcing the estimates to have certain characteristics regardless of whether the data favor those characteristics. Probably the most commonly used restrictions are the restrictions in the “multi-level budgeting” or the “utility tree” model (see, e.g. Deaton and Muellbauer, 1980; which is largely based on the work of Terrence Gorman). Then, our representative agent

- First allocates expenditures among certain “upper level” expenditure groups.
- Then it allocates each group’s expenditure among the goods within the group (without reference to which goods are purchased in other groups).

Details: Multilevel Budgeting

If we have a two level system with K groups multilevel budgeting will reduce the number of parameters by about a factor of K . There can be more than two levels. An example follows.

- Upper level: food, entertainment, transportation, . . .
- Within transportation: air, rail, bus, vehicles,
- Within vehicles: cars, minivans, sport utility vehicles, pickup trucks,
- Within cars: 2-seaters, 4-seaters, family cars, station wagons . . .
- Within family cars, Taurus, Accord . . .

Two Types of Restrictions Implicit in Multi-level Budgeting.

- To make the lower level maximization make sense, the preference orderings for the goods within a group cannot depend on the precise quantities of the goods we consume from other groups (else we would have to choose all goods jointly)
 - The allocation of goods into groups is done “a priori” (and is usually based on the characteristics of the product). Often a single good could be classed in different groups (is the land rover a luxury vehicle or a sport utility?) - and will be depending on the goals of the research. This makes it difficult to compare across estimates.

- The intuitive basis of the classification scheme is made difficult by the representative agent framework, but in a more micro world it is fairly simple: Do my preferences among family cars (which presumably depends on passenger size of the cars) depend on the (passenger size) of the sport utility that I own? If so, then sport utilities and family cars should be grouped together.
- To make the upper level maximization make sense we must be able to construct “aggregate” goods and associated “prices” in a way that allows us to choose the quantities for these aggregates by maximizing a utility function defined on them subject to the usual budget constraint. Without further restrictions (see below), this requires all goods in a subgroup to increase proportionately as expenditures in the subgroup

increase; for then the aggregate is just a constant fraction of each good.

- This implies that there can be no goods in the group which receive higher fractions of expenditures as incomes rise (no “luxuries”).

Stated in this way, the restrictions implicit in multilevel budgeting look rather strong.

Detail on the Restrictions of Multi-level Budgeting.

- Weak separability and cross price elasticities. The response of the quantity of a good from group r to the change of a price of a good from group q can only depend on the groups the goods belong to and the responses of the two goods to changes in expenditures on their respective groups.
- Formally, for the aggregation conditions to be met we require the subutility functions to be one of the “Gorman’s polar forms”. The easiest “polar form” is one which requires that within every group, the fraction of the expenditure that go to the various goods in the group are independent of the total quantity of expenditures in the group.

The group price index is then just the appropriate weighted average of the prices of the goods in the group.

- Gorman's second polar form allows the straight lines not to go through the origin (so the shares of some goods in the group can decrease as group expenditures go down) but then constrains the cross price derivatives between two goods even further; they only depend on the expenditure derivatives for the goods in question – i.e. they cannot even depend on the groups from which the goods are chosen.

Details on Multilevel Budgeting: Approximations

Recall where we started. We wanted to develop a utility maximizing theory, which abided by what seemed a sensible way of making decisions (multilevel budgeting) and then use the restrictions implied by that decision making process to reduce the number of parameters that need to be estimated to analyze a demand system. One way of interpreting Gorman's results is that this cannot be done in a way that does not restrict the utility function, and hence the implied demand patterns, in "unacceptable" ways.

One reaction to this is to find a system which is "acceptable" in terms of the restriction it imposes, but only approximately satisfies utility maximizing theory. Predominant among these systems, are systems that maintain the weak

separability assumptions, but do not really attempt to do the upper level allocation exactly. In particular, they do not produce exact quantity and price aggregates for the groups, thus allowing for a wider range of within group Engel curves. Instead they use approximate aggregates, price and quantity indices, which may be exact at some point in their sample, but are only approximations at other points. Provided weak separability is indeed appropriate, the error implied by the approximation should be confined to the expenditure allocations to the different lower level groups, and one has to judge whether this can seriously harm the estimates of the quantities of interest.

General Impressions: Multilevel Budgeting.

Whether a particular set of estimates are to be “trusted” depend on

- What they are to be used for, and
- Whether the assumptions made in getting them leads to a misrepresentation of what the data has to say on the issue at hand.

If, as in the analysis above, we are interested in estimating own and cross price elasticities, all we really care about is that we fit a functional form which is flexible enough to account for the differences in demand patterns generated by the different prices observed in the data. In this context, I believe the major worries should be,

- That our attempts to keep down the number of parameters estimated does not restrict the cross price elasticities in unreasonable ways (and if the division of products into groups, or the weak separability assumption, is inappropriate, this will tend to happen), and
- That the estimation method is appropriate for the data at hand, and will, therefore, generate a “good” estimate of the restricted price response surface.

The latter point is meant to emphasize taking proper account of estimation problems (like simultaneity, and mismeasured variables), problems which I have spent a good fraction of my working life on, but which I will not have time to cover in my discussion today.

Details: Further Caveats on Multi-level Budgeting

Other issues here that you might want to be aware of.

- There may be other questions you want to analyze with the demand system; for example the demand impacts of the introduction, or the discontinuance of different goods. Here the multilevel budget system becomes more suspect.
 - It does not allow us to do “ex ante” analysis of product introductions (since it does not allow us to analyze the determinants of demand or price responses for a product that doesn't exist).
 - “ex post” analysis will also be difficult because now we are not just making

small changes and trying to mimic the data, but to extrapolate demand in unexplored regions of the data. The representative agent model is also suspect here. The benefits to increases in variety often result from different consumers liking different goods.

- Another, less explicit attack on the “two many parameter” problem, is to simply define markets narrowly (dropping goods out of the analysis where it seems possible), and goods broadly (not differentiating between minor variants of the good in question; this can be justified by Hick’s aggregation theorem if the prices of the variants always move in proportion.) It goes without saying that this too can cause problems.

Demand in Characteristic Space

The Basics of the Model.

- Products are bundles of characteristics (a car is described by its size, horsepower, miles per gallon, quality of its service network, etc.)
- Consumer preferences are defined directly on these characteristics (and not on the product per se),
- Each consuming unit chooses the product that it most prefers (We use the discrete choice framework; see McFadden, 1981)
- Aggregate demand is obtained by explicitly summing over the demands of the different households.

Notes on the Basic Model

- A consumer's preferences over the various product characteristics depends on the consumer's characteristics (thus large families might prefer large cars, etc.),
- The consumer may choose the "outside" alternative, the alternative of not choosing any of the products in this market and spending all of its income on goods in other markets (this generates aggregate demand)

The Major Advantages of the Characteristics Model

- Main advantage: there is often a substantial reduction in the number of parameters that need to be estimated.
- To see this note that all we need to estimate is the distribution of preferences over characteristics. E.g. if there were C “important” characteristics, and it was sufficient to estimate the mean and the variance of the preferences for each (i.e. of the β 's multiplying each), then we would only need to estimate $2C$ parameters. Usually we want much more detail than that (we want to let preferences for product characteristics depend on household characteristics, etc.).

- However, it is clear that the number of parameters we need estimate is independent of the number of products, or J . Further once we estimate these parameters they determine all J^2 cross price elasticities – *no matter J*.

As a result we can typically include many more products in demand system estimated in characteristic (in contrast to product) space, and still obtain fairly precise parameter estimates.

Details: Further Properties of the Characteristic Model.

- Another advantage of this system is it allows us to analyze the likely impacts of new products *before* they are introduced.
- Though in principle it is straightforward to develop a model which allows individuals to purchase more than one unit of the product in every period (we would need preferences for “couples” or “n-tuples” of goods for each household), for computational reasons almost all models that are in current use are models that assume that each household buys at most one unit of the good in a given period.
- We expect the quality of the estimates from a characteristics based system to depend

on the extent to which one can capture preferences for products by preferences for the observable characteristics of those products (which is likely to vary with product group), though we will show how to incorporate unobservable product characteristics below.

Early Discrete Choice Models.

The early models in this spirit are exemplified by the simple logit model.

The reference for the current generation of models is Berry, Levinsohn, and Pakes, 1995, who simply

- Modify the simple logit in two ways. These are designed to
 - Generate more sensible price elasticities, and to
 - Take account of the fact that we seldom have information on all the relevant characteristics.
- Points out ways of increasing the precision of the estimators from the model.

- Provides a computationally practical estimation algorithm.

I will not discuss computational issues (they are based on a simulation technique introduced by Pakes,1986, and an inversion technique introduced by Berry,1994), and will only briefly consider estimation problems.

The Simplest Logit Model.

Each product is endowed with a mean utility which is a function of its characteristics and price, say

$$\delta_j = \delta(x_j, p_j).$$

The preferences of consumers are assumed to be distributed independently about this mean. I.e. the utility of consumer i for good j was written as

$$U_{ij} = \delta(x_j, p_j) + \epsilon_{ij},$$

where the ϵ_{ij} are distributed independently across products and across consumers for a given product.

A simple “two step” estimation algorithm could be used. First estimate each of the $\{\delta_j\}$. Then regress these estimates on (x_j, p_j) .

Problems With The Simple Logit Model: Price Elasticities

- δ_j is the only characteristic of good j that determines demand. Consequently the model constrains all price elasticities to depend only on the δ_j .
- Since there is a one to one relationship between the δ_j and the shares of sales, any two goods with the same shares are constrained to have the same cross price elasticity with any third good, and are also constrained to have the same own price elasticity (if they do not, there is an error in the computer code).

- E.g. Consider the Yugo (which was the lowest priced car in our sample), and the Mercedes. They both had about the same shares. So estimates from the simple logit model *would have to* imply
 - That both cars have the same cross price elasticity with any third car, say with the Geo Metro. I.e. the estimated demand system will imply that as I raise the price of a Geo Metro and consumers substitute out of that car, just as many substitute to the Mercedes as to the Yugo.
 - The “semi-price elasticities” of the two cars are identical. Recall that in Nash pricing equilibrium these semi elasticities determine markups. So our model will imply that the Yugo (which sold in 1990 for about \$3500) had the same markup as a \$65000 Mercedes.

Obtaining Reasonable Cross Price Elasticities.

- We allow different households to prefer different characteristics differently

$$U_{ij} = \sum x_j \beta_i^x - p_j \beta_i^p + \epsilon_{ij},$$

where the β_i differ among consumers (we come back to how they are determined below).

- Now increase the price of good j . Some people who purchased good j will substitute out to other cars. The people who purchased good j were people who preferred the characteristics of good j and hence will tend to substitute to other cars which are similar to good j .

- Also the people who buy high priced cars will tend to be consumers with low β_i^p , and this will tend to generate low semi-price elasticities, and hence high markups, for high priced cars.

The Determinants of the β_i

We allow the β_i to depend on both measured and unmeasured household characteristics. For e.g. we allow families with high measured incomes to be less responsive to price (have a lower β_i^p). Also, since, there may be sources of wealth that we do not measure and different degrees of price sensitivity among households with the same wealth, we allow for an interaction of price with an unmeasured variable whose mean and variance are estimated. Precisely how you do this depends on the data that is available. For a discussion see Berry Levinsohn and Pakes, 1997.

Unobserved Product Characteristics.

- Because we seldom have information on the whole range of characteristics that effect people's behavior, we allow for an unobserved (ξ_j), as well as for the observed (x_{kj}) product characteristics.
- It is easy to see the role of unobserved characteristics in the simple logit model, in which case all we need to do is let the δ_j (the mean utility) to depend upon ξ_j as well as on the (x_j, p_j) , i.e

$$\delta_j = \sum_k x_{kj} \beta_k^x - p_j \beta^p + \xi_j.$$

- The ξ_j play the role of the disturbance in this equation (which is estimated in the simple logit model). Thus if cars with higher unmeasured qualities (higher ξ_j) also charge

higher prices (as our model of price setting implies), then ξ will be correlated with p and OLS estimation algorithms are inappropriate.

- E.g. In the car example, if we only included standard characteristics of cars like size, horsepower, number of airbags, . . . , the Lexus 400 would look quite similar to many upper level family cars. It has a higher price, and this is presumably because both the consumers and the retailers know that the Lexus is better along the dimensions of quality that we do not measure.
- This is the same simultaneity problem as we discussed in the traditional demand analysis and we have to modify our estimation algorithm in similar ways. If we do not do this modification we will tend to get much lower price elasticities than is appropriate.

Dynamics: What is “Doable” .

My overall view of what is possible here is different than in the static analysis. There we could attempt to get fairly realistic estimates of the quantities of interest. Here I think of what we are doing as a reality check on the rather fuzzy lines of reasoning that are often used in the dynamic analysis of mergers.

We are interested in evaluating the likely impact of the merger on entry and investment decisions; and then analyze

- the extent to which the changes in market structure that result from the entry and investment decisions, imply post merger prices that are different than the static prediction for post merger prices given above.

- the change in entry and investment costs (as well as in non-price aspects of consumer benefits), that are likely to occur if the merger is permitted.

On Cost Savings, Mergers, and Investment.

- The analysis of possible cost savings from mergers has largely focused on synergies in production costs. Whether or not there will be cost savings of this form depend on details of the production process; details that are often proprietary and difficult to verify.
- On the other hand simple economics tells us that the merger changes incentives for investment, and, at least for the merging firms, almost always changes it in the same

direction. Premerger the firm ought to invest until their marginal gain just equals the marginal cost of investment. Part of the gain to one firm's investment is a result of business taken away from other firms, (either through increased quality or decreased costs and then prices; see for eg. Mankiw and Winston, 1986). The merged firm will internalize this externality (as should society), and invest less. This is a real cost savings.

- There may also be a decrease in consumer surplus as a result of a smaller amount of investment, so we cannot say, without further analysis, whether society should prefer to have the lower investment path. On the other hand the analysis that has been done to date suggests that often the gains from the fall in investment costs far exceed whatever losses in consumer benefits

(E.g. Pakes and McGuire,1994, or Berry and Pakes,1993).

- Note also that investment cost savings can (and probably often do) provide the incentive to merge. This is a very different reason for merging than the reason traditionally associated with mergers (where it is often presumed that mergers are designed to increase market power and price), and should be treated differently by society.

Mode of Analysis.

I am going to assume that we know

- The results of a static analysis something like the analysis described above (or at least can provide a simple characterization of them).
- The likely costs and time required both to enter, and to expand existing facilities (or at least provide a range of possible values for them).

Then I am then going to describe a computer program which

- Takes as inputs different guesses for these magnitudes.

- Assumes that firms make investment, entry and exit decisions to maximize the expected discounted value of their future net cash flows.
- Computes equilibrium investment, entry, and exit policies for all firms. (Here again equilibrium is in the Nash sense; i.e. each firm is doing the best it can do given the actions of its competitors).
- Takes the equilibrium policies and computes their implications for prices and costs, and therefore for profits and consumer and producer benefits, over time.

The purpose of the algorithm is to allow us to compute what different primitives imply for the objects of interest; price cost margins, entry and investment costs, consumer and producer benefits, etc.

Details: Accessing the Algorithm, and its History

The algorithm used here was first developed by Ariel Pakes and Paul McGuire (1994). It makes extensive use of theoretical environments developed in Makin and Tirole (1988), and Ericson and Pakes(1995).

A public use version of the algorithm was written by Gautum Gowrisankaran and Ariel Pakes. To access it FTP to "econ.yale.edu", use "anonymous" to login, and your own name as the "password". The user should change directory to "pub/mrkv/eqm" and copy all the needed files. There is a "readme" file to start you off; and a description of the program in postscript. Alternatively, you can find both a link that will allow you to download the needed files, and a description of the program, on my web site.

Detail: The Nature of Equilibrium.

Each firm chooses its policies to maximize the expected discounted value of its future net cash flows. The future earnings a firm will make as a result of its investment decisions are determined, in part, by the likely states of its competitors, or the “market structure”, in the future. Thus in order to make optimal investment decisions the firm must have some perceptions of what the market structure is likely to be tomorrow conditional on current information. Given a perceived distribution of market structures, each incumbent can make its exit and investment decision, and each potential entrant can decide whether to enter, by maximizing their expected future values. These investment, exit and entry decisions generate an objective distribution of future market structures. Equilibrium occurs when the perceived distribution of future market structures generate an objective distribution which coincides with the perceptions.

Different Choices in the Program.

This is a click on program (directions for accessing it are in your handout). It has three parts or modules. The details of each module can be modified to fit the environment being analyzed (more on this below); but I will not spend time on any but the most basic variations.

The first module has the user specifying which of three static I.O. paradigms seems most appropriate, and then setting the parameters needed to analyze dynamic behavior in that framework. These are

- A differentiated product model (similar to that discussed above) in which;
 - Each good has a quality,

- Investment is designed to improve that quality,
 - Prices in each period are the Nash in prices (or Bertrand) equilibrium prices for the distribution of qualities in that period.
- A homogeneous product model in which;
 - There are differences in capacities among firms,
 - Investment increases capacity,
 - Quantities in each period are the Cournot (or Nash in quantities) equilibrium quantities for the distribution of capacities in that period.
- A homogeneous product model in which;

- There are differences in marginal costs across firms,
- Firms invest in capital to reduce marginal cost,
- Quantities in each period are the Cournot (or Nash in quantities) equilibrium quantities for the distribution of capital stocks in that period (the same as above).

Detail Module 1: Calculating Profits, The Differentiated Products Example.

Let (x_1, \dots, x_J) represent the qualities of the products of the J firms active. Then, if for expositional simplicity we assume constant marginal cost, we have from our previous discussion

$$p_j = mc_j(x_j) + \frac{1}{[\partial q_j(x_1, \dots, x_J, p_1, \dots, p_J) / \partial p_j] / q_j}.$$

These are J equations in the J unknown prices. The program uses them to solve for

$$p_j = p_j(x_1, \dots, x_j).$$

These p_j 's and x_j 's are substituted into the demand system to determine quantities [the q_j 's] as a function of the x 's. Profits are then determined as

$$\pi_j(x_1, \dots, x_j) \equiv q_j(\cdot)[p_j(\cdot) - mc(x_j)].$$

This is done for every possible value of the (x_1, \dots, x_J) tuple.

Module 2: Entry, Exit, and Investment Policies.

The second module of the program takes the profit function as input, and produces equilibrium entry, exit, and investment policies as output. Again these are calculated for all incumbents and all potential entrants at each possible quality tuple.

- Entry (the simplest model). The entrant sinks e dollars. If the entrant enters it enters in the next period at location x^e with probability p^e . A potential entrant enters only if the expected discounted value of net cash flows generated by entering exceeds the cost of entry (x^e).
- Exit. Each incumbent calculates its continuation value (its value if it would continue), and continues only if the continuation value is larger than ϕ dollars (the

selloff value; or the quantity that the firm could get were it to use its assets in another activity).

- Investment. If the firm continues, it chooses an amount to invest which maximizes firm value. The more investment the firm makes the more likely it is that the firm's state variable (the quality of its product, its capital, or its capacity, depending on the model) increases in the next period.

Details: Module Two.

- The user needs to specify the parameters of the entry process (e, x^e, p^e) , the alternative value of the firm ϕ , a distribution for the increase in its state (for $x_{t+1} - x_t$) conditional on investment, and a cost of investment function. The reasonableness of certain choices can often be judged by the relationship between the parameters of the model (thus we can consider entry costs relative to the total profits typically earned per period, etc.).
- One version of the model differentiates potential entrants by differing costs of entry, and the entrant that appears during the period is chosen by a random draw from that distribution.

- We could, and probably should, modify the program to allow for different combinations of sunk cost and marginal cost (or quality of the good) when we enter. This would allow us to model the distinction between “committed” and “uncommitted” potential entrants that is in the guidelines.

Module 3: The Output Program.

The third module takes the investment, entry and exit policies generated by the second module as input, then simulates market outcomes from an initial condition that the user supplies, and collects statistics of interest that result from the simulation.

The “initial condition” specifies where the market is today. For merger analysis we will want to compare the output from the program when we specify the initial condition to be

- The condition that the market would be in, if the merger were not to take place, to
- The condition the market would be in were the merger to proceed.

From each initial condition, the program

- Uses the investment, entry and exit decisions calculated by the second module to generate likely sequences of future industry structures, and then
- Accumulates data from the market structures that are generated in this fashion. The data include measures of: price cost margins, concentration ratios, entry, consumer welfare, profits, investment and more.

Some Current Limitations of the Framework.

- We do not allow for the possibility of collusion (for prices or quantities not to be set by “unilateral” behavior, but rather by either explicit or implicit agreement among the agents). It is possible to analyze the possibilities for collusion in this dynamic framework, and Chaim Fershtman and I have made some progress in this direction, but this “module” will not be ready for some time.
- The framework does not allow for learning by doing, or other factors which make demand or cost in year t depend on price or quantity decisions in year $t - 1$. Again, the framework is being modified to allow

for this possibility (by Lanier Benkard), but this is not yet ready for use.

- The framework ignores the possibilities of future mergers, and this can impact on the benefits and the costs of current merger activity. Again, some progress has been made on modifying this framework to endogenize merger activity by Gowrisankaran(1995), but this program, though available from that author, is still quite difficult to use.

Current Limitations of the Framework: Computational.

- Though the public use program does allow the user several options, it is easy to think of modifications that one might want to undertake in order to better approximate a given institutional environment. In order to make such modifications a reasonably computer literate person will have to spend some time familiarizing him or herself with both the program and the Pakes Mcguire article.
- The public use program is currently set up to handle problems where firms are distinguished from each other by only one state variable (quality of product, or capital stock, . . . , but not quality of product

and capital stock, or two quality dimensions). It is easy to modify the program to allow for two or more state variables per firm (and we have already done it in our own research), but at least right now those modifications are not included in the public use version.

- Part of the reason richer sets of assumptions are not included in the public use version of the algorithm, is that once richer sets of assumptions are used it becomes important to use more efficient computational techniques than those employed in that algorithm. In particular the memory and time requirements of the public use version have limited the experiments we do on it to single state variable per firm problems with seven or less firms active, and two state variables per firm problems with

five or less firm's active. These programs run in about two hours on our Sun Ultra's (smaller problems run much quicker, and larger problems much slower).

- Two more powerful algorithms have been developed. First Gowrisankaran(1997) has produced a modification to the original program which allows it to compute problems for multiproduct firms more efficiently. You can contact him for more details.
- There is now a second Pakes McGuire algorithm (Ariel Pakes and Paul McGuire,1996), and it is both quite a bit more powerful than the first and much easier to program. Indeed it is so easy to program that it was not obvious to us that we needed a public use version of it (though we are now thinking of producing one). On the other hand

this algorithm uses very different computational principles, principles that might have to be explained to your programmer (and we would be happy to do so.)