NBER WORKING PAPER SERIES

A DYNAMIC ANALYSIS OF THE MARKET FOR WIDE-BODIED COMMERCIAL AIRCRAFT

C. Lanier Benkard

Working Paper 7710
http://www.nber.org/papers/w7710

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2000

I would like to thank Ariel Pakes, Steven Berry, Donald Andrews, and John Rust for their invaluable advice. I also thank Daniel Ackerberg, Patrick Bajari, Tom Crawford, Dennis Epple, Gautam Gowrisankaran, Phillip Leslie, Michael Riordan, and many seminar participants for their helpful comments. Financial assistance was provided by the Alfred P. Sloan Foundation. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft
C. Lanier Benkard
NBER Working Paper No. 7710
May 2000
JEL No. L13, L11, L62, L50

ABSTRACT

This paper develops a multi-agent dynamic model of the commercial aircraft industry and then uses that model to analyze industry pricing, industry performance, and optimal industry policy. In the model, firms are differentiated in their products and cost structure, and entry, exit, prices, and quantity sold are endogenously determined in dynamic equilibrium. Re ecting the focus of the paper, demand and supply are modeled structurally, while investment is modeled in reduced form. The model utilizes a cost model of commercial aircraft production developed and estimated in a previous paper (Benkard (2000)), and a discrete choice model of commercial aircraft demand to determine static profits. I find that many unusual aspects of the aircraft data, such as high concentration and pricing below the level of static marginal cost, are explained by this model. The model also replicates the stochastic evolution of the industry well. Many of these properties could not be explained with a static model. These results provide support for the structural dynamic modeling approach in general. I also find that the unconstrained Markov perfect equilibrium is quite efficient from a social perspective, providing only 9% less welfare on average than a social planner would obtain, but that the Markov perfect equilibrium shifts a substantial amount of welfare from consumers to producers. Finally, I provide simulation evidence that an anti-trust policy in the form of a concentration restriction would be welfare reducing with high probability.

C. Lanier Benkard
Graduate School of Business
Stanford University
518 Memorial Way
Stanford, CA 94305-5015
and NBER
lanierb@leland.stanford.edu

1 Introduction

Industrial policy in the commercial aircraft industry is a subject of recurring debate. Many previous authors have argued that, due to learning by doing and massive entry costs, the conventional wisdoms regarding trade and antitrust policies may not apply to the aircraft industry. However, an analysis of policy alternatives that properly accounts for these properties requires a significantly more complex model than has typically been used. The goal of this paper is to take a first step in that direction.

More specifically, because of strong intertemporal incentives due to learning by doing, the commercial aircraft industry exhibits features that are not consistent with static optimization. For example, it is not uncommon to observe prices below the level of static marginal cost, particularly upon introduction of a new product. Thus, static models are severely limited in their ability to provide us with a meaningful analysis of the aircraft industry. Since there has been little progress to date in making dynamic models tractable enough to apply to empirical problems, in evaluating industry performance and policy alternatives the previous literature has primarily relied on the predictions of relatively simple theoretical models.

There are two weaknesses of the theoretical approach. The first is that the models tend to be quite stylized, not reflecting any industry particularly well. The second is that many of their predictions are ambiguous in practice. Policy prescriptions in theory depend on the exact parameter values of the model. Both of these shortcomings point to the need for empirical work to provide a more definitive analysis.

This paper represents a first attempt at analyzing the aircraft industry using a dynamic equilibrium model that is estimated consistently from observed data. Since the goal of the paper is to bring the model to data, the model is tailored to include key features of the commercial aircraft industry. The model focuses specifically on the market for wide-bodied commercial jets since that market contains a more tractable number of products than the commercial aircraft market more generally.[1] The wide-body market is also the largest[2] and most prominent segment of the commercial aircraft market. The model allows for closed-loop strategic interaction between firms that are differentiated in their products and their production cost. The structure of the model — including the cost function, the demand function, and the evolution of the firms' endogenous state variables — attempts to reflect the actual structure of the industry, and the supply and demand systems are estimated (separately) using industry data in a manner consistent with the underlying dynamic model. I then insert the obtained parameter estimates into the dynamic model and numerically compute the equilibrium of this model.

The advantages of this approach are two-fold. First, because equilibrium is not enforced in the estimation procedures, consistency of parameter estimates does not depend on the particular equilibrium assumptions made and is therefore robust to a wide set of possible assumptions. Second, there is no need to solve the dynamic programming problem during estimation, which greatly reduces the computational burden of the estimation procedures. The disadvantage of the approach is that, if the equilibrium assumptions are true, then greater efficiency could be obtained in the estimates by enforcing equilibrium during estimation.

I find that, despite some simplifications, the dynamic model predicts many aspects of equilibrium behavior well, particularly those that have been the focus of the past theoretical literature. It improves vastly on previous attempts at modeling aircraft industry pricing. For example, even though observed markups vary over a wide range, the model predicts both price levels and price movements that are similar to those observed, including many instances of below static marginal cost pricing. The model tends to predict equilibrium prices and markups that are slightly higher than those observed, but I do not feel that this tendency is a shortcoming of the theoretical model. Rather it is largely attributable to an arbitrary dimensional restriction placed on the model for computational reasons. The model also represents many aspects of the industry dynamics well, generating entry, exit, concentration ratios, plane value, and plane type distributions that are similar to those observed.

[1] The term "wide-body" refers to a plane with more than one aisle separating seats. The first wide-body, the 747, was introduced in 1969. As of 1997 there are seven wide-bodied jets in the market.

[2] The wide-body market is largest in value terms but not in volume terms.

The model is also well suited to a detailed analysis of the industry, including alternative market structures and industrial policies. I have considered three alternative market structures, single-product firms, a multi-product monopolist, and a multi-product social planner. The results from this comparison suggest that the single-product firm Markov perfect equilibrium (MPE) is quite efficient from a social perspective, providing only 9% less total welfare on average than the social planner could obtain. However, relative to the social planner, the MPE shifts a substantial portion of total surplus from consumers to producers. I also find that an unconstrained multi-product monopolist with no threat of entry would lead to large inefficiencies from a social perspective.

I go on to consider an anti-trust policy which places a *per se* restriction on the highest market share any single firm may attain. I find that such a policy would be welfare reducing with very high probability, particularly hurting consumers.

2 The Commercial Aircraft Industry: Some Background and Motivation

Total commercial aircraft industry revenue for 1997 was approximately $60 billion, of which $40 billion is attributable to U.S. producers. In many years commercial aircraft has been the U.S.'s largest net export, with trade surpluses averaging about $25 billion annually over the early 1990s. The commercial jet aircraft industry has existed since 1956, but the first wide-body (the 747) was not introduced until 1969. Sales of wide-bodies have grown steadily since then so that in 1997 they accounted for approximately 60% of total industry revenue (30% of units).

The commercial aircraft industry, and aerospace more generally, has seen much merger activity in recent years which has led to increased concentration. For example, since 1980 Lockheed-Martin (which no longer produces commercial aircraft but is a major military producer) has absorbed 17 other companies, two of them major aerospace industry players.

Frequently, though not always, mergers have come about when the future viability of a weak firm has come into question. The recent Boeing–McDonnell-Douglas merger would likely have been blocked if this had not been the case. The resulting industry for commercial jets of more than 100 seats consists of only two major producers.

Finally, commercial aircraft is regarded by many countries as a "strategic" industry, meaning that its presence is essential to the nation's well-being. As such it has frequently been the target of industrial policy, most notably in Europe, where government supported efforts at developing a viable industry suffered many failures before finally experiencing success with the Airbus consortium.

2.1 Dynamic Model Motivation: Commercial Aircraft Pricing

As an example of industry pricing policy, Figure 1 graphs estimates of price [P] and average variable cost [AVC] for the Lockheed L-1011. The price series shown is a nearest neighbor smoothed transaction price series constructed from a data set that contained sales prices for approximately 60% of the units. The variable cost series shown was constructed using data from Lockheed's annual reports. With the exception of the first data point which covers two years, all observations for the cost series are annual averages. It should be noted that due to incomplete accounting this series contains more error than the price series.[3] Also, because the first data point covers more than 50 units, it does not show very well that variable cost for the first few units produced was much greater than the later second peak in cost that occurs near unit 160.[4] However, even with these faults, the data exhibits the two traits that I would like to emphasize.

First, note that AVC exceeded P for much of the 14 year period that the plane was produced.

[3]In particular, costs seem to have been recorded in such a way as to "front-load" the data, i.e., costs in a given year actually somewhat reflect units that were shipped in the next even though some considerable effort was undertaken to eliminate this feature.

[4]Using the production data provided by Lockheed, it was possible to make a good estimate of the average variable cost of the first few planes. This method suggests that AVC for the first ten planes was approximately $220 million per unit, which is about three times the AVC for 1979, where the second peak in cost occurs.

Figure 1: Lockheed L-1011: Price Vs. AVC 1972-1985

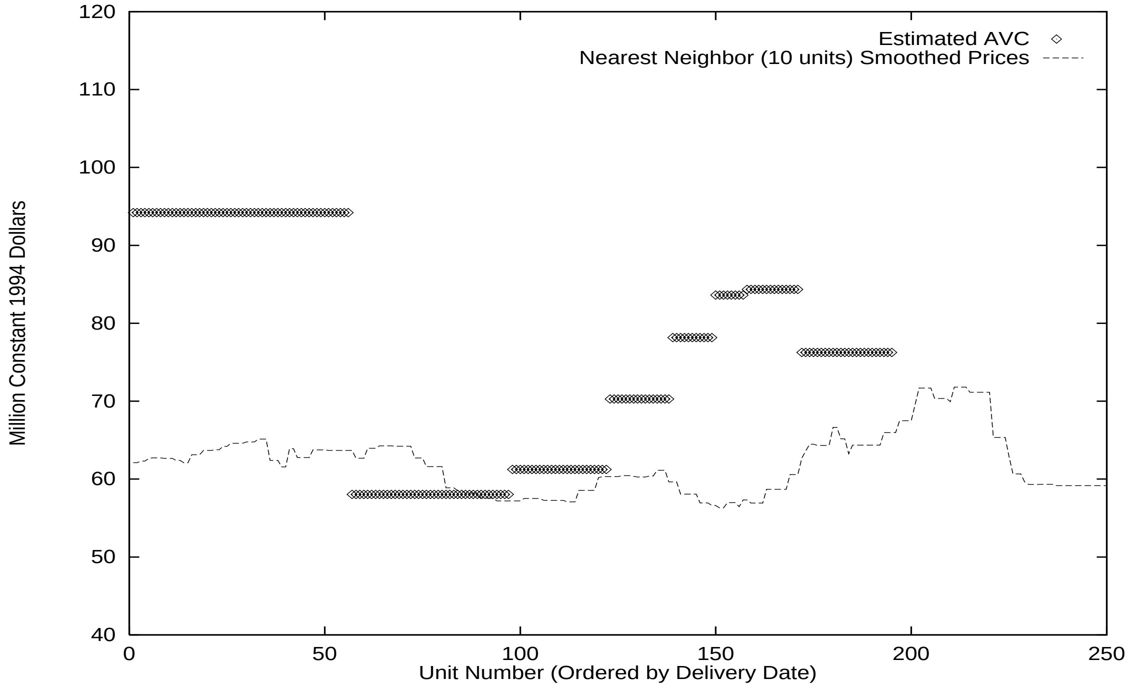

This example illustrates the need for a dynamic model in order to effectively analyze this industry. No static model could rationalize the prices that we observe in the data. In fact, without further knowledge of the industry, a summary inspection of this graph might lead to the conclusion that Lockheed may not have been acting optimally in both pricing the L-1011 and in remaining in the market so long.

Second, the graph shows that there is much greater variance in cost than in price. This property is assumed to hold for all products in the market since observed prices exhibit yearly variance of no more than 10-20% while many authors (e.g., Benkard (2000)) have shown that, due to learning curves, the first few planes can be as much as five to six times more costly to produce than the one-hundredth plane.

The dynamic model presented below replicates both of these properties. Section 7 also shows that equilibrium prices for the L-1011 predicted by the model are quite similar to those in Figure 1, which both lends support to the model and helps to explain why this kind of pricing behavior may in fact be optimal in certain circumstances.

3 Previous Work

Learning curves have been found to be important in many industries,[5] and there is a large literature documenting learning curves in aircraft production (e.g., Wright (1936), Alchian (1963), Asher (1956), Gulledge and Womer (1986), Benkard (2000), et al.).

The theoretical literature on competition with learning curves is comparatively sparse, but it does provide significant insight into industries with learning, specifically that learning curves can provide strong strategic incentives to firms. Several authors (Fudenburg and Tirole (1983), Dasgupta and Stiglitz (1988), Cabral and Riordan (1994), et al.) have shown that learning curves can lead to pricing below the level of static marginal cost, high industry concentration, and aggressive competition even when industry concentration is high. Stylistically, all three of these properties have been exhibited by the commercial aircraft industry at one time or another.

Fudenburg and Tirole (1983), show that in a duopoly with learning the link between current price and current cost is very loose. Specifically they show that as firms work down their learning curves, prices may actually rise. Cabral and Riordan (1994) show that in a duopoly with learning there is increasing dominance, i.e., there is a tendency for the firm with lower costs to increase its lead. This paper is more in the spirit of Dasgupta and Stiglitz (1988), which considers the evolution of an industry's structure, industry performance, and optimal policy. They show that when learning is strong, an oligopoly with learning may tend toward monopoly, an unrestrained monopolist may be socially preferable to any market with more than one firm (supporting restrained anti-trust policy), and that it may increase a country's welfare to protect an infant industry.

However, the theoretical learning models are quite specialized, making it difficult to apply them to any specific policy question for a specific industry. Fudenburg and Tirole (1983), Dasgupta and Stiglitz (1988), and several other authors cover the cases of Cournot and

[5]Argote and Epple (1990) cover over 100 studies documenting learning curves in many widely varying industries including both manufacturing and services.

Bertrand oligopoly. Cabral and Riordan (1994) utilize a more general price-setting differentiated products duopoly. Furthermore, many of the models' predictions are ambiguous in practice, and even in cases where these models predict an unambiguous outcome, it can be difficult to tell the magnitude of the effect. Thus there remains a gap in the literature between the empirical models that document the existence and extent of learning and the theoretical models analyzing oligopolies with learning. This paper attempts to bridge this gap. In the model presented here the key features of the theoretical learning models are retained, but the parameters of the model are estimated econometrically.[6]

The model used here also relies heavily on recent advances toward the development of an empirical framework for dynamic oligopoly models put forward by Ericson and Pakes (1995). Ericson and Pakes (1995) introduced a class of multi-agent dynamic models that can be solved computationally and are well suited to a variety of empirical problems. Gowrisankaran and Town (1997) also apply a model of this type to the hospital industry. Contrary to the learning literature, the currently existing models in this class contain only static pricing equilibria. Learning curves have the consequence that current prices and quantities influence future costs, and hence that dynamic equilibrium is not consistent with either static price-setting or static quantity-setting, the two examples that have been used in previous models of this type. The dynamic model presented here is similar to the Ericson-Pakes class of models in its use of dynamic equilibrium in a multi-agent setting. But, prices and quantities are endogenously determined in a *dynamic* quantity-setting equilibrium in the tradition of the theoretical literature on learning curves.

4 The Model

This paper essentially adapts the theoretical learning models from the literature to an empirical framework similar to that of Ericson and Pakes (1995). The Ericson-Pakes framework

[6]Baldwin and Krugman (1988) is also somewhat similar in spirit to this paper except that they calibrate their model rather than estimate it. This allows them to use a much simpler single-product duopoly learning model similar to those used in the theory literature. They also solve their model for an equilibrium with precommitment (meaning that the equilibrium is not subgame perfect).

was used because it is easily expanded to cover a variety of empirical dynamic problems and because it facilitates computation of the model equilibrium. However, the specific model used here differs from Ericson-Pakes quite substantially in that prices and quantities are endogenously determined in dynamic equilibrium, rather than being a by-product of the spot market. Additionally, in order to make the model rich enough to match observed data the model contains several exogenous state variables.

The model describes an infinite-horizon discrete-time industry with endogenous entry, exit, and quantity-setting, where firms choose strategies in order to maximize the expected discounted value [EDV] of their net future profits given their information set. Prices are endogenously determined through a differentiated products demand system. Investment in product quality is modeled in reduced form.

In this model, industry structures are represented by states that summarize all currently available information relevant to current and future payoffs. Each active firm is assumed to have three state variables, its experience level (with respect to the learning process), and two state variables representing the quality of its product described in more detail below. There is also one industry-wide state which determines the overall level of demand in each period.

Firms' quality states and the aggregate demand state evolve exogenously according to fixed Markov transition matrices. These Markov processes represent reduced forms for investment and aggregate demand respectively. The firm's experience state evolves endogenously according to a Markov process that depends on its own value last period and last period's production. The specification used is the one introduced by Argote, Beckman, and Epple (1990) and used in Benkard (2000), which allows for the hypothesis of organizational forgetting but encompasses both learning models:

$$E_t = \delta E_{t-1} + q_{t-1} \qquad \text{and} \quad E_1 = 1 \tag{1}$$

E_t is the firm's experience at time t, q_t is the firm's production in time t, and δ is the periodic retention rate for experience. If $\delta = 1$ then the traditional learning specification results, while

$0 \leq \delta < 1$ results in organizational forgetting. This relationship is discussed further in section 6.1.

The dynamic model consists of three stages within each period, with exit, output choice and entry, and production occurring in order. At the beginning of each period, firms simultaneously make their exit decisions. They each observe their potential scrap value, Φ_{jt}, which they receive if they choose to exit. Firms exit if their continuation EDV from remaining in the industry is less than Φ_{jt}. After the exit phase, conditional on the entrant's entry policy the remaining incumbents simultaneously make their production decisions. Given the production choices of its competitors, each firm's current production determines both its current profit and the evolution of its experience.

This model has quantity as the strategic variable for several reasons. Aircraft producers fix their production schedules a year or more in advance, and even with that lead they are constrained as to how much they can change production rates from past levels.[7] Hence, in the short term aircraft producers are clearly capacity constrained. Aircraft contracts result from complex bargaining arrangements that usually specify both quantity and price, so the commercial aircraft market is probably not well represented by either a price-setting or a quantity-setting game. However, the existence of strict capacity constraints set in advance seems to indicate that quantity is the primary strategic variable. Baldwin and Krugman (1988) also come to this conclusion.

Simultaneously with the production choices, one potential entrant observes what quality state it may enter at and the development cost it will have to pay in order to enter with that product. If the entrant pays its development cost draw, then it will enter and begin production in the next period. It is assumed that the entrant cannot produce in the period in which it enters, which is assumed to be the development period. It always enters at the lowest experience level $E = 1$.

[7] Constraints in changing production rates are partly technological, since increasing production rates can require some reorganization of the plant. However, the most important constraint is work force training. Due to the learning curve it can take a year or two for new workers to become fully productive. The length of training period significantly limits aircraft firms' ability to increase production rates.

The last stage to take place within each period is production. All incumbents that chose not to exit produce the quantities chosen in the second stage and receive profits. An individual firm's current-period profit is a function of the product qualities for all active products, the quantity produced by all active firms, its own current experience level, and current-period aggregate demand. Prices are determined by the inverse demand function while marginal cost, which is assumed to be constant within the period, is determined by the cost function. The cost function also includes a fixed cost. The demand and cost functions are discussed below.

The equilibrium concept used is symmetric Markov perfect nash equilibrium [MPE], where the strategy space includes the quantity, entry, and exit decisions. MPE, as defined by Maskin and Tirole (1988), picks out those subgame perfect equilibria where actions are a function only of payoff relevant state variables, and thus eliminates many of the vast multiplicity of subgame perfect equilibria that would normally exist in this type of model. Firms maximize their EDV of profits conditional on expectations about the evolution of present and potential future competitors. Equilibrium occurs when all firms' expectations are consistent with the process generated by the optimal policies of their competitors.

The next two subsections discuss some more technical aspects of the model including some of its theoretical properties and its computation. Some readers may at this point wish to skip to section 6, which discusses the estimation and parameterization of the model.

4.1 Value Functions:

The model as outlined above results in the following Bellman's equations for incumbent firms:

$$V(i,s,M) = \max\left\{\Phi_{it}, \sup_{q_{it}\geq 0}\left[\pi(i,s,q,M) + \right.\right. \tag{2}$$

$$\left.\left.\beta \sum V(i',s',M')\,\mathcal{P}(i',s' \mid i,s,q,M)\,\mathcal{P}(M' \mid M)\right]\right\} \tag{3}$$

where i is the firm's own state triple $(\{E_{it}, \mu_{it}, \xi_{it}\})$; s is the industry structure, a vector containing the number of firms at each possible state triple; M is the aggregate market size common to all firms; q is the vector of quantities chosen by each firm; and Φ_{it} is the firm's exit value.

Entrants' Bellman's equations are similar:

$$V^e(s, M) = \max \Big\{ 0, \tag{4}$$

$$-x_k^e + \beta \sum V(i', s', M') \mathcal{P}(i', s' \mid i_k^e, s + e(i_k^e), q, M) \mathcal{P}(M' \mid M) \Big\} \tag{5}$$

where k is the type of firm entering, x_k^e is the random entry cost; i_k^e is the entry state; and $e(i_k^e)$ is a vector of zeros with a one in the i_k^e th spot, and $q_i = 0$ (the entrant produces nothing in the entry and development period).

4.2 Computation of the Model Equilibrium

Proof that equilibrium exists for this model is straightforward and is essentially identical to the proof in Ericson and Pakes (1995) altered to include a random entry cost.[8] It is not possible to solve for the MPE of the model analytically. However, the equilibrium can be solved for numerically on the computer. In particular, Pakes and McGuire (1994, 1997) provide two computational algorithms that can be adapted to solve for the equilibria of dynamic games like this one. The primary algorithm used here is an asynchronous parallel Gauss-Seidel value iteration algorithm adapted to a game analogously to Pakes and McGuire (1994).[9]

The algorithm essentially iterates dynamic programming steps, testing for convergence at each step. When the value and policy functions do not change very much point-wise between iterations, the algorithm is assumed to have converged. The algorithm updates each firm's

[8] Available from the author upon request.

[9] Pakes and McGuire (1994) use a synchronous Gauss-Jacoby algorithm but their general approach, which is to use iterated best responses in the stage game, is the same.

value and policy functions over the entire state space at each iteration and hence must be performed conditional on a maximum number of firms, N, in the industry in order that the state space be finite. The algorithm is not guaranteed to be a contraction mapping. However, in practice the algorithm has generally converged to an equilibrium. Though non-convergence is not necessarily evidence against the existence of an equilibrium, convergence of the algorithm is sufficient for the existence of an equilibrium for a specific parameterization of the model.[10]

A more difficult problem than non-existence is that of multiple equilibria. This problem is handled in part by looking for equilibria that satisfy several "nice" properties. The most important property imposed is a weak form of symmetry, i.e., two firms that are at identical states *and* are identically situated (with the same set of competitors) are restricted to follow the same strategies. This form of symmetry is not a restriction to the model as it was described above. Rather, an equilibrium in which this weak symmetry did not hold would require additional state variables that essentially serve to label firms so that each firm knows what its competitors' policies are at every state. The symmetry assumption renders firm labels irrelevant and serves to reduce the set of equilibria.

I tested for the existence of multiple symmetric equilibria in two ways. First, attempts were made to solve for different equilibria by starting the solution algorithm from random starting locations.[11] Second, in order to test if the algorithm was somehow selecting a particular equilibrium, an entirely different computational algorithm was used to solve for the equilibrium of the model.[12] Using these two techniques, no case was identified where there was more than one symmetric equilibrium of the kind described above.

In the process of working with this model, several advances to the algorithms have been made that allow for more rapid calculation of the equilibrium. These primarily include adaptations that allow for asynchronous computation of equilibria in parallel. The model

[10] To be precise, convergence of the algorithm is sufficient for the existence of an ϵ-equilibrium.

[11] In the interest of saving time, these tests were run on a version of the model in which the computational burden was slightly reduced.

[12] The second algorithm used stochastic approximation based on Pakes and McGuire (1997).

as parameterized in section 6 currently requires 100 CPU-days to solve on a Sun Ultra 400 processor, depending on the exact parameterization and the associated convergence problems encountered. Efficient parallelization of the solution algorithm divides the run-time essentially by the number of processors used, reducing computation to a more reasonable time-frame. However, the computational burden of solving a dynamic game of this magnitude is clearly massive and cannot be overlooked. While it is the intention of this paper to show the value of the dynamic approach to economic problems of this sort, despite many theoretical and technological advances in recent years, the computational burden of the approach remains clearly its biggest obstacle.

5 Data

In estimating the parameters of the dynamic model, I rely on a wide variety of industry sources. They are discussed in order below.

5.1 Cost Data

This paper relies heavily on the data set obtained by Benkard (2000) for production of the Lockheed L-1011. This data set contains labor requirements per unit for 238 of the 250 L-1011's produced between 1970 and 1984. The data refers to direct man hours incurred by Lockheed itself in the production of each plane including detail fabrication in Burbank, Burbank assembly, Palmdale final assembly and flight test. Unfortunately, at the present time very little data (annual reports, newspaper articles etc.) is available to document the cost of other inputs to production such as capital investment and materials.

Additional data was obtained from the 1995 edition of the *Jet Airliner Production List Vol. 2* that lists each plane's model, serial number, and entire ownership history including first flight date, which is taken to be the plane's date of production.

With these two data sets, it was possible to compile the production schedule and sales pattern of the L-1011 across time and models. Please see Benkard (2000) for a more detailed discussion of the L-1011 labor cost data.

The remaining L-1011 cost data was gleaned from Lockheed's annual reports, which contain very detailed data on yearly development costs as well as initial tooling. Lockheed also reports "Standby Production Costs" in years of low production, which was used to estimate Lockheed's fixed costs in production of the L-1011.

5.2 Demand Data

The annual fleet and deliveries data used in the demand estimation comes from Boeing's *World Jet-Airplane Inventory Year-End 1993* with supplemental information for 1994 that came from the continuation of that publication published by Jet Information Services. Aircraft characteristics come from various years of *Jane's All the World's Aircraft*.

Key to estimating the demand system is relating prices to the exact set of characteristics for each plane. Since individual planes differed significantly with respect to equipment, for each of the eight wide-bodied planes offered over the time period 1969-1994 I chose the yearly modal plane in terms of model and equipment (primarily engine type). Then, using a transaction price dataset provided by Avmark and the characteristics data from Jane's, I constructed a series of modal characteristics matched to the average price for that set of characteristics for each plane in each year.

6 Parameterization of the Model

This section discusses the estimation and parameterization of the model using industry data. Where possible, all parameters were estimated econometrically. The remaining parameters are matched to observed data. Units are 1994 million dollars throughout.

6.1 Cost Function

A goal of this paper is to match the industry model to the commercial aircraft industry as closely as possible. The cost structure of the industry drives much of the interesting strategic behavior that we observe among aircraft producers. Therefore, in specifying the cost function for the industry model it was critical to work with data that came directly from the *commercial* aircraft industry, rather than using more widely available military production data. Much of the work toward that end was accomplished in a previous paper (Benkard (2000)). I do not go into as great detail here, but rather direct readers to that paper in the event that they desire further clarification or technical detail.

Benkard (2000) lists the assumptions needed to derive the following labor requirements equation:

$$\ln L_{it} = \ln A + \theta \ln E_t + \gamma \ln S_t + \epsilon_{it} \tag{6}$$

where L_{it} is the labor input per-unit; A is a constant; E_t is experience; ϵ_{it} is a plane-specific productivity shock; and S_t is line-speed, a measure of the current production rate.

Many authors have estimated similar learning curve specifications to (6) using data from countless other industries and defining experience as cumulative production ($E_t = \sum_{i=0}^{t} q_t$). I refer to this specification as the traditional learning curve because a similar specification was originally applied by Wright (1936) in the first published paper recognizing the existence of learning curves. The main contribution of Benkard (2000) is to show that, due to high variance in output rates for commercial aircraft production, the traditional learning curve does not explain costs for commercial producers particularly well. However, a similar learning model that incorporates the hypothesis of organizational forgetting fits the data extremely well, while simultaneously providing a very satisfying economically intuition as to why this might be.

In the organizational forgetting model, experience evolves as follows. At time $t-1$ a firm has a stock of experience E_{t-1}. The firm then chooses its current production rate q_{t-1}. Between periods $t-1$ and t, the firm's existing stock of experience depreciates by a factor δ, while new experience equal to q_{t-1} is acquired. This process is summarized by the following equation:

$$E_t = \delta E_{t-1} + q_{t-1} \qquad \text{and} \quad E_1 = 1 \tag{7}$$

The specification described in (7) is also very intuitive. Production experience in the aircraft industry is embodied in the actual workers. It refers to the workers' ability to perform their tasks efficiently. Hence, an aircraft producer's stock of production experience is constantly being eroded by turnover, layoffs, and simple losses of proficiency at seldom repeated tasks. When producers cut back output, this erosion can even outpace new learning, causing the stock of experience to decrease, as was the case in the dataset presented in Benkard (2000). In that event, production costs will rise. To reduce costs back to their previous levels, producers must maintain higher production rates for a long enough period that experience gains outweigh declines, and the former experience level is reached again. Another intuition exactly analogous to the depreciation story of equation (7) is that recent past production should be more important in determining current production costs than distant past production. For a plane like the 747 that has been produced for almost three decades, it is hard to imagine that production in the early 1970's is much of a factor in current production costs. In the traditional learning model, all production experience is treated equally regardless of how old it is. Please see Benkard (2000) for a complete discussion of the organizational forgetting model in general as well as in the context of aircraft production.

6.1.1 Cost Parameters

Figure 2, derived from the results in Benkard (2000), shows the actual labor requirements of the Lockheed L-1011 versus the fitted equation (6).[13] Note that the data does not monotonically decrease as $\delta = 1$ would imply, but instead contains two turning points. As discussed in

[13] The labor requirements equation is estimated using GMM with a non-parametric heteroskedasticity and autocorrelation consistent (HAC) estimator for the weighting matrix. This technique was developed by Andrews (1991) and is described in detail in Benkard (2000). The fitted equation shown in Figure 2 also accounts

Benkard (2000), the organizational forgetting model captures these two turning points almost perfectly.

This near-perfect fit of the cost system is very important to the overall dynamic model since it is believed that the learning/forgetting dynamics are the most important factor driving strategic interaction in the aircraft industry. Thus, capturing this aspect of the problem so well should be expected to translate into better overall results.

Figure 2: Organizational Forgetting Regression: IV with HAC Weighting Matrix

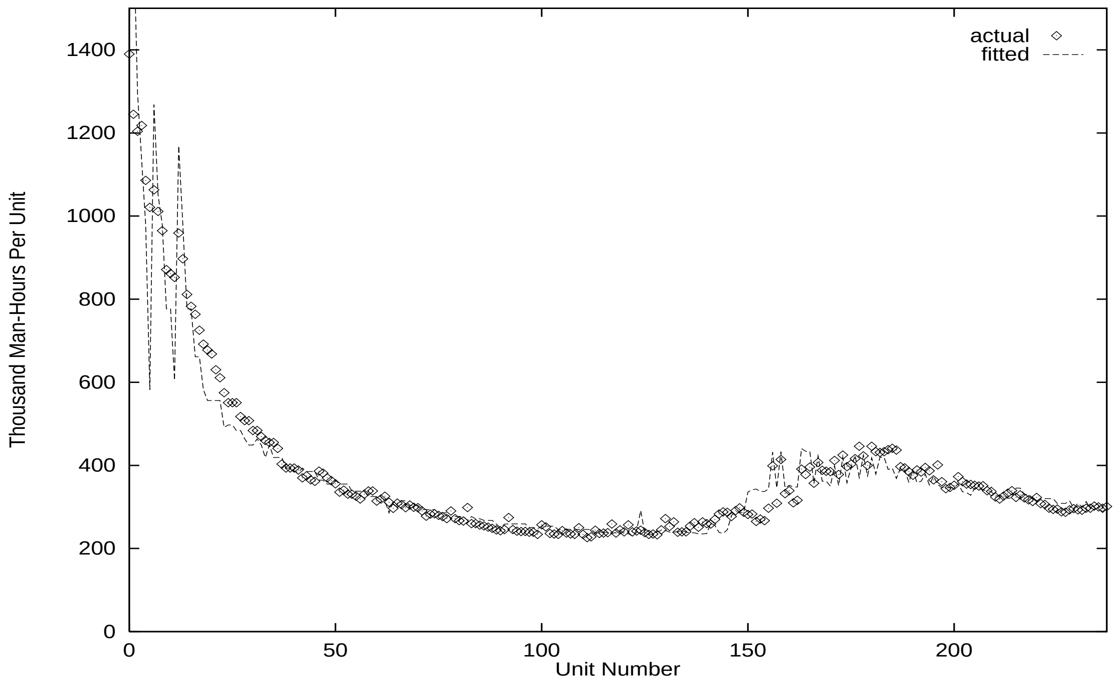

The monthly retention rate of experience (δ) is estimated in Benkard (2000) to be 0.960 with a standard error of 0.003, corresponding to the annual rate of 0.613 listed in Table 1. The hypothesis that $\delta = 1$ is therefore strongly rejected. The implication here is that an aircraft producer loses about 40% of its previous stock of experience every year. This number may at first seem high, since the forgetting is believed to be driven by labor turnover and company-level turnover rates do not approach this level. However, there is evidence to suggest that position-level turnover rates may in fact be this high due to a unique contract structure with

for the fact that there was more than one model of L-1011 produced by allowing for incomplete spill-overs of production experience between models.

the aircraft workers union (IAM). See Benkard (2000) for an elaboration.

Table 1: Cost Parameters

Parameter	Explanation	Value
δ	Depreciation of Experience	0.613 (0.023)
A	Labor Cost Intercept	7.73 (0.01)
θ	Learning Parameter	-0.63 (0.03)
	(Implied Learning Rate)	36%
W	Wage Rate	\$20 / hr
FC	Fixed Costs	\$200 Mil / yr
TCF	Total Variable Cost / Labor Cost	6.0
TCC	Total Variable Cost Intercept	36.2
	Cost/Plane-Size Ratio	1.0
x_1^l	Lowest Entry Cost for Type 1	\$2.5 Billion
x_1^h	Highest Entry Cost for Type 1	\$3.5 Billion

The implied learning rate with respect to experience is quite rapid at 36%. Note, however, that the interpretation of the learning rate in this model differs from that of the traditional learning model ($\delta = 1$) since production rates also matter. The learning rate implies that if experience were doubled, then labor requirements would fall by 36%. Whether or not this reduction is attainable depends on the current experience level and future production rates.

Since the cost model is estimated on data for the L-1011, some further assumptions were necessary to calculate costs for other plane types. Specifically, production cost is assumed to depend on the two quality state variables for each firm. For reference, the quality states are discussed in more detail in the following section.

The unobservable quality state (the unobservable product characteristic ξ_{jt} as estimated in the demand system — see below) was included in the cost regressions and, despite its having quite large variance in the sample, it was found not to affect production cost, so that assumption will be maintained in the cost function. Note that the unobservable portion of product quality may be most representative of characteristics like suitability to the current airline route network, which would not necessarily influence marginal cost.

Unfortunately it was not possible to estimate the impact of the observable quality state (the observable product characteristics μ_{jt} as estimated in the demand system) on cost since the cost data only covers one plane and, as is typical, the observable characteristics had very little variance in the data for this plane. However, based on materials inputs alone it is quite obvious that larger planes have higher costs in relation to their size. Therefore it seems logical to assume that the variable cost of a larger plane is greater than that of an L-1011 in exact proportion to its relative size.[14]

The wage rate was set equal to the wage rate for aircraft workers in 1994 and then total variable cost from Lockheed's annual reports was regressed on total labor costs to obtain the two variable cost parameters. Fixed costs are assumed to be constant across plane types and were also estimated using annual reports. Development (entry) costs are assumed to scale up with size similarly to variable cost, and were based on Lockheed's development costs.[15]

6.2 Demand System

Demand for commercial aircraft is very complex, and estimating the demand for aircraft is a formidable research agenda in itself. Therefore, in modeling aircraft demand the goal was to find a model that is theoretically appealing and fit the data well, without adding greatly to the computational burden of the dynamic model.

The most important feature of aircraft demand that differentiates it from standard discrete

[14] Here I measure the size ratio as the average between the ratio of seats and the ratio of volume. Volume ought to give a good measure of relative capital and materials inputs, but seats should be a better measure of relative labor input.

[15] Lockheed developed the L-1011 at a cost of $2.52 billion. However, industry sources agree that the biggest structural change in the aircraft industry in the last 20 years has been the escalation of development costs. Hence, the Lockheed figure was chosen as a lower bound for the entry cost distribution.

Note also that modeling entry costs as similar in magnitude to those experienced by a current aircraft producer amounts to assuming that potential entrants are current aircraft producers. This seems like a reasonable assumption in view of the fact that the only outside entrant into the industry in the last thirty years was the Airbus consortium, which was funded by several European governments at great cost.

choice frameworks is that aircraft are durable goods. Therefore, rather than modeling aircraft purchases as occurring only at the time of a transaction, I assume that each airline optimally reallocates its entire aircraft fleet each year, choosing from all available new and used planes at the going market prices. This assumption, which amounts to treating aircraft purchases as rentals, relies on the fact that the market for used commercial aircraft is very efficient so there is little cost in conducting a transaction.[16]

I model yearly aircraft demand using a standard characteristics based approach. I use a nested logit discrete choice model with several observable characteristics (number of seats, number of engines, etc) and one unobserved characteristic, similarly to Berry (1994). The unobserved product characteristic represents the unobservable aspects of an aircraft's quality, such as reliability, or suitability to current route structures, and is estimated using the data. By construction, with the addition of the unobserved product characteristic the demand model fits the data exactly.

There are two groups (nests) in the model, one that includes all new planes in the market, and one that includes only the outside good, which is defined to be all new narrow-bodied jet planes and all used jet planes. Individual aircraft purchases are assumed to be independent decisions even if undertaken by the same airline. This assumption is not likely to hold. However, relaxing it has proven to be quite difficult and the literature on multiple discrete choice is sparse, so I maintain that assumption here for lack of a better alternative.

The nested logit model is a great improvement over the standard logit model because it allows for the estimation of a "within group" correlation of utilities. The implication here is that an airline's preference for each of the new wide-bodied planes is correlated, and that this correlation will be estimated. This correlation allows for more reasonable substitution patterns than the standard logit model because inside goods (new wide-bodies) are not constrained to substitute with the outside good in relation to its share as they are in the standard logit model. Indeed, this feature of the model was found to be important in fitting the data well.

[16] Alfred Kahn, former Chairman of the Civil Aeronautics Board and architect of airline deregulation in the 1980s, once referred to aircraft as "...nothing but a marginal cost with wings."

Airlines' utility functions thus look as follows:

$$u_{ijt} = x_{jt}\beta - \alpha p_{jt} + \xi_{jt} + \zeta_{igt} + \lambda\epsilon_{ijt} \tag{8}$$

where x_{jt} are observed characteristics of product j in period t, ξ_{jt} is an unobserved characteristic of j, ζ_{igt} and ϵ_{ijt} are the random group- and plane-specific tastes respectively, and $0 \leq \lambda \leq 1$ is a parameter representing the within group correlation of utilities.

Solving for the aggregate market shares and inverting gives:

$$\xi_{jt} = \ln(s_{jt}) - \ln(s_0) - x_{jt}\beta + \alpha p_{jt} - (1-\lambda)\ln(s_{jt/g}) \tag{9}$$

where s_{jt} is the overall share of good j, s_0 is the share of the outside good, and $s_{jt/g}$ is the within-group share of j. From this equation it is easy to see that the within group correlation of utilities (λ) is identified by covariation between the within group market share of the good ($s_{jt/g}$) and its total market share (s_{jt}).

Utilizing the following moment condition assumption:

$$E\left[\xi_{jt} \mid Z_{jt}, \theta_0\right] = 0 \tag{10}$$

for an appropriate set of instruments Z_{jt}, and similar assumptions and procedures to those used in Benkard (2000)[17], consistent estimates of the parameter vector θ are obtained.

Instruments used include plane characteristics (and model dummies), the hourly wage in manufacturing, the number of years a model has been on the market (to proxy learning while maintaining uncorrelatedness with ξ_j), price of aluminum, and dummies for the MD-11 in 1990 and the A310 in 1994, which were years of supply disruptions for those models.

[17]GMM with an optimal weight matrix.

Table 2: Demand Function Estimates

Variable	Estimate	S.E.	HAC
Constant	-4.81	0.16	0.15
Seats/100	1.10	0.21	0.23
Freighter	2.45	0.24	0.26
# Engines	-0.30	0.53	0.46
Price/100	-2.40	0.21	0.30
MD-11	-0.35	0.49	0.35
L-1011	-0.14	1.39	1.01
B-747	-0.40	0.49	0.42
B-767	-0.61	0.53	0.53
A-300	-0.91	0.34	0.32
A-310	-0.40	0.78	0.71
Last Year Dummy	-0.90	0.37	0.38
Trend	0.25	0.43	0.58
λ	0.23	0.61	0.63

6.2.1 Demand System Estimates

The demand system was estimated for the period 1975-1994.[18] A total of eight models are observed over the estimation period, leading to 98 model-year observations. Parameter estimates are shown in Table 2. The column labeled "HAC" refers to heteroskedasticity and autocorrelation consistent standard errors. Robust standard errors were calculated in case the unobserved product characteristics were serially correlated or correlated with each other. For the most part the coefficient estimates are as expected. "Number of Engines" is a proxy for fuel efficiency since, given plane size, more engines create more drag. "Last Year Dummy" is a dummy that is one in the last year that a plane was sold.

Only two of the estimated parameters, the within group correlation of utilities (λ) and the price coefficient (α), are relevant to the dynamic model since the remaining parameters are aggregated into the two product quality states. The parameter λ is estimated to be close to zero, which means that the within group correlation of utilities is high for this market.

[18]Parameter estimates including the period 1969-1974 were similar. However, the market was immature and contained very few products in its first five years. The first wide-body available was the 747 which was introduced in 1969. It was followed by the DC-10, the L-1011, and the A-300, all of which were available by 1975.

The implication is that new wide-bodies substitute much more highly with each other than they do with other types of aircraft. This finding is quite consistent with intuition. Airlines often play producers against each other in bargaining, seeming to care as much about the deal reached as which plane they end up purchasing. In addition, with the exception of the 747 wide-bodies are quite highly substitutable with one another in use. The relatively high standard error on λ may be caused in part by the changing nature of the wide-body market over time as it matured. However, the estimated value of this parameter was found to be robust to alternative specifications, which included changing the time-period of estimation.[19]

On the other hand, the coefficient on price (α), is estimated very precisely. Taken together, the two parameters lead to own price elasticities in the 5-13 range for 1994, which seems appropriate for this market given the above discussion. These elasticities are also consistent with Newhouse's (1982) anecdotal accounts of the industry.

6.2.2 Stochastic Processes for Quality

The detailed plane characteristics used in estimation of the demand parameters are collapsed down to two dimensions for the purposes of the industry model. Those dimensions are plane type and plane quality. Table 3 summarizes the estimated Markov processes for plane type and quality.

The first dimension, aircraft type (μ_{jt}), corresponds to the observable characteristics of the plane ($x_{jt}\beta$ previously). In the dynamic model a plane retains the same type for its lifetime. This assumption simply rules out Boeing turning a 747 into some other type of plane (e.g. DC-10). Within the context of the model such a large overhaul of the plane's design would be treated instead as a new product introduction. This assumption also reflects actual practice as none of the aircraft in the data have undergone significant changes in size. I allow for three types of planes in the model, corresponding to the three levels for μ_{jt} in the table. These three types can be thought of as small (L-1011, A300, etc.), medium (MD-11, A330, etc.),

[19] In all the specifications estimated the standard error on λ remained high.

Table 3: Demand and Other Parameters

Parameter	Explanation	Value
λ	Group Corr. Parameter	0.234 (0.613)
α	Price Coefficient	-0.024 (0.002)
μ	Discrete Plane Types (Small, Medium, Large)	$\left\{ -2.6, \quad -2.2, \quad -1.6 \right\}$
$P(\mu^e)$	Entry Type Distribution (Small, Medium, Large)	$\begin{pmatrix} 0.50 & 0.38 & 0.12 \end{pmatrix}$
ξ	Discrete Plane Qualities	$\left\{ -0.90, \quad -0.40, \quad 0.11 \quad 0.61 \right\}$
$\Delta\xi$	Transition Matrix for Quality	$\begin{pmatrix} 1.00 & 0.04 & 0.033 & 0.000 \\ 0.00 & 0.44 & 0.233 & 0.200 \\ 0.00 & 0.48 & 0.667 & 0.800 \\ 0.00 & 0.04 & 0.067 & 0.000 \end{pmatrix}$
M	Discrete Market Sizes	$\begin{pmatrix} 10339 & 10929 & 11519 \end{pmatrix}$
ΔM	Transition Matrix for Market Size	$\begin{pmatrix} 0.895 & 0.143 & 0.000 \\ 0.105 & 0.786 & 0.200 \\ 0.000 & 0.071 & 0.800 \end{pmatrix}$
β	Firm's Discount Rate	0.925
(Φ^l, Φ^h)	Range of Scrap Values	$(\$300M, \$700M)$

and large (747).

The second dimension, plane quality (ξ_{jt}), corresponds to the unobservable characteristic in the demand system. The plane's quality moves according to a discrete Markov process which was estimated nonparametrically using the estimated values of ξ_{jt} from the demand system. In the data (and also in the model), it is this unobserved characteristic that accounts for most of the year-to-year variance in product quality, since major changes in an aircraft's characteristics occur infrequently. The estimated Markov process should be viewed as a reduced form for the outcomes associated with the firm's product-level investment process.

Finally, planes are assumed to enter at the second highest quality level because that corresponded to the observed entry value for seven of the eight planes in the sample. In addition, the empirical distribution of plane entry types is used as the entry type distribution, listed in the table as $P(\mu^e)$.

6.3 Stochastic Process for Aggregate Demand

In order to reduce the complexity of the problem, and because the steady growth in market size was deemed of second order importance relative to business cycle fluctuations, the market size state variable is de-trended to reflect 1994 values. This has the effect of making all the state variables of the model finite and stationary, which facilitates computation of the model equilibria,[20] but retains business cycle fluctuations in the form of booms and busts. This market evolution adds interesting dynamics in the organizational forgetting case, where an extended recession can result in significant productivity losses.

The Markov process for demand fluctuations (shown in Table 3) was discretized to three points and estimated nonparametrically using market size data for the complete history of the commercial jet aircraft industry (1956-1994).

6.4 Other Parameters

Table 3 also lists values for the two remaining parameters of the model. The firm's discount rate, β, was set to 0.925, which corresponds to a standard annual interest rate.[21]

The scrap value of a production facility is inherently very difficult to measure. However,

[20] An alternative suggested by several seminar participants would be to allow some growth in the market that would eventually cease. I have solved for versions of the model with this feature, but found that there were no qualitative differences in the results. Thus, in the interest of keeping computational burden to a minimum this feature was taken out of the model.

[21] Changes in this value within a reasonable range did not result in significant changes to the model results. No attempt was made to estimate this parameter as past results suggest that it is typically not identified in the data.

Lockheed does report significant detail on setup costs for the L-1011 and, in particular, the "initial tooling" portion of L-1011 development costs was about $1.0 billion. Since much of this initial tooling is design-specific, and since the scrap value should vary across firms depending among other things on whether the firm was going to continue producing other kinds of aircraft,[22] the scrap value distribution was chosen to be close to $500 million.

7 Results: Properties of the Equilibrium and Comparison with Historical Data

The results presented in this section are taken from the symmetric MPE of the dynamic model with the industry restricted to a maximum of four single product firms. Ideally this restriction would have been relaxed to the point that it was no longer binding. However, computational considerations precluded that. As above, the industry structure in each period is summarized by aggregate demand (M_t), plane type for each active firm (μ_{jt}), unobserved quality for each active firm (ξ_{jt}), and experience for each active firm (E_{jt}). All states are discretized for the purposes of computing the model equilibrium. See section A for a discussion of the discretization of the experience state. The model contains a total of 13 state variables and approximately seven million states.

7.1 Pricing Policies in Equilibrium

7.1.1 Introductory Pricing

Figure 3 graphs the equilibrium price-cost ratios for a newly introduced small (L-1011 sized) plane with three equal rivals. According to the model, in every state in which a new product is introduced, introductory pricing is at a level below static marginal cost. There are theoretical models in the literature that predict below static marginal cost pricing, but to my knowledge

[22] Most former commercial jet producers have continued to produce military or other smaller commercial planes.

this is the first empirical model to successfully capture this fundamental feature of aircraft pricing.

Figure 3: Introductory P/MC Ratios For a Small Aircraft Entrant with Three Equal Rivals

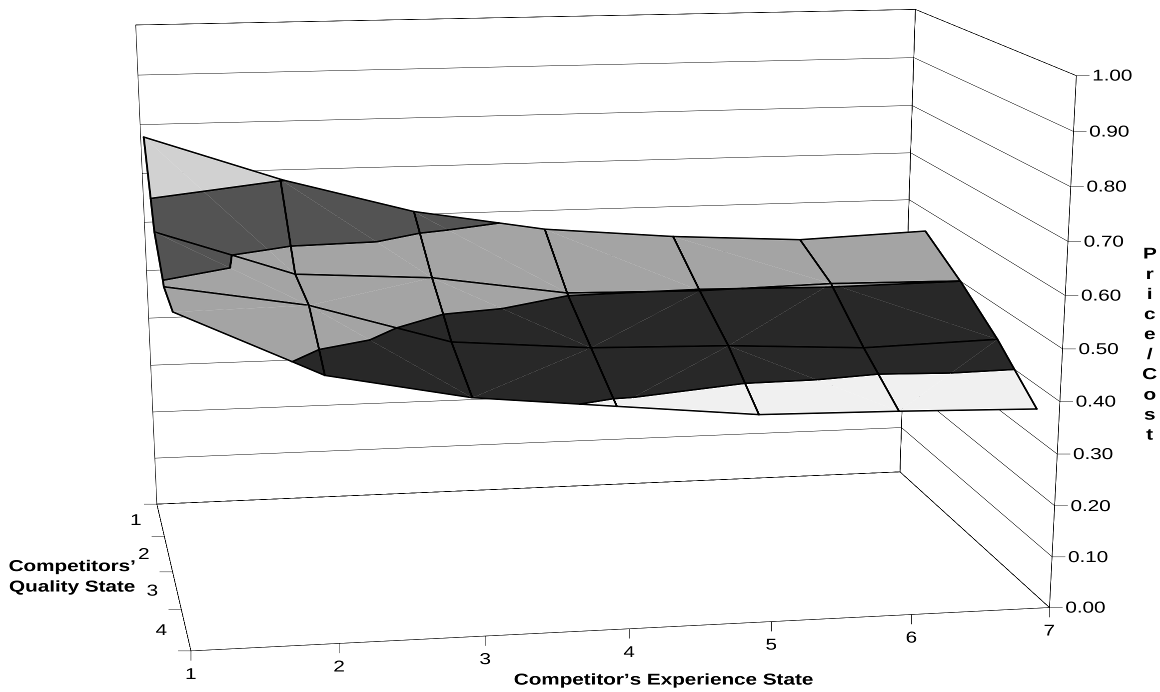

Qualitatively, the introductory price-cost ratios predicted by the model match the industry well. The predicted price-cost ratios cover a wide range (0.33-0.79), which shows that pricing depends critically on the nature of the competition. The model predicts introductory price-cost ratios that are typically lower in three cases (cases 2 and 3 can be seen in Figure 3): 1) when there are more competitors in the market, 2) when incumbent products are higher quality, 3) when incumbent firms are further down their learning curves. The strongest of the three effects is the learning curve. The model sometimes predicts high markups in states where there are many high-quality competitors, if the competitors also have high cost, but always predicts low markups when there is even one low-cost competitor.

In the past we have only observed entry of new products under conditions where there was relatively strong competition, so to the extent that introductory price-cost ratios are observable, they have generally been quite low. When the L-1011 entered the market there were two competitors, the Boeing 747 and the McDonnell-Douglas DC-10. At this state, the model predicts a price-cost ratio of 0.49. The actual observed price-cost ratio for the L-1011 in 1972 was very close to this level at 0.48.

7.1.2 Markups

As a measure of how well the model is capturing industry pricing behavior I now compare the equilibrium pricing policies predicted by the dynamic model with those observed for the L-1011. This is an extremely rigorous test. The historical price data was used to estimate the parameters of the demand system but, aside from these parameters, the model contains no direct information about prices. In the model, prices are generated endogenously through the equilibrium assumption. Additionally, because demand and supply were estimated separately, all parameters were estimated without using any information about markups that is contained in the data. The prices and markups in the model equilibrium are thus generated largely by the structure of the dynamic model. Note also that observed prices have been shown to be very different from contemporaneous marginal cost, so the near perfect fit of the cost system does not in any way guarantee that prices will be predicted well.

In order to make the comparison it was first necessary to calculate the closest discretized industry structures to those that actually occurred, a simple task given the parameter estimates and the observed data. Figure 4 graphs observed price-cost ratios for the L-1011 against equilibrium price-cost ratios for a small (L-1011 sized) wide-body in the model located at the industry structures actually observed from 1972-1985.[23]

Generally, the price-cost ratios predicted by the model are quite similar to those observed. The two series are very similar in both overall shape and year-to-year variation, and the model correctly predicts negative markups for most of the period. These results suggest that the dynamic equilibrium assumption is doing much to capture the influence of intertemporal

[23]Computation of the model equilibrium was restricted to four firms. Thus, in calculating the predicted prices and price-cost ratios for states with more than four firms, I made the assumption that each firm cares only about its three strongest competitors.

maximization on current prices. It should also be noted that, while some members of the press have characterized Lockheed's persistent below static marginal cost pricing as irrational, this model seems to suggest otherwise.

Figure 4: Predicted vs Actual Price/Cost Ratio for L-1011: 1972-1985

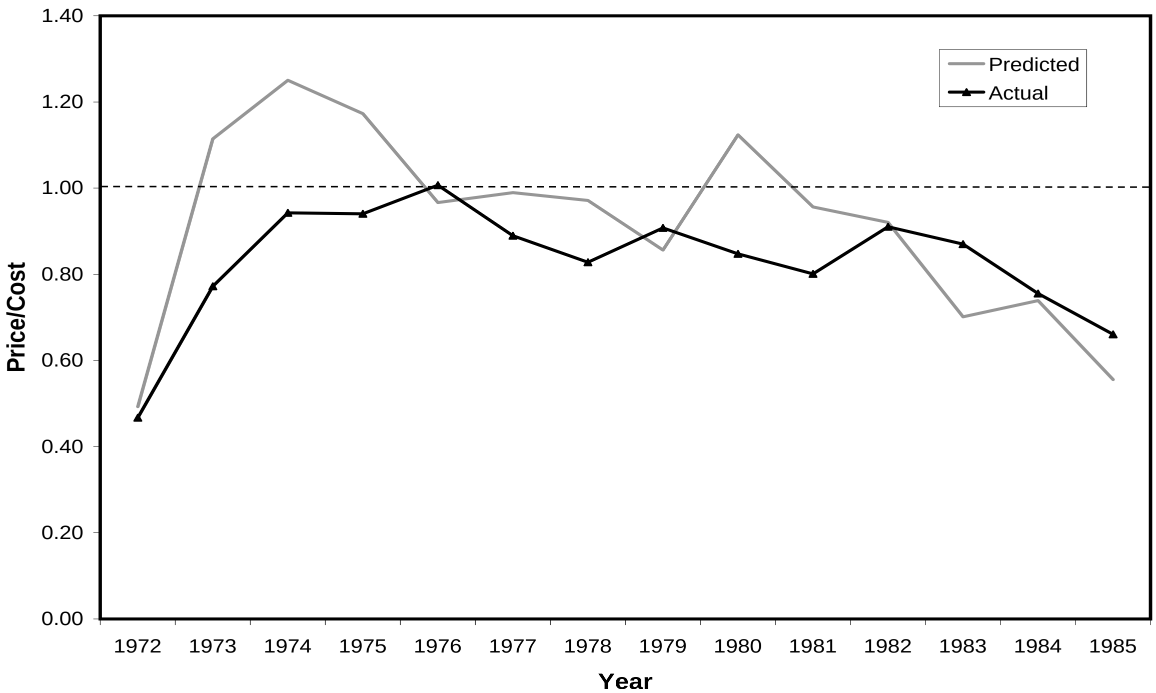

The most notable discrepancy (see also Table 4) between the two series occurs in the period immediately after the L-1011's introduction (1973-1975), where the model predicts higher markups and prices than those observed. I believe that this overprediction is largely attributable to the use of a simplified demand system. The demand system underestimates the high degree of competition that took place between the L-1011 and DC-10. In the nested logit model substitution between products is based on the inside share of the products without separately accounting for the proximity of products in characteristic space. An alternative interpretation is that Lockheed was pricing at a level that was lower than optimal as suggested by the dynamic model.[24]

[24]Recall that consistency of the estimation of the model does not require optimizing behavior on the part of firms.

Markups and prices over the rest of the period (1976-1984) are on average correct and fall within the observed range of prices in almost every year (see Table 4). While markups are not observable for other aircraft, comparisons of predicted versus actual prices for other aircraft are universally similar.

Table 4: Predicted L-1011 Prices and Observed Price Range

Year	Predicted Average	Observed Average Modal	Observed Min	Observed Max
1972	99.3	62.6	59.4	64.7
1973	82.5	64.0	58.3	77.1
1974	75.8	60.9	52.9	76.8
1975	71.1	57.2	54.9	58.0
1976	58.6	62.0	55.2	73.6
1977	59.7	57.6	56.9	59.2
1978	58.9	55.5	55.0	63.1
1979	63.4	67.4	42.9	66.5
1980	83.2	67.0	50.4	82.7
1981	70.8	57.6	57.0	86.4
1982	55.8	62.1	57.5	63.2
1983	51.9	63.5*	NA	NA
1984	54.7	64.2*	NA	NA
1985	NA**	65.4*	54.0	54.0

* Estimated Sales-Weighted Prices
** Model Predicts Exit

In the past it has been difficult to come up with a model that explains aircraft industry pricing policies well. Note that any static model would necessarily predict positive markups and thus overpredict prices and markups for the L-1011 in every period. The only previous attempt that I know of at modeling the aircraft industry using a dynamic model is Baldwin and Krugman (1988). Their primary intent was to evaluate the merits of strategic trade policy, however the pricing policies in a precommitment equilibrium in their model did not reflect observed prices very well. Thus, despite many simplifications in the model, the equilibrium pricing policies predicted by the model are much closer to observed policies than previous models in the literature. I consider this to be one important contribution of this paper.

Furthermore, the model predicts that the variance in price is much lower than the variance in cost, and that there is widespread below static marginal cost pricing. Both of these features were shown in section 2 to be present in the observed data, and I know of no other empirical model in the literature that has replicated them.

7.2 Industry Dynamics

The dynamic model also replicates observed industry dynamics well in many respects. Table 5 shows actual (1969-1994) and simulated statistics for the wide-body market. Two simulations were used to calculate these statistics, one which reflects the initial condition for the industry in 1969 and one which reflects the long run invariant distribution. The first simulation ("I.C. Simulation") shows statistics from 1000 26-period simulations of the dynamic model with initial condition equal to the actual initial state of the industry in 1969, i.e., one large (747-sized) plane. The model generates an ergodic Markov process of industry states, so for long enough simulations the initial condition is irrelevant. However the observed data corresponds to a certain initial condition and this condition is likely to affect industry dynamics in the short run. I do not believe that 26 periods is long enough to exhaust the memory of the process, so I believe that the initial condition simulation is a better point of comparison for the observed data than the long run invariant distribution.

However, there are also some statistics of interest which are difficult to collect with any accuracy from such a short simulation period. For example, since many of the firms that entered in this period have not exited yet, it would be difficult to compile statistics for firm value and lifetime distributions without a longer simulation. Thus, for the firm lifetime and value distributions a much longer simulation of 10000 periods ("Invariant Distribution Simulation") was used. The statistics collected reflect the unique invariant distribution of states, so the initial condition for the second simulation is irrelevant.

7.2.1 Concentration and Market Size

In both the observed data and the initial condition simulations there are initially few firms in the market, so initial concentration is very high. Then, as more firms enter, market concentration falls and stabilizes at approximately the levels represented by the invariant distribution simulation. The simulated one- and two-plane concentration ratios from the model thus appear to closely match the observed ratios. Firms in the model are single-product producers (or, equivalently, firms are setting prices independently across products), so the model does not make predictions about firm-level concentration ratios.

The total market size distribution generated by the model is slightly smaller than that observed, most probably reflecting the artificial dimensional restriction in the model used to limit the computational burden of the problem. Concentration ratios and market size are closely matched in distribution as well as in mean, implying that the model is also doing quite well at replicating the underlying stochastic process of industry states.

7.2.2 Product Type, Value and Lifetime Distributions

Table 5 also lists observed and simulated plane type distributions. The distribution of plane types generated by the initial conditions simulation is very close to the observed distribution. The high percentage of large planes over the historical period reflects the early entry and continued market participation of the 747. This feature of the data is captured nearly perfectly by the initial condition simulation. However, of the three product types, large planes are also the least likely to enter. Thus, according to the invariant distribution simulations we should expect to see a market made up of more small planes and fewer large planes in the future. If the present is any reflection of the future, that prediction seems correct. There have been many entrants in the small and mid-sized wide-body classes, but as yet the 747 has no competition in its class.

Note that these results are in part driven by the fact that the entry type distribution was

Table 5: Model Simulations and Historical Industry Characteristics 1969-1994

Concentration Ratios:	*Observed*	*I.C. Simulation*	*Invariant Distribution*
1-Plane	0.44	0.47	0.40
S.D.	0.20	0.17	0.10
2-Plane	0.68	0.75	0.69
S.D.	0.14	0.13	0.11
1-Firm	0.55	–	–
S.D.	0.17		
2-Firm	0.82	–	–
S.D.	0.12		
Market Size:			
# Planes	4.4	3.5	3.8
S.D.	1.2	0.8	0.4
# Firms	3.4	–	–
S.D.	0.7		

Distribution of Plane Types:	*Observed*	*I.C. Simulation*	*Invariant Distribution*
Small	0.56	0.54	0.72
Medium	0.23	0.22	0.23
Large	0.21	0.24	0.05

Table 6: Invariant Distribution of Plane Values and Lifetimes

Distribution of Plane Values: (*Invariant Distribution Only*)			
Median	1027	Min	-8593
Mean	832.6	Max	15756
S.D.	4000	% Positive	54.5

Distribution of Plane Lifetimes: (*Invariant Distribution Only*)			
Median	21	Min	2
Mean	30.4	Max	261
S.D.	29.6		

parameterized to exactly match the observed distribution of entry types, a modeling convenience designed to limit the number of potential entrants each period to one. However, since entrants can choose whether or not to enter given their draw on product type, the distribution of products generated by the model remains fully endogenous, and should reflect the relative profitability of each product type rather than the parameterized entry distribution. In fact, the invariant distribution of plane types is quite different from the parameterized entry distribution (see Table 3 for comparison), and the invariant distribution reflects the observed data quite closely, while the parameterized entry distribution does not.

Due to incomplete disclosure and the short history of the industry it would be very difficult to calculate observed plane values.[25] However, the distribution of values generated by the model has several features that qualitatively match the industry. High variance is one feature that is predicted by the model that is most definitely present in the industry. Some planes lose a great deal of money, while others are very successful. On the other hand, the model predicts that a slim majority (55%) of planes in the market are profitable, which at least two (Newhouse (1982), Seitz and Steele (1985)) authors have claimed is not the case, though both of these publications are now more than a decade old and there is evidence that more products have been profitable recently than were in the past. The predicted median value of $1027 Million is within a reasonable range, but because program-level data is highly guarded there is no corresponding observable to compare it to. The model also predicts that the value distribution has a thick right tail, which seems to reflect observation. There are a few planes, e.g., the Boeing 747, that have been extremely profitable.

Product lifetimes in the invariant distribution are very left skewed, with a median of 21 years. Again there is no corresponding observable to compare this figure to, but based on observation to date and our knowledge of the industry, the lifetime distribution seems reasonable. Of course with only 26 years of history and only two aircraft that have exited to date there is little information in the data that would help to identify the right tail of the distribution, so we should not expect the model simulations to be too accurate there.

[25]Most companies do not give any public accounting of development costs, making it very difficult to determine whether or not the plane broke even in the long run. Furthermore, the majority of wide-bodied aircraft are still being produced today.

The product lifetime distribution is driven by the equilibrium exit policy function as well as the Markov process of industry states. The equilibrium exit policies generated by the model are also quite consistent with the observed history. For example, while the L-1011 actually exited in 1986 (defining exit consistently with the dynamic model, 1986 was the first year in which zero L-1011's were delivered so it is the year in which Lockheed exited rather than produce), the model suggests that it would have been optimal for the L-1011 to have exited in 1985. However, the distinction between the two years is essentially a technical one. By 1985 Lockheed had ceased to produce the L-1011. The two aircraft sales that were made reflected unsold inventory from the previous year. The model also suggests that the DC-10 should have technically exited one year prior to its actual exit in 1990, but again in this case the one aircraft sale that took place in 1989 reflected the remaining inventory from the year prior.

8 Representative Twenty-Year Simulation

This section will use a typical twenty period industry simulation to display several important features of the model simulations. Figures 5-8 describe a typical 20 year model simulation with initial condition as above, i.e., a market with one large 747-style plane. During this period, five firms are observed, the initial large plane (firm 1), three small plane entrants, and one medium sized entrant (firm 4). Firm 3, which is a small sized plane, enters in period 4 and exits in period 8. The remaining firms remain active at the end of the simulation period.

This simulation shows three major points. The first is that, according to the model, prices generally do not reflect costs. Cost curves follow a standard looking learning curve (despite the presence of forgetting in the cost function) and thus cost varies over a wide range. Prices, on the other hand, are relatively constant to changes in the market. The first three entrants have slightly higher initial prices due to the fact that they have fewer rival firms, and more importantly, no rivals that have reached the bottom of their learning curves. However, once

Figure 5: **Twenty-Year Simulation: Prices**

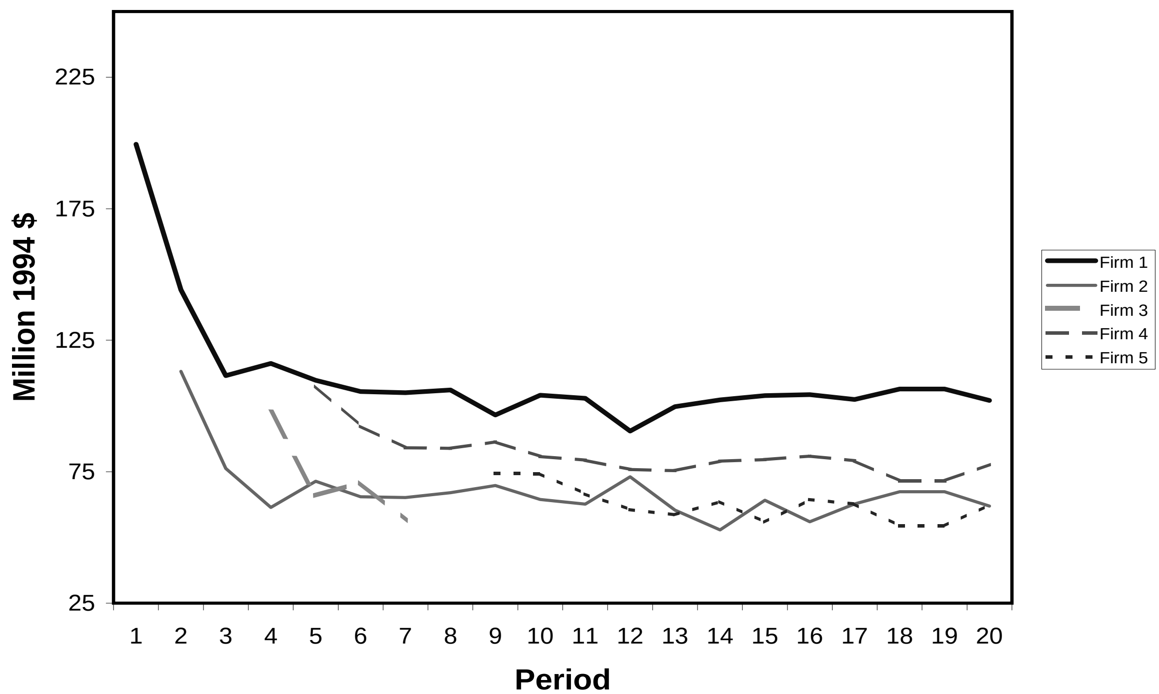

Figure 6: **Twenty-Year Simulation: Cost Curves**

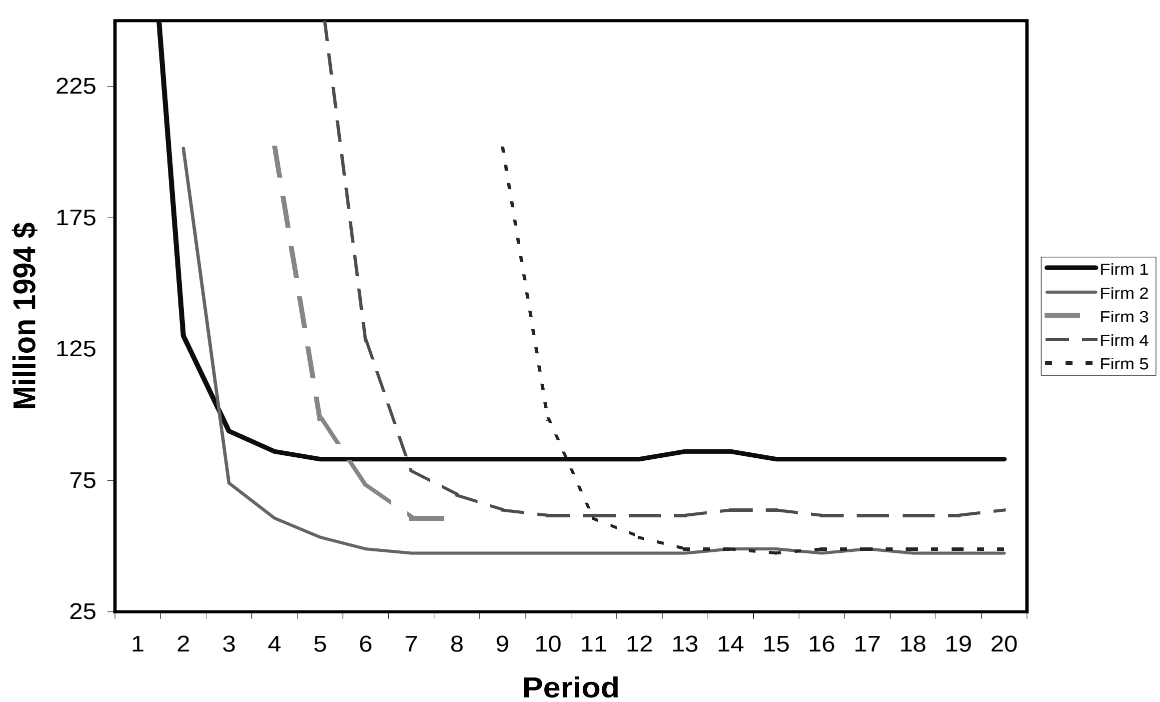

Figure 7: **Twenty-Year Simulation: Units Produced**

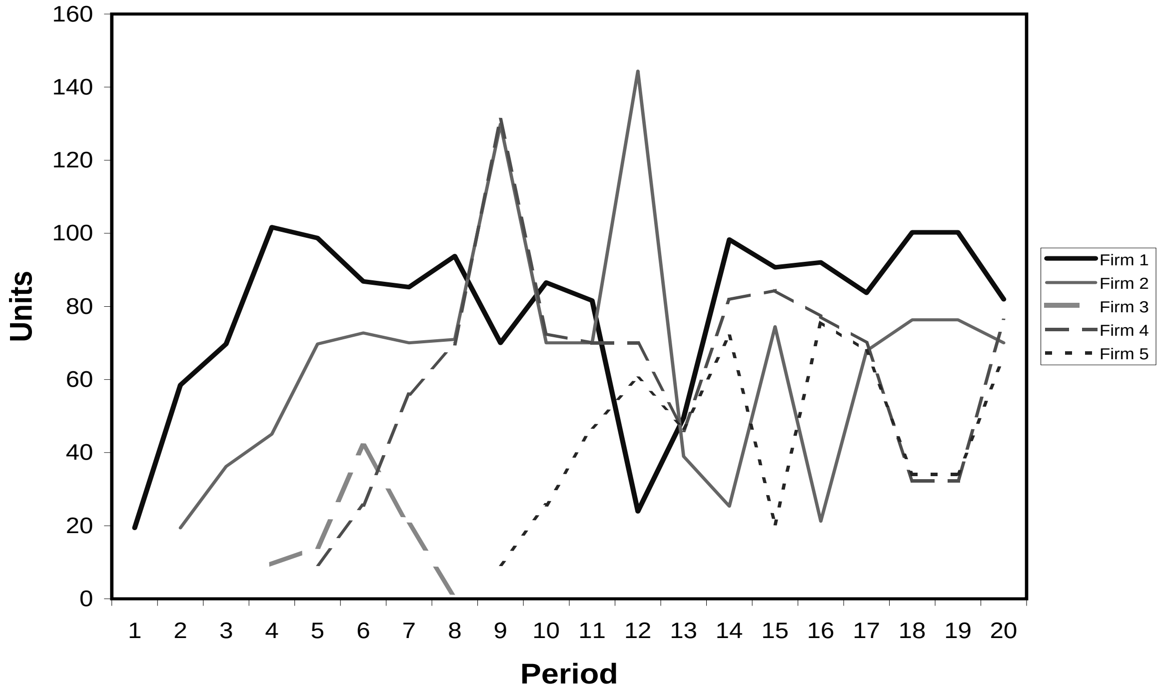

Figure 8: **Twenty-Year Simulation: Realized Discounted Value of the Firm**

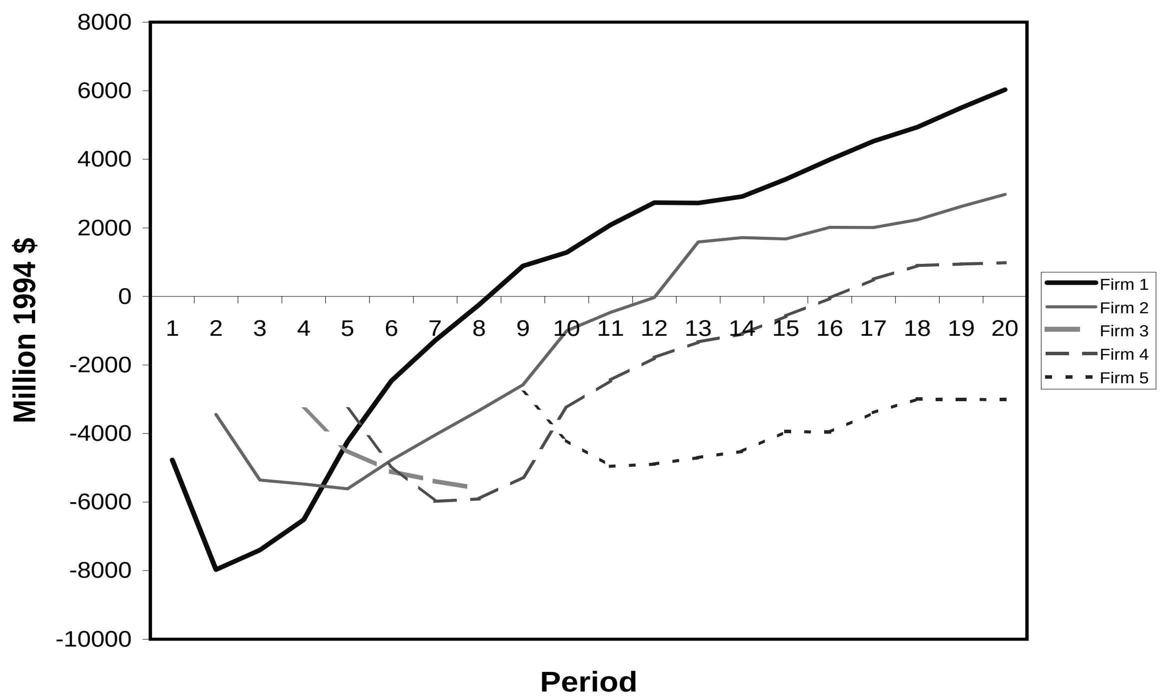

the industry reaches maturity, firm 4 and especially firm 5 enter with essentially their long run product price.

The second point exhibited by the simulations is that profit realizations have very high variance in this model. Firm 3 makes losses in *every* period that it operates from the time that it enters (see Figure 8). In that sense, this firm is reminiscent of the L-1011 and particularly the price versus average-variable-cost graph shown in the beginning of the paper. The model tells us that in expectation it is optimal for the firm to enter and it is optimal for it to remain in the market for all five years. It happens that, despite acting in an optimal manner, this firm receives bad market realizations which cause it to make large losses totaling about $6 billion dollars. In hindsight and without an intimate knowledge of the industry the firm's actions may appear to have been suboptimal, but the model tells us that this is not the case.

A third point is also exhibited particularly well by Figure 8 and that is that firms always start out by losing money in early periods, even net of development costs. In most cases, firms go on to make profits in future periods, though often (45%) these profits are not large enough to make the program an overall success. See, for example, firm 5 in the simulation. It should be noted that the shape of the cash flow curves in Figure 8 looks remarkably like cash flow charts published by aircraft industry firms. Firms in the model generally reach profitability within 10-15 years if they are going to at all. This feature is also very consistent with industry norms.

9 Industry Performance: Alternative Market Structures

In this section, the base market structure (MPE) is compared with two alternatives: a multi-product monopolist (M) and a multi-product social planner (SP). To accomplish this comparison it was necessary to calculate a new equilibrium under each of the two alternative market structures using the same parameters as in the base model. The primary difference to the model is that in each case there is now only one optimizing agent. Therefore, for each of the two alternatives there is a unique value function and associated policy function and the

Table 7: Invariant Distribution Under Alternative Market Structures

Market Structure:		**Social-Planner**	**MPE**	**Monopolist**
Concentration Ratios:				
1-Plane		0.68	0.40	0.98
(S.D.)		0.18	0.15	0.08
2-Plane		0.94	0.69	1.0
(S.D.)		0.09	0.11	0.0
Market Size:				
# Firms		2.4	3.8	1.03
(S.D.)		0.4	0.4	0.03
Avg. Quantity Per Period:				
small		91	174	69
medium		120	51	21
large		129	6	0.3
Avg. Price Per Unit:				
small		47.6	66.1	89.5
medium		61.8	81.6	104.1
large		83.2	105.4	127.4
Avg. MC Per Unit:				
small	Lowest=47.3	48.9	50.5	49.2
medium	Lowest=61.6	63.3	65.5	64.1
large	Lowest=83.1	84.7	86.2	89.6
Avg. (Price/MC):		0.94	1.24	1.76
(Avg. Price)/(Avg. TC):		0.95	1.19	1.66
EDV of New Product Investment:				
Mean		25011	38940	6800
(S.D.)		8404	10227	4499
EDV of Consumer Surplus:				
Mean		209353	135325	73526
(S.D.)		12840	6957	7171
Min		171585	110286	43307
Max		243133	152224	89919
EDV of Producer Surplus:				
Mean		-14966	42363	61374
(S.D.)		2094	3754	5237
Min		-22550	30704	35207
Max		-8299	52668	74820
EDV of Total Surplus:				
Mean		194387	177689	134337
(S.D.)		12434	10333	12144
Min		153563	142133	78514
Max		228081	204182	163767

Table 8: Distribution of Gain Under Alternative Market Structures

Distribution of Gain for:	**SP over MPE**	**SP over M**	**MPE over M**
Consumer Surplus Gain:			
Mean	74026	135826	61800
(S.D.)	9276	12076	8446
Min	38825	101103	34707
Max	102943	172481	91294
% Pos.	100%	100%	100%
Producer Surplus Gain:			
Mean	-57329	-75777	-18448
(S.D.)	4408	5623	5423
Min	-70598	-91154	-34277
Max	-42687	-46760	3176
% Pos.	0%	0%	0.3%
Total Surplus Gain:			
Mean	16698	60049	43352
(S.D.)	8465	13478	13301
Min	-20936	19033	1701
Max	44358	105097	93620
% Pos.	96%	100%	100%
Mean Consumer Gain:	55%	187%	86%
Mean Producer Gain:	-136%	-125%	-30%
Mean Total Gain:	9.4%	46%	33%

solution algorithm is a contraction mapping. Tables 7 and 8 show statistics for the invariant distribution of under each of the three market structures.

According to the simulations, the MPE is quite efficient from a social perspective. On average, the social planner increases total surplus by just 9% ($17 Billion) over the MPE. However, consumers are a great deal better off and producers a great deal worse off with the social planner. The monopolist, on the other hand, provides much lower social welfare than either the social planner or the MPE, at great expense to consumers.

While the social planner does have the lowest production costs on average, surprisingly the welfare improvements from the social planner are not driven primarily by lower marginal costs through learning as suggested by the theoretical models. All three market structures lead to fairly efficient production. Instead, welfare gains under the social planner result primarily from more standard sources. The social planner sets price approximately equal to marginal cost and produces about 40% more total output per period on average than the competitive firms. The competitive firms in turn produce about 2.5 times as much total output as the monopolist.

The second area of welfare savings under the social planner results from concentrating output among just 2.3 firms on average, as compared with 3.8 in the competitive case, which leads to approximately a 40% reduction in new product investment. It seems that in the competitive case there is excess investment in development of new planes and wasted investment in learning to produce these products efficiently. Thus, while marginal costs and prices are lowest on average under the social planner, concentration ratios are higher than in the MPE case.

10 Policy Experiment: Restricting Concentration

In a recent article, *The New York Times* referred to Boeing as "essentially a government-sanctioned monopoly". Theoretically, there is reason to believe that high concentration may be socially beneficial in industries with strong learning curves. In the absence of perfect spillovers of experience between firms, it is always cost-minimizing to concentrate production as much as possible. However, the standard welfare reducing effects of monopoly are also present: in an unconstrained monopoly there tend to be fewer products and lower total production, both of which reduce welfare.

An advantage to having such a detailed model is that it becomes possible to evaluate which effect will dominate in this particular industry, and hence to determine whether the current policy is the correct one. Specifically, I consider alternative anti-trust policies under which firms are punished if they become "too large" as measured by the industry one-firm concentration ratio (a *per se* restriction on concentration as considered in Dasgupta and Stiglitz (1988)). Under such a policy no single firm will choose to produce more than a certain percentage of the aircraft sold in a given year. Note that with the policy in place equilibrium strategies differ from those described above. Firms know that the policy exists and therefore, since the policy changes payoffs in certain states, equilibrium strategies must also change. Thus, in order to evaluate alternative policies it was necessary to re-solve the model for a new equilibrium in each case.[26]

Table 9 lists some summary statistics drawn from industry simulations under the base case (no restriction) and two alternative policies. Simulations were initiated at the observed industry structure for 1994 and use identical random draws for each policy. The experiment is thus analogous to implementation of the given policy alternative beginning in that year. Figures reported are present discounted values from 1000 simulations of 100 periods each.

Table 9 shows that the impact of the two alternative policies on the predicted distribution of concentration ratios is only slight. In that sense, the concentration restriction policy is somewhat unsuccessful in that its effect seems to be largely limited to those states in which

[26] It was also necessary to make an assumption about what occurs in states where there is outright monopoly, since at these states the one-firm concentration ratio would always be one. In this case I assumed that the firm would be a regulated monopolist and must set price equal to marginal cost, but would be reimbursed for its fixed costs. Thus in monopoly states the firm makes zero current profit (and consumers benefit). Note that in equilibrium such states are only reached 0.01% of the time so the exact assumption made is irrelevant to the results.

Table 9: Statistics from 1000 Industry Simulations Under Alternative Policies

	Invariant Distribution		
Maximum Concentration:	**100%**	**60%**	**51%**
Concentration Ratios:			
1-Firm/Plane	0.40	0.40	0.40
(S.D.)	0.10	0.09	0.08
2-Firm/Plane	0.70	0.70	0.70
(S.D.)	0.11	0.11	0.11
Consumer Surplus:			
Mean	135326	134899	133927
(S.D.)	6957	7235	7428
Min	110286	102774	99164
Max	152224	152521	153024
Producer Surplus:			
Mean	42363	42351	42364
(S.D.)	3754	3790	3780
Min	30704	30078	30702
Max	52668	52691	52502
Total Surplus:			
Mean	177689	177250	176291
(S.D.)	10333	10597	10747
Min	142133	135772	132253
Max	204182	204096	205060

Table 10: Distribution of Harm Under Alternative Policies

Maximum Concentration:	**60%**	**51%**
Consumer Surplus Harm:		
Mean	427	1399
(S.D.)	1029	1864
Min	-8912	-8216
Max	9065	13614
% Pos.	95%	94%
Producer Surplus Harm:		
Mean	11.7	-1.2
(S.D.)	523	885
Min	-3036	-5463
Max	6320	8961
% Pos.	29%	32%
Total Surplus Harm:		
Mean	439	1398
(S.D.)	1339	2427
Min	-10718	-10080
Max	13338	19401
% Pos.	94%	92%
Mean Consumer Harm:	0.3%	1.1%
Mean Producer Harm:	0.0%	0.0%
Mean Total Harm:	0.3%	0.8%

the policy actually binds, without having too great an impact on other states. Table 9 also shows that the consumer and producer surplus distributions are lower the stronger the policy alternative. However, differences are quite small compared with the variance in these distributions making it difficult to conclude how harmful the policies are.

Table 10 lists the distribution of welfare gains/losses across the 1000 simulations, making it clear that welfare losses are in fact systematic. Both policies reduce consumer surplus and total surplus in over 90% of the simulations. Mean total harm in the 51% policy is 0.8%, or $1.4 Billion in present value terms. While the effect is fairly small relative to its standard deviation of $2.4 Billion, again confirming that the policy is not binding in many states, these results ought to be indicative of what might occur if such a policy was instituted in a market with multi-product firms.

The concentration restriction primarily binds in states where one firm has high quality and low cost and the others do not. In these states, the primary effect of the policy is to restrict the dominant firm's output. Within the period, the negative effects of the policy are rather straightforward since reduced sales by the dominant firm tend to lower consumer surplus. However, there is also an opposing positive effect in that the policy causes weaker firms to react by increasing their output. These two within-period effects also have dynamic implications since the dominant firm is less likely to remain a low-cost producer while weaker firms are more likely to move down their learning curves. The evidence in Table 10 suggests that once all effects are accounted for, both alternative policies lead to welfare losses overall.

The distributional effects of the policy are somewhat complicated by the stochastic nature of the dynamics in the model (which reflect the industry itself). In spite of the fact that welfare losses occur in over 90% of the simulations and that losses are quite large in some cases, in one case the 60% policy improves total welfare by approximately 7% ($13 Billion). The reason for this outcome is quite complex. There are certain sequences in which the firm that is dominant in the market at the start of the simulations receives bad draws on quality very early, while smaller firms simultaneously receive good ones. In such sequences, a government policy which hurts this dominant firm in early periods and helps smaller ones is welfare enhancing. Of course there is no way that any government could possibly foresee this occurrence. Moreover, there exist sequences where the opposite occurs and a concentration restriction can result in as much as a 17% ($11 Billion) welfare loss. The conclusion that a concentration restriction would be welfare reducing thus holds only in expectation. As noted in Table 10, approximately 6-8% of the time such a policy would increase welfare. This result is quite intuitive and underscores the richness of the overall model.

It is also interesting that according to the surplus figures in Table 10 producers as a whole should be more indifferent to this policy than consumers since it hurts them very little on average. This result may at first seem counter-intuitive since anti-trust policy is usually thought of as pro-consumer. However, due to the presence of learning curves, concentrating production lowers cost far enough that consumers may actually experience lower prices in situations where concentration is high. Furthermore the conclusion that the policy does not harm producers does not account for the fact that the distribution of the policy's effects is highly skewed. All producer losses from the policy in any given period are experienced by only one firm. Thus, if such a policy were proposed, according to the model we should expect consumers (airlines) and the dominant firm to oppose this policy and weaker competing firms to support it.

Finally, note that in evaluating this policy we have kept firms' investment in product quality, which is modeled in reduced form, fixed. In actuality, since the concentration restriction reduces producer surplus in high quality states, investment in quality would likely fall on average under the policy. In that case, since producers do not account for the social benefit of increased investment, it also seems likely that further harm would result from the policy, so that the results actually represent a lower bound to the harm distribution.

11 Conclusions

This paper represents a first attempt at building an empirical dynamic equilibrium model of an industry with learning by doing. It is also a first attempt at constructing an empirical multi-agent model that generates prices and quantities endogenously in dynamic equilibrium.

Despite many simplifying assumptions, the model predicts many aspects of the historical data well, particularly the periodic pricing of aircraft below the level of static marginal cost. The model also replicates many aspects of observed industry dynamics, including entry, exit, concentration ratios, plane value, and plane type distributions.

While computation of the model equilibria has proven to be quite computationally intensive, the evidence presented here provides tentative support for the use of this class of dynamic models in empirical work more generally. This is an important finding since in the past analysis has been difficult, if not impossible, in industries where dynamics play an important role. Indeed, despite many important policy implications, to my knowledge noone has undertaken such a detailed analysis of the aircraft industry previously, primarily due to the intractability of the problem.

At the same time, the aircraft industry is one of the simpler cases to work with because the small number of firms and products in the industry provide some relief from the curse of dimensionality. Given the great computational burden of modeling even a small industry such as this one, the outlook may at first seem grim with respect to tackling larger problems. However, there are extensions to the general class of multi-agent dynamic models used here that would allow application to other more complex and higher dimensional industries. Such extensions include modeling several dominant firms individually and treating remaining firms as acting together through one combined agent. In many industries where there are only a few leading firms and a large number of "fringe" firms, such an assumption would not be unreasonable.

Having such a detailed model is a great advantage because it is well-suited to analyzing various policies by simply re-solving the model with alternative institutions in place. I have evaluated three alternative market structures, with results suggesting that the single-product MPE is on average quite efficient from a social perspective. I also evaluated a policy which would restrict one-firm concentration in the aircraft industry, with the conclusion that such a policy would reduce total welfare with high probability. Both of these policy evaluations suggest caution with respect to government intervention in the aircraft industry. However, the model also suggests that an *uncontested* monopolist producer would lead to a large loss in social efficiency, an outcome which should be avoided if possible.

With some simple extensions, the model could also be extended to look at other anti-trust alternatives such as the break-up of a multi-product firm, regulation strategies, and various strategic trade policies, all of which are of current relevance in the commercial aircraft industry. With some further extensions, such as the explicit modeling of the firm's investment in product quality, the effects of R&D subsidies could also be considered. None of these types of policy simulations would be possible without a fully specified model.

A Discretized Experience Process

Define e_t to be the experience state number for a given firm at time t, and discretize experience into seven points:

$$E[] = \{1, 10, 20, 40, 70, 110, 165\}$$

Thus e_t can take on values in the range [1..7], and $E_t = E[e_t]$. These seven points characterize approximately the observed range of experience in the data. Experience was discretized at closer intervals for low levels of experience because cost changes more quickly when experience is low.

As described in the text, experience is assumed to evolve as follows:

$$E_t = \delta E_{t-1} + q_{t-1} \qquad \text{and} \quad E_1 = 1 \tag{11}$$

However, equation (11) describes a deterministic continuous process that I need to transform into a stochastic discrete process for the purposes of the dynamic model.

To accomplish this, I first calculate E_t from E_{t-1} and q_{t-1} using (11). Then, I compare E_t to the discretized points $E[]$ to see what range it falls in. Define e_d and e_u to be the two closest discretized points such that $E[e_d] \leq E_t \leq E[e_u]$. Then the distribution of e_t given e_{t-1} and q_{t-1} is defined as follows:

$$e_t = \begin{cases} e_d & 1 - \frac{E_t - E[e_d]}{E[e_u] - E[e_d]} \\ e_u & \frac{E_t - E[e_d]}{E[e_u] - E[e_d]} \end{cases} \tag{12}$$

Hence, in theory, e_t communicates with all other values of e, but in equilibrium it only communicates with two: e_d and e_u.

A.1 Deterministic vs Stochastic Experience Evolution

The cost function estimates from Benkard (2000) assume a deterministic progression for experience. However, it seems more likely that in reality experience evolves stochastically, i.e. when you build a plane, sometimes you learn and sometimes you do not. Hence, the stochastic progression of experience used in the industry model may in fact be more realistic than a deterministic one.

Moreover, re-estimating the model from Benkard (2000) under the assumption that experience follows a stochastic process does not change the results. Without looking at further implications of the stochastic model, such as the fact that it predicts greater variance in experience as time goes on, it is not possible for the data to distinguish which model is the correct one. Thus, assuming stochastic evolution of experience in the industry model is justified.

A.2 Discrete vs Continuous Experience

Assuming that experience must make large discrete movements rather than small continuous ones has the effect of convexifying the value function within these ranges. Where the integral of the value function would otherwise be a smooth function of quantity, it now is essentially piecewise linear, with a quantity derivative that makes a finite number of discrete jumps.

The main impact of these discrete jumps on the dynamic model is that, in the maximization process, firms frequently optimally choose to produce a quantity level that puts them at a certain level of experience with probability one. This result arises because the marginal loss from increasing quantity is smooth while the marginal benefit is piecewise linear as described above. Hence the point of equality (or tangency) often occurs at a cusp, where the derivative of the marginal benefit is undefined.

B References

Alchian, A. (1963), "Reliability of Progress Curves in Airframe Production," *Econometrica*, **31:4**, 679-694.

Andrews, D. W. K. (1991), "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, **59:3**, 817-858.

Argote, L. and D. Epple (1990), "Learning Curves in Manufacturing," *Science*, **247**, 924.

Argote, L., Beckman, S. and D. Epple (1990), "The Persistence and Transfer of Learning in Industrial Settings," *Management Science*, **36:2**, 140-154.

Arrow, K. J. (1962), "The Economic Implications of Learning by Doing," *Review of Economic Studies*, **29**, 155-173.

Asher, H. (1956), *Cost-Quantity Relationships in the Airframe Industry*, R-291, Santa Monica, Ca: The RAND Corporation.

Aviation Week & Space Technology, "Tri-Star Production Costs Offset Lockheed Profits," (October 15, 1979), 32.

Bailey, C. D. (1989), "Forgetting and the Learning Curve: A Laboratory Study," *Management Science*, **35:3**, 340-352.

Baldwin, R. E. (1969), "The Case Against Infant-Industry Tariff Protection," *Journal of Political Economy*, **77**, 295-305.

Baldwin, R. and P. Krugman (1988), "Industrial Policy and International Competition in Wide-bodied Jet Aircraft," in R. E. Baldwin (ed.) *Trade Policy and Empirical Analysis*, (University of Chicago Press for the NBER), 45-71.

Benkard, C. L. (2000), "Learning and Forgetting: The Dynamics of Commercial Aircraft Production," *American Economic Review*, forthcoming.

Berry, S. (1994), "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, **25:2**, 242-262.

Berry, S., J. Levinsohn, and A. Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, **63:4**, 841-89.

Boston Consulting Group (1972), *Perspectives on Experience*, Boston: Boston Consulting Group.

Budd, C., C. Harris, and J. Vickers (1993), "A Model of the Evolution of Duopoly: Does the Asymmetry between Firms Tend to Increase or Decrease?," *Review of Economic Studies*, **60**, 543-573.

Cabral, L. M. B., and M. Riordan (1994), "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, **62:5**, 1115-1140.

Dasgupta, P., and J. Stiglitz (1988), "Learning-by-Doing, Market Structure and Industrial and Trade Policies," *Oxford Economic Papers*, **40**, 246-268.

Davidson, R. and J. G. MacKinnon (1993), *Estimation and Inference in Econometrics* Oxford: Oxford University Press.

Ericson, R., and A. Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, **62:1**, 53-83.

Flaherty, M. T. (1980), "Industry Structure and Cost-Reducing Investment," *Econometrica*, **48**, 1187-1210.

Fudenburg, D., and J. Tirole (1983), "Learning-by-Doing and Market Performance," *Bell Journal of Economics*, **14**, 522-530.

Ghemawat, P., and M. Spence (1985), "Learning Curve Spillovers and Market Performance," *Quarterly Journal of Economics*, **100**, 839-852

Gilbert, R. J., and R. G. Harris (1981), "Investment Decisions with Economies of Scale and Learning," *American Economic Review*, Papers and Proceedings, **71**, 172-177.

Gowrisankaran, G. (1997), "Antitrust Policy Implications or a Dynamic Merger Model," Mimeo, University of Minnesota.

Gowrisankaran, G. (1997), "Dynamic Equilibrium in the Hospital Industry," *Journal of Economics and Management Strategy*, **6:1**, 45-74.

Gruber, H. (1992), "The Learning Curve in the Production of Semiconductor Memory Chips," *Applied Economics*, **24**, 885-894.

Gulledge, T. R., and N. K. Womer (1986), *The Economics of Made-to-Order Production*, New York: Springer-Verlag.

Hirsch, W. Z. (1952), "Progress Functions of Machine Tool Manufacturing," *Econometrica*, **20:1**, 81:82.

Hirsch, W. Z. (1956), "Firm Progress Ratios," *Econometrica*, **24**, 136-143.

Irwin, D. A., and P. J. Klenow (1994), "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, **102:6**, 1201-1227.

Judd, K. L. (1992), "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory*, **58:2**, 410-452.

Klepper, G. (1990), "Entry into the Market for Large Transport Aircraft," *European Economic Review*, **34**, 775-803.

Maskin, E., and J. Tirole (1988a), "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica*, **56:3**, 549-569.

Maskin, E., and J. Tirole (1988b), "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, **56:3**, 571-599.

McFadden, D. (1995), "Computing Willingness-To-Pay in Random Utility Models," Mimeo, University of California, Berkeley.

Mookherjee, D. and D. Ray (1991), "Collusive Market Structure Under Learning-By- Doing and Increasing Returns," *Review of Economic Studies*, **58**, 993-1009.

Newhouse, J. (1982), *The Sporty Game*, New York: Alfred A. Knopf Inc.

New York Times, "Aircraft GIant Tries to Recover From Hard Landing" (November 14, 1997).

Nye, W. (1996), "Firm-Specific Learning-by-Doing in Semiconductor Production: Some Evidence from the 1986 Trade Agreement," *Review of Industrial Organization*, **11**, 383-394.

Olley, G. S., and A. Pakes (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, **64:6**, 1263-1298.

Pakes, A. (1991), "Dynamic Structural Models: Problems and Prospects. Mixed Continuous Discrete Controls and Market Interactions," in C. A. Sims (ed.) *Advances in Econometrics Sixth World Congress*, (Cambridge University Press), Vol 2, 171-259.

Pakes, A., G. Gowrisankaran, and P. McGuire, "Implementing the Pakes-McGuire Algorithm for Computing Markov Perfect Equilibria in Gauss," Mimeo, Yale University.

Pakes, A., and P. McGuire (1994), "Computing Markov Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, **25:4**, 555-589.

Pakes, A., and P. McGuire (1997), "Stochastic Approximation for Dynamic Models: Markov Perfect Equilibrium and the 'Curse' of Dimensionality," Mimeo, Yale University.

Rapping, L. (1965), "Learning and World War II Production Functions," *Review of Economics and Statistics*, **47**, 81-86.

Scherer, F. M., and D. Ross (1990), *Industrial Market Structure and Economic Performance*, 3rd ed. Boston: Houghton Mifflen.

Searle, A. D. and C. S. Gody (1945), "Productivity Increases in Selected Wartime Shipbuilding Programs," *Monthly Labor Review*, **61:6**, 1132-1147.

Seitz, F., and L. Steele (1985), *The Competitive Status of the U.S. Civil Aviation Manufacturing Industry*, Washington, DC: National Academy Press.

Smiley, R.H., and S. A. Ravid (1983), "The Importance of Being First: Learning Price and Strategy," *The Quarterly Journal of Economics*, **98:2**, 353-362.

Spence, A. M. (1981), "The Learning Curve and Competition," *Bell Journal of Economics*, **12**, 49-70.

Stokey, N. L., R. E. Lucas [with E. C. Prescott] (1989), *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press.

Tirole, J. (1988), *The Theory of Industrial Organization*, Cambridge, MA: MIT Press.

Wall Street Journal, "Lockheed is Planning a Write-Down of $600 Million on Its Tri-Star in a Complex Financial Restructuring," (June 3, 1974), 1.

Wall Street Journal, "Lockheed Net Soared 108% in 3rd Period but Outlook for Tri-Star Sales is Dismal," (November 7, 1975), 1.

Webbink, D. W. (1977) *The Semiconductor Industry: A Survey of Structure, Conduct, and Performance*, Washington, DC: US Government Printing Office.

Womer, N. (1979), "Learning Curves, Production Rate, and Program Costs," *Management Science*, **25:4**, 313-319.

Womer, N. (1984), "Estimating Learning Curves from Aggregate Monthly Data," *Management Science*, **30:8**, 982-992.

Wright, T. P. (1936), "Factors Affecting the Cost of Airplanes," *Journal of the Aeronautical Sciences*, **3**, 122-128.

Zhang, R. (1996), "An Economic Analysis of Commercial Jet Aircraft Prices," Ph.D. Dissertation, Boston University.

Zimmerman, M. B.. (1982), "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power," *Bell Journal of Economics*, **13**, 297-310.