

# A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft

C. LANIER BENKARD

*Graduate School of Business, Stanford University*

*First version received May 2000; final version accepted March 2003 (Eds.)*

This paper uses an empirical dynamic oligopoly model of the commercial aircraft industry to analyse industry pricing, industry performance, and optimal industry policy. A novel feature of the model with respect to the previous literature is that entry, exit, prices, and quantities are endogenously determined in Markov perfect equilibrium (MPE). We find that many unusual aspects of the aircraft data, such as high concentration and persistent pricing below static marginal cost, are explained by this model. We also find that the unconstrained MPE is quite efficient from a social perspective, providing only 10% less welfare on average than a social planner would obtain. Finally, we provide simulation evidence that an anti-trust policy in the form of a concentration restriction would be welfare reducing.

## 1. INTRODUCTION

Due to steep learning curves and high entry costs, the commercial aircraft industry is an industry in which many conventional wisdoms of static economic theories do not apply. For example, as firms work down their learning curves it is common to observe prices that are well below static marginal cost, a result that is inconsistent with static profit maximization. The aircraft industry has also been a frequent target of industrial policy, most notably in Europe, and is subject to several special trade agreements, including the 1979 GATT agreement and the 1992 bilateral agreement between the U.S. and Europe. These issues have been addressed in part in the recent theory literature (Spence (1981), Fudenberg and Tirole (1983), Dasgupta and Stiglitz (1988), Cabral and Riordan (1994) and others) as well as in the recent international trade literature (Baldwin and Krugman (1988), Klepper (1990), Tyson (1992)). The approach of this paper is to use a fully specified empirical dynamic model. Our hope is to gain a better understanding of strategic interaction in the industry, as well as to develop a framework in which to better evaluate trade and anti-trust policies.

The past theoretical literature on learning by doing, which we use as guidance, has focused on two main areas. First, the literature has found that learning curves have important strategic implications. Pathbreaking papers by Spence (1981) and Fudenberg and Tirole (1983) showed that learning curves can lead to a dynamic incentive to set the price below the level of static marginal cost, especially upon the introduction of a new product. Dasgupta and Stiglitz (1988) showed that industries in which learning curves are steep may tend to become highly concentrated over time. The paper's first goal is to test these theoretical models on industry data. We do this by estimating the primitives of a dynamic oligopoly model and then comparing the equilibrium predictions of the model with the observed data.

Additionally, we address the following striking empirical fact: data obtained for the Lockheed L-1011 shows that price was below static marginal cost for essentially its entire 14-year production run (see Figure 1), for a total variable loss of approximately \$2.5 billion




FIGURE 1  
Lockheed L-1011: price vs. AVC 1972-1985

current dollars.<sup>1</sup> While this is the only aircraft model for which detailed cost data is available, data on profits for other firms suggests that many other aircraft have also been sold at a variable profit loss. We investigate whether a dynamic oligopoly model with learning parameters that are estimated from aircraft production data can explain this pricing behaviour, or whether, as suggested by some (*e.g.* Reinhardt, 1973), a rational firm should have set a higher price or exited.

The second area that the past theoretical literature has focused on is optimal policy. Dasgupta and Stiglitz (1988) showed that learning curves may lead to a specialized policy environment in which standard policy prescriptions may perform poorly. They showed that there are situations in which an unrestrained monopolist may be socially preferable to any market with more than one firm, which suggests that a blind application of existing anti-trust policies may not be desirable in the commercial aircraft industry. They also showed that when learning is strong it may increase a country's welfare to protect an infant industry.

While the theory literature has been very insightful and has vastly increased our understanding of industries like this one, there are two weaknesses of the theoretical approach. The first is that the theory 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 typically depend on the values of important model parameters such as learning rates and demand conditions. Both of these shortcomings point to the need for empirical work to provide a more definitive analysis. Thus, the second goal of this paper is to use our empirical model to investigate optimal policy.

1. Source: author's calculations based on Lockheed's annual reports. This figure does not include development costs of another \$2.5 billion.

While even the simplest models of dynamic oligopoly can be quite complex, an empirical analysis of our questions of interest would not be possible without such a model. Because firms must introduce new aircraft models at prices below the level of static marginal cost while they work down their learning curves, a static model would have difficulty justifying production in many cases, let alone explaining why rational firms would pay entry costs in the billions of dollars. Furthermore, understanding the initial post-entry period, in which large amounts of surplus are transferred from producers to consumers, is important to understanding the industry's overall performance. Therefore, static models are not likely to provide a very useful analysis. We also show that the observed mark-ups are not consistent with single-agent optimization, but are consistent with equilibrium behaviour in oligopoly, so it would not be possible to simplify the analysis by using a single-agent model. Finally, because dynamic profit incentives are of primary importance in the industry, a complete analysis of a policy change requires analysis of the equilibrium response of firms to the policy, which is only possible in an equilibrium model.

In developing the oligopoly model, we build on recent theoretical advances in modelling dynamic oligopoly by Pakes and McGuire (1994, 2001) and Ericson and Pakes (1995). Due to the curse of dimensionality inherent to dynamic programming models, it would not normally be feasible to apply these discrete time, discrete state space, dynamic game models to a real world industry. However, the fact that the aircraft industry is a small oligopoly, with only a handful of firms and products, makes the approach computationally feasible.

This paper also makes several methodological contributions. The paper is among the first empirical applications of a fully specified dynamic oligopoly model (cf. Gowrisankaran and Town, 1997). Our model is also a significant extension to the literature on computable discrete time dynamic games (*e.g.* Ericson and Pakes, 1995), as it is the first model of this type to incorporate dynamics in the product market equilibrium. In the previous literature (including Berry and Pakes (1993), Pakes and McGuire (1994, 2001), Gowrisankaran (1999)), the product market was modelled as static price or quantity setting, with dynamics coming through firms' investment decisions. We feel that the extension of dynamics to the product market is potentially quite important for the I.O. literature because it covers several important cases in addition to learning by doing, including durable goods, network effects, switching costs and others. Note that Fershtman and Pakes (2000) also allow for current prices to influence future states through a state variable that reflects past cooperation in a collusive scheme.

Our empirical strategy involves three steps. First, we write down a dynamic model that we feel represents the industry well, maintaining the key features of the industry such as learning curves, product differentiation, large entry costs and closed loop strategic interaction, while abstracting away several less important features such as multi-unit demand. Second, we estimate the primitives of the model. The model has both a supply and a demand system, which we estimate separately, in a manner that is consistent with the underlying dynamic model. We simultaneously estimate the stochastic processes determining the transition processes for the state variables. Third, we numerically compute the equilibrium of the model, test the model predictions against the data, and evaluate several counterfactual policies.

There are several advantages to this three-step approach. Because equilibrium is not enforced in the estimation procedures, consistency of the parameter estimates does not depend on the equilibrium assumptions and is therefore robust to a wide set of possible assumptions. In addition, there is no need to solve the dynamic programming problem during estimation, which greatly reduces the computational burden of the estimation procedures. Finally, since the equilibrium behaviour of firms is not used in estimation, the equilibrium predictions of the model can be used as a test of the theoretical model. If the equilibrium model fits the data well, we take that as evidence in support of the model. The main disadvantage of our approach is that, if the equilibrium assumptions are true, then greater efficiency could be obtained in the estimates by

enforcing equilibrium during estimation. Another disadvantage of this approach is that it does not easily provide a formal test of alternative models.

We find that, despite some simplifications, the dynamic model predicts many aspects of equilibrium behaviour well, particularly those that have been the focus of the past theoretical literature. It improves vastly on previous attempts at modelling aircraft industry pricing (cf. Baldwin and Krugman (1988), Klepper (1990)). For example, even though observed mark-ups 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 mark-ups that are slightly higher than those observed, but we 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.

We also find that the model approximately replicates Lockheed's pricing and exit strategies for the L-1011, suggesting that Lockheed's behaviour was in fact consistent with profit maximization. The results of the model suggest that Lockheed always had reason to expect a reasonable chance of future success that, coupled with the incentive to move down its learning curve, led to an incentive to continually price below static marginal cost, but not to exit.

The model is also well suited to a detailed analysis of industry performance. In Section 7 we consider three alternative market structures: single-product firms, a multi-product monopolist and a multi-product social planner (SP). The results from this comparison suggest that the single-product firm Markov perfect equilibrium (MPE) is quite efficient from a social perspective, providing only 10% less total welfare on average than the SP could obtain. However, relative to the SP, the MPE shifts a substantial portion of total surplus from consumers to producers. We also find that an unconstrained multi-product monopolist with no threat of entry would lead to large inefficiencies from a social perspective. We go on to consider an anti-trust policy that places a per se restriction on the highest market share any single firm may attain. This policy is equivalent to the type of policy considered in Dasgupta and Stiglitz (1988). We find that such a policy would be welfare reducing with very high probability, particularly hurting consumers.

## 2. THE COMMERCIAL AIRCRAFT INDUSTRY

### *2.1. Industry background*

The commercial jet aircraft industry has existed since 1956, with the first wide-body introduced in 1969. In the empirical analysis, we focus on the market for wide-bodied commercial jets since that market contains a more computationally tractable number of products than the entire commercial aircraft market. 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). Total commercial aircraft industry revenue for 1997 was approximately \$60 billion, of which \$40 billion is attributable to U.S. producers. Commercial aircraft is also among the U.S.'s largest net exports, with trade surpluses averaging about \$25 billion annually over the early 1990s. The commercial aircraft industry, and aerospace more generally, has seen much merger activity in recent years that has led to increased concentration. Currently, only two major producers remain for commercial jets of more than 100 seats.

Many countries regard the commercial aircraft industry as a "strategic" industry. 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. The 1979 GATT Agreement on Trade

in Civil Aircraft substantially liberalized trade in aircraft and aircraft parts, as well as partially exempting aircraft from anti-dumping regulations, allowing them to be sold below short run marginal cost but not below long run marginal cost. The 1992 bilateral agreement between the European Community and the U.S. places limits on government financing of aircraft purchases and government subsidies.<sup>2</sup>

### *2.2. Commercial aircraft pricing: two striking facts*

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 neighbour 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. Observations for the AVC series are annual averages with the exception of the first data point which covers 2 years.<sup>3</sup>

The first striking fact is that AVC exceeded P for essentially the whole 14-year period that the plane was produced. To casual observers and economists alike, this pricing behaviour may at first seem pathological. While it is known that learning curves may sometimes lead to prices below static marginal cost, in the standard theoretical models of learning by doing such periods appear to be short and only occur in the initial post-entry periods. A question that we address in this paper is whether the observed behaviour of prolonged below marginal cost pricing can possibly be explained with a more detailed model.

The second striking fact that can be seen in the graph is that there is much greater variance in cost than in price. This property seems to hold for all products in the market. While costs are known to change drastically over time due to learning curves, prices are quite flat over time for all products in the market, suggesting that mark-ups move to offset the movement in costs.

While neither of these properties could be explained with a static model, we have found that the dynamic model presented in what follows replicates both. Section 5 also shows, perhaps surprisingly, that equilibrium prices for the L-1011 predicted by the model are quite similar in level to those in Figure 1.

### *2.3. Product differentiation in wide-bodied aircraft*

Figure 2 shows a summary of the nine wide-bodies in our data-set, graphed by range and seats. As is true for many aircraft characteristics, range and seats are highly correlated, with a correlation coefficient of 0.95 in our data. While seats and range are not the only important characteristics, the graph shows that most of the aircraft in the market have fairly close competitors. This suggests that these aircraft should on average be highly substitutable in practice, a feature that is also supported by anecdotal accounts of the industry (see Newhouse, 1982). The one exception to this is the 747, which is a great deal larger than its nearest competitor. However, there are smaller planes with as great range as the 747 and an airline always has the option of flying two planes in place of one.

Models of entry for differentiated products are not yet well understood for even just a two-dimensional space. Thus, for the purposes of the dynamic model, we abstract the level of product differentiation somewhat by dividing the market into three groups, by the number of seats (see Figure 2). Due to the high correlation between characteristics, it is implicit that the smaller planes

2. See Tyson (1992) for a more detailed discussion of these agreements.

3. Because the first data point covers 2 years and more than 50 units, the AVC series shown does not precisely reveal the average costs for the first 10 units produced. Using the labour requirements data, we have estimated that AVC averaged \$220 million for the first 10 units.




FIGURE 2

Product differentiation among wide-bodies

also have shorter range, smaller cabins, fewer engines, etc. Throughout the paper we refer to the three types as “small”, “medium” and “large”. We also allow planes to vary in another dimension that we call “quality” (details in Section 3).

### 3. THE THEORETICAL MODEL

#### 3.1. Model and notation

We index products by  $j \in 1 \dots \infty$  and time periods by  $t \in 1 \dots \infty$ . For the purposes of the empirical application, a time period is assumed to be 1 year. In our notation, the product label,  $j$ , remains constant for the same product across time periods.

There are three state variables per active product (plane) in the model. A firm is defined by the set of products it produces and thus may have varying numbers of state variables. Reflecting learning by doing, the first product-specific state is the firm’s production experience with respect to the product, denoted  $E_{jt} \in \mathcal{E}$ . The set  $\mathcal{E}$  is the set of all possible experience levels, which is assumed to be finite. The experience state evolves endogenously as a function of the firm’s past experience and the firm’s current production. More precise definitions of all of the state variables with respect to the empirical application are postponed to the estimation sections of the paper.

The second and third product-specific states are denoted  $\mu_j \in \mathcal{A}$  and  $\xi_{jt} \in \mathcal{X}$ , and reflect the product’s “type” and “quality”, respectively. The sets  $\mathcal{A}$  and  $\mathcal{X}$  refer to the sets of all possible plane types and plane qualities, both assumed to be finite. The plane’s type refers to one of the three size-range classes shown in Figure 2, while the plane’s quality refers to more subjective factors such as reliability, or suitability to the current route network.

Reflecting industry practice, a plane’s type ( $\mu_j$ ) is determined upon entry, after which planes retain the same type over the whole of the product lifetime. Plane quality ( $\xi_{jt}$ ) evolves stochastically over time according to a first-order Markov process. Such a process might result from random shocks such as accidents or changes in the policy environment, or could be the result of random outcomes to investment. While it would be possible, both theoretically and

computationally, to model the quality process in more detail including, for example, a detailed model of firms' investments in product quality, we determined that the amount of investment data available was not sufficient to attempt this.

We assume that the number of products in the market is bounded above. Ericson and Pakes (1995) prove this result for their model. A similar proof holds for this model with similar restrictions on the parameters of the per-period profit function. This result is stated here because of its implication that the state space of the model is finite.

Some further notation is necessary in order to simplify the writing out of firms' value functions. Let  $i \in 1 \dots \infty$  refer to a firm. Let  $\mathcal{J}_i$  be the set of products owned by firm  $i$ . Let  $\omega_{it}$  refer to firm  $i$ 's ( $3 \times \#\mathcal{J}_i$ -dimensional) vector of product-specific state variables and denote the set of all possible firm-specific state vectors as  $\Omega$ . Because the number of products in the market is bounded and each state variable lies in a finite set,  $\Omega$  is finite. Let the industry structure at time  $t$ ,  $s_t$ , be defined as a vector of length  $\#\Omega$  that lists for each  $\omega_0 \in \Omega$  the number of firms at time  $t$  with  $\omega_{it} = \omega_0$ . In a slight duplication of notation, we use the variable  $i_t \in 1 \dots \#\Omega$  to represent firm  $i$ 's position in  $s_t$  at time  $t$ .

There is also one common state variable,  $M_t \in \mathcal{M}$ , that represents aggregate plane demand in period  $t$ . The set  $\mathcal{M}$  is the set of all possible aggregate demand states, which is also assumed to be finite.

The dynamic model consists of two stages within each period. At the beginning of each period, active firms simultaneously make their exit decisions for each of their products. They first observe a random scrap value,  $\Phi_{jt}$ , for each product. They receive the scrap value if they choose to exit the product. Firms choose which products to exit so as to maximize the expected discounted value (EDV) of future profits from their entire portfolio of products.

After the exit phase, each active firm and one potential outside entrant makes an entry choice. Each firm (including the outside entrant) can enter at most one product per period, which can be of any type. The random development cost associated with each product type is observed prior to the entry choice. Production for newly entered products does not begin until the following period. All entering products enter with no production experience.

Simultaneously with the entry choice, the incumbents 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 state variables.

The control variables of each active firm are thus the exit rules,  $\chi_{jt} \in \{0, 1\}$ , and the production choices,  $q_{jt} \in \mathbb{R}^+$ , for each of the firm's products. In addition, each active firm and the potential outside entrant at time  $t$  have the entry rules  $\chi_{it}^e \in \{0, 1, 2, 3\}$ , reflecting no entry and entry of types 1–3 (small, medium or large).

The model as outlined above can be characterized by the following Bellman's equations for active firms  $i$  (subscript  $t$  suppressed):

$$\begin{aligned}
 V(i, s, M) = \max_{\chi_i^e, \chi_j, q_j \forall j \in \mathcal{J}_i} & \left\{ - \sum_{k=1}^3 1\{\chi_i^e = k\} x_k^e \right. \\
 & + \sum_{j \in \mathcal{J}_i} [\chi_j \Phi_{jt} + (1 - \chi_j) \pi_j(i, s, q, M)] \\
 & \left. + \beta \sum_{i', s', M'} V(i', s', M') \mathcal{P}(i', s', M' | i, s, q, M, \chi, \chi^e) \right\} \quad (1)
 \end{aligned}$$

where the unsubscripted variables  $q$ ,  $\chi$  and  $\chi^e$  represent vectors of the controls for all products;  $\pi_j$  is the per-period return function for product  $j$ ;  $k$  indexes the product types;  $x_k^e$  is the random entry cost associated with each product type;  $i'$ ,  $s'$  and  $M'$  are values of the state variables one period into the future; and  $\mathcal{P}$  is the probability distribution generating the transition probabilities of the states.

The potential outside entrant's Bellman's equation is similar:

$$V^e(s, M) = \max_{\chi_i^e \in \{0,1,2,3\}} - \sum_{k=1}^3 1\{\chi_i^e = k\}x_k^e + \beta \sum_{i',s',M'} V(i^e, s', M')\mathcal{P}(i^e, s', M' | s, q, M, \chi, \chi^e) \quad (2)$$

where  $i^e$  is the entry state next period, and  $q_j = 0$  for the entrant's product in the current period.

To complete the model, it is still necessary to specify the per-period profit function,  $\pi$ , and the transition probabilities,  $\mathcal{P}$ , as functions of the state and control variables. The per-period profit function,  $\pi_j$ , takes the form

$$\pi_j(i, s, q, M) = p_j(i, s, q, M)q_j - c_j(i, q_j) \quad (3)$$

where  $p_j$  is the inverse demand function for product  $j$ , and  $c_j$  is the cost function for product  $j$ . The inverse demand function depends on all products' types, qualities and quantities, as well as aggregate demand. Its detailed specification and the stochastic processes for the exogenous state variables are given in Section 4.2. The cost function depends only on the firm's own experience at producing product  $j$ , the type and quality of the product, and the quantity produced. Its specification and the stochastic processes for the endogenous state variables are given in Section 4.1.

### 3.2. Equilibrium concept and discussion

The equilibrium concept used is symmetric Markov perfect Nash equilibrium (MPE), where the strategy space includes quantity, entry and exit. MPE, as defined by Maskin and Tirole (1988a,b), picks out those subgame perfect equilibria where actions are a function only of pay-off 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 their expectations of 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.

Implicit in the notation of (1) is the assumption that quantity is the primary strategic variable. Since aircraft contracts typically specify both quantity and price, the commercial aircraft market is probably neither a strictly price-setting nor a strictly quantity-setting game. However, since the difference between the two extreme cases of price setting and quantity setting was found to be small in preliminary investigations, and implementation of a mixed aversion of the game (Judd, 1996) would further complicate the model, we chose to implement one of the two extreme cases instead. Since aircraft producers fix their production schedules a year or more in advance, and even with that lead they are constrained as to how quickly they can change production rates, we concluded that quantity setting would give the best approximation to the industry. Baldwin and Krugman (1988) and Neven and Seabright (1995) also come to this conclusion. Because quantity is the control variable in the model, prices are determined endogenously in dynamic equilibrium. They are generated by the inverse demand function as the price that sets current demand equal to current supply.

Proof that equilibrium exists for this model is straightforward. We omit the proof from this paper both because it would replicate previous work and because such a proof would be redundant given that our approach in this paper is to solve numerically for equilibrium once the parameters of the model are known. In the event that the numerical algorithm converges, that is sufficient for existence of equilibrium for a specific set of parameters.<sup>4</sup>

4. To be precise, convergence of the numerical algorithm is sufficient for the existence of an  $\varepsilon$ -equilibrium.



A much greater difficulty in dynamic games of this type is that of multiple equilibria. In this paper, we restrict our attention to 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 symmetry did not hold would require additional state variables that would serve to label firms.

The assumption of symmetric equilibrium, while standard in the recent literature on Markov perfect games, is not entirely innocuous. It is likely that there are cases where asymmetric equilibria exist in addition to the symmetric equilibrium that we consider. For example, in our analysis we restrict two similarly situated firms to produce the same quantity. However, it is likely that, in some cases, there would also be an asymmetric equilibrium in which one firm produces more and the other produces less.<sup>5</sup> We make the symmetry assumption specifically to reduce the set of equilibria that we consider.

Our second approach to handling multiple equilibria is to do a direct numerical search. The easiest way to do this is to start the computational algorithm at random locations. In addition, in order to test if the computational algorithm itself was somehow selecting a particular equilibrium, a second and entirely different computational algorithm was also used to search for equilibria. Using these two techniques, no case was identified where there was more than one equilibrium satisfying the restrictions listed above.<sup>6</sup> Thus, our analysis in the results section of the paper is limited to a single equilibrium in all cases.

### 3.3. Computational limitations to the theoretical model

The computational algorithms used to solve for the equilibrium of the model are discussed in a supplementary section.<sup>7</sup> The model as parametrized in Section 4 requires over 100 CPU-days to solve on a Sun Ultra 400 processor, depending on the exact parametrization 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 timeframe. However, the computational burden of solving a dynamic game of this magnitude is 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.

As a result, we found it computationally infeasible to calculate the equilibrium of the model as described above for the most general case of a multi-product firm oligopoly. Thus, in this paper we consider three simpler market structures. Our base case model is a single product firm oligopoly, which we use to analyse the industry. This model ignores joint profit maximization across products owned by the same firm. We also consider a multi-product monopoly, and a multi-product SP. These three cases also simplify the computational algorithm because in each case the number of potential entering products is at most one per period.

## 4. ESTIMATION OF THE MODEL PARAMETERS

This section discusses the specification of the profit function and transition probabilities, as well as the estimation of the model parameters. The section is divided into three main subsections,

5. Similar sets of asymmetric equilibria are also likely with respect to exit.

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

7. Supplementary material available on *The Review's* website.

the cost function, the demand function and other parameters. Data sources are discussed in detail in a supplementary section. Units are 1994 million dollars throughout.

#### 4.1. Cost function

**4.1.1. Labour requirements equation.** The learning by doing cost function described in this section contains the primary dynamic parameters of the model. Therefore, in specifying the cost function we felt it was crucial to work with data that came directly from the commercial aircraft industry, rather than using the more widely available military production data. Much of the work toward that end was accomplished in Benkard (2000). To avoid replication, we only briefly summarize the labour requirements equation here and direct readers to that paper for further details. Estimation of the remaining cost function parameters is discussed in what follows.

Benkard (2000) lists a set of assumptions that lead to the following labour requirements equation:

$$\ln L_{lt} = \ln A + \theta \ln E_t + \gamma \ln S_t + \varepsilon_{lt} \quad (4)$$

where  $L_{lt}$  is the labour input for unit  $l$  in period  $t$ ,  $A$  is a constant,  $E_t$  is experience,  $\varepsilon_{lt}$  is a plane-specific productivity shock, and  $S_t$  is line-speed, a measure of the current production rate.

Many papers have estimated similar learning curve specifications to (4), where experience is defined as cumulative past production ( $E_t = \sum_{i=0}^t q_i$ ). See Argote and Epple (1990) for a good summary of this literature. There is also a relatively large empirical literature on learning curves in aircraft production using this specification, including Wright (1936), Asher (1956), Alchian (1963), Gullledge and Womer (1986) and others. A contribution of Benkard (2000) was 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 well. A similar learning model that also incorporates depreciation of experience capital (“organizational forgetting”) was found to fit the data much better.

In the organizational forgetting model, experience evolves according to a first-order deterministic process. At time  $t$  a firm has experience  $E_t$ . The firm then chooses its current production rate  $q_t$ . Between periods  $t$  and  $t+1$ , the firm’s existing stock of experience depreciates by a factor  $\delta$ , while new experience equal to  $q_t$  is acquired. This process is summarized by the following equation:

$$E_{t+1} = \delta E_t + q_t \quad \text{and} \quad E_1 = 1. \quad (5)$$

The intuition for this specification is as follows. Production experience in the aircraft industry is a form of human capital that is embodied in the 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, lay offs 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 is the case in the L-1011 data. Note that the traditional learning model is a special case of this model with no depreciation ( $\delta = 1$ ).

The primary dynamic parameters of our model are the learning parameter  $\theta$  and the depreciation factor  $\delta$ . In Benkard (2000), the monthly depreciation factor ( $\delta$ ) was estimated to be 0.960. For the purposes of the dynamic model, which uses annual time periods, this number was translated to an annual rate of 0.613 ( $= 0.96^{12}$ ), with a standard error of 0.023. This estimate implies that an aircraft producer loses about 40% of its previous stock of experience every year.<sup>8</sup> This number is large enough that it could potentially lead to somewhat different properties than

8. Please see Benkard (2000) for further discussion of the magnitude of the depreciation parameter.

those of pure learning models. However, one of our findings is that this does not usually seem to be the case. According to the model simulations, most of the time firms exit before large experience declines would occur.

The learning parameter,  $\theta$ , was estimated to be  $-0.63$  with a standard error of  $0.03$ . This estimate implies that the learning rate with respect to experience is a quite rapid  $36\%$ . 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 labour requirements would fall by  $36\%$ . Whether or not this reduction is attainable depends on the current experience level and future production rates.

The returns to scale parameter,  $\gamma$ , was estimated to be  $0.11$  with a standard error of  $0.17$ . The point estimate reflects increasing returns to scale, but the effect is small and insignificant.

The cost specification was also found to fit the data very closely (see Figure 4 of Benkard, 2000), which is important for this paper because the learning/forgetting process is central to our dynamic model. Furthermore, the L-1011 data allowed very precise estimation of the dynamic parameters,  $\theta$  and  $\delta$ .

**4.1.2. Experience process.** The estimated cost function parameters determine not only the static production cost for the firm, but also the transition of the experience state variables over time. The experience process described by (5) is deterministic and real valued. This process could be approximated arbitrarily well in the dynamic model with enough discrete experience states. However, due to computational limitations, we were forced to limit the number of experience states to seven, which we chose to be

$$\mathcal{E} = \{1, 10, 20, 40, 70, 110, 165\}.$$

The range of states in  $\mathcal{E}$  ensures that the entire equilibrium range of experience is covered, and the geometric spacing of the states ensures that the discretization error in cost is always small.

Because computational constraints limited us to a coarse discretization of experience, in order to better approximate the process in (5) we made the evolution of experience stochastic. Let  $E_{t+1}^*$  be the experience level that would be achieved using equation (5),  $E_d$  the largest experience state less than  $E_{t+1}$  and  $E_u$  the smallest experience state greater than  $E_{t+1}$ . Then the stochastic process for  $E_{t+1}$  is specified as a binomial random variable,

$$E_{t+1} = \begin{cases} E_u & \frac{E_{t+1}^* - E_d}{E_u - E_d} \\ E_d & 1 - \frac{E_{t+1}^* - E_d}{E_u - E_d} \end{cases}. \tag{6}$$

By making the experience process stochastic, the effects of the coarse discretization on the value and policy functions are minimized.<sup>9</sup>

**4.1.3. Other cost parameters.** Since the cost model is estimated on data for the L-1011, some further assumptions were necessary to calculate costs for other plane types. The most important assumption we make is that the learning and forgetting parameters,  $\theta$ ,  $\delta$  and  $\gamma$ , are the same for all producers. Since there are no other studies of commercial aircraft production in the literature, this assumption is difficult to test. However, results from the literature on learning curves in military production largely support our assumption. It is widely accepted in the industry that there is a “20% learning curve” in aircraft production, which suggests that the learning

9. The approximation error will depend on the curvature of the value function. In areas where the value function is approximately linear in experience, the discretization has no effect on the value and policy functions. We found that the value function did have some curvature, particularly at the lower experience levels, which is why we chose to discretize the experience state more finely for lower experience levels.

parameters do not vary much across production lines. The only paper that specifically addresses the issue is Alchian (1963), which compares production data for nine bombers, eight fighters and three trainers. Alchian's conclusions relevant to this paper can be summarized as (i) there is no statistically significant difference in the learning parameters across the three aircraft types, (ii) the hypothesis that all 20 aircraft have the same learning parameter is statistically rejected, but (iii) the error when using an aircraft's own production data to forecast future labour requirements is the same magnitude as that if the industry-wide average learning parameters are used.<sup>10</sup> Based upon this evidence, and the precision of our estimates of the learning parameters, we feel that any differences in the learning parameters across aircraft are small enough that this assumption is largely innocuous.

We also need to estimate the level of the labour input,  $A$ , for the different aircraft. However, since labour input data is not available for aircraft other than the L-1011, direct estimation of the constant term for each aircraft was not possible. The standard approach to this problem in the literature is to estimate the labour requirements per pound of the aircraft and then to hold this constant across aircraft. Our approach is similar. We assume that the level of the labour input depends on the type and quality state variables for each product, which are measures of the aircraft characteristics. Specifically, we assume that the variable cost of a larger plane is greater than that of an L-1011 in exact proportion to their relative size.<sup>11</sup>

We also investigated whether the plane quality state affected production cost. The values of the plane quality state variables, which change over time for each product, are estimated with the demand system in Section 4.2, and are therefore known. Despite its having quite large variance in the sample, when added to the cost specification above quality was found not to affect production cost. Since product quality may be most representative of characteristics like suitability of the plane to the current airline route network, which would not necessarily influence marginal cost, this finding is not surprising.

Total variable cost is calculated by multiplying labour hours from (4) by the real wage, and then adding in other variable costs. Because real wages have historically been quite flat, in the dynamic model, we fix the real wage rate at \$20.<sup>12</sup> We then regressed total variable cost from Lockheed's annual reports (see Figure 1) on total labour costs to obtain the two variable cost parameters listed in the table.<sup>13</sup>

Fixed costs are assumed to be constant across plane types and were also estimated using annual reports data. Because Lockheed produced very few L-1011's in 1977 and 1978, the maintenance costs of the L-1011 plants were reported in the annual reports for 1977 and 1978. These figures were converted to 1994 dollars and taken as the plant's fixed costs.

The only remaining cost parameters are those of the entry cost distribution. Lockheed developed the L-1011 at a cost of \$2.52 billion.<sup>14</sup> The only other available data is for the 747 (\$3.6 billion) and 777 (\$4.7 billion),<sup>15</sup> while development costs for the A-380 are forecast to be approximately \$10 billion.<sup>16</sup> It is widely argued in the industry press that in the last 20 years there has been an escalation of development costs. Hence, the L-1011 and 747 figures were

10. The exact figures are 22% error when using the aircraft's own production data to predict future costs vs. 25% error when using the industry-wide average curve.

11. Because we do not have accurate data on airframe weight, we measure the size ratio of the planes as the average between the ratio of seats and the ratio of volume.

12. Over the estimation period, the real wage rate varied between a minimum of \$18.03 in 1995 and a maximum of \$20.53 in 1978.

13. In support of our approach, our "fully burdened wage rate" of \$120 (=  $6.0 \times \$20$ ) compares closely with those used by the aircraft costing experts at RAND Corp. See, for example, Resetar, Rogers and Hess (1991, Table 20).

14. Source: author's calculations based upon Lockheed's annual reports.

15. Source: Tyson (1992). Figures converted to 1994 dollars.

16. Cole, Jeff. "Flight of Fancy: Airbus Prepares to 'Bet the Company' as It Builds a Huge New Jet", *The Wall Street Journal*, 3 November, 1999, A1. Figures converted to 1994 dollars.

TABLE 1  
*Cost parameters*

Parameter	Explanation	Value
A	Labour cost intercept	7.73 (0.01)
$\gamma$	Returns to scale	0.11 (0.17)
$\delta$	Depreciation of experience	0.613 (0.023)
$\theta$	Learning parameter	-0.63 (0.03)
	(Implied learning rate)	36%
W	Wage rate	\$20/h
FC	Fixed costs	\$200 million/year
TCF	Total variable cost/labour cost	6.0
TCC	Total variable cost intercept	36.2
	Cost/plane-size ratio	1.0
$x_1^l, x_1^h$	Type 1: entry cost distribution	\$2.5-\$3.5 billion
$x_2^l, x_2^h$	Type 2: entry cost distribution	\$3.3-\$4.6 billion
$x_3^l, x_3^h$	Type 3: entry cost distribution	\$4.4-\$6.2 billion

chosen as lower bounds for the entry cost distributions. Based upon this data, the entry cost for a type I (L-1011) aircraft was set to be uniform between \$2.5 and \$3.5 billion. The entry costs for type II and III aircraft were scaled up according to their size (see Table 1). The ranges used are consistent with the figures given above. Since the entry cost parameters are likely to be quite important to the dynamic model and there is very little data on entry costs, in Section 9 we investigate the robustness of our results to changes in these parameters. Note that by setting entry costs at a magnitude similar to those experienced by current aircraft producers, we are assuming that potential entrants are current aircraft producers. This seems like a reasonable assumption in view of the fact that the only outside entrants into the commercial aircraft industry in the last 30 years have been military aircraft contractors.

#### 4.2. Demand function

Because aircraft are differentiated products, a natural way to model aircraft demand is to use a characteristics-based discrete choice approach. However, one important feature of aircraft demand that differentiates it from standard discrete choice frameworks is that aircraft are durable goods. Therefore, rather than modelling aircraft demand as a static decision in a standard discrete choice framework, which we feel would be incorrect, we chose the opposite extreme. We instead 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. The aggregate market size, which is used in estimation and is a state variable of the dynamic model, is therefore all new and used planes in use in a given year.

Our approach, which amounts to treating aircraft purchases as rentals, relies on the fact that the market for used commercial aircraft is very efficient, with small transaction costs. There is substantial evidence in our data in support of this. In our data for wide-bodies (1969–1994), we found that 10–20% of the airline fleet changes operators each year. Approximately 40% of all individual aircraft in our data-set had more than one operator during the period, while 60% of

all aircraft more than 10 years old had more than one operator. Of those aircraft that changed operators at least once, the average aircraft had three operators. One aircraft changed operators 22 times between six airlines.<sup>17</sup> There are even seasonal shifts in the world aircraft fleet.<sup>18</sup> These facts suggest that the transaction costs for changing operators are low.

One drawback to the price data is that we only observe new aircraft prices, whereas our approach of treating aircraft as rentals suggests that we would instead want to use implicit aircraft rental prices. If the two are roughly proportional, this distinction should not matter as the difference would be absorbed into the price coefficient. Implicit rental prices should depend on interest rates, depreciation, and the expected change in the price of the good. Interest rates and depreciation can be considered to be roughly constant, but one worry is that expected price changes may potentially vary over time. While the use of new prices is therefore not perfect, in support of our using the new price data are two facts: first, new aircraft prices are remarkably constant over time, suggesting that expected price changes also do not vary much over time. Second, across product variation in rental prices is driven much more by variation in the levels of new prices than it is by variation in expected price changes.

We model yearly aircraft demand using a nested logit discrete choice model with several observable characteristics (number of seats, number of engines, etc.) and we also allow for one characteristic to be unobserved to the econometrician, which we call quality. All characteristics are assumed to be known to aircraft consumers. The unobserved product quality represents the unobservable aspects of an aircraft such as reliability, or suitability to current route structures. It is assumed to change over time, and is to be estimated. The product quality characteristic serves a role similar to that of a demand error and thus, with its addition, the demand model fits the data exactly.<sup>19</sup>

One consequence of using the nested logit model is that all individual aircraft purchases are assumed to be independent decisions even if undertaken by the same airline. This assumption is not likely to hold in this data. However, it is not clear that there are any consequences of making it. In a previous paper (Benkard, 1996) we estimated a multiple discrete choice model similar to that of Hendel (1999) on a more detailed micro-data-set on aircraft purchases and obtained results at the aggregate demand level that are similar to those here. The primary reason we do not use that model in this paper is that it would greatly add to the computational burden of the dynamic model, without clear benefits.

We allow for 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. The nested logit model is a great improvement over the standard logit model because it allows for airlines' preferences for the new wide-bodied planes to be correlated. 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. This feature of the model was found to be important in fitting the data.

A weakness of the nested logit structure is that it assumes that the correlation of utilities across all inside goods is the same. Thus, it may not be able to fully capture, for example, the uniqueness of the 747 relative to other products, except through the higher mean utility assigned to the 747. We discuss the implications of this property further in Section 5. We also rule out other dynamic effects such as switching costs and network effects, which we believe are small in commercial aircraft.

17. L-1011 serial number 1013.

18. Many aircraft in our sample change operator every spring and every fall and there is a worldwide shift of the operating fleet into Muslim countries yearly for the period of the Haaj.

19. See Berry (1994) for a discussion of this modelling approach.

TABLE 2  
Demand function estimates

Variable	Estimate	S.E.	Robust S.E.'s
Constant	-4.81	0.16	0.15
Seats/100	1.10	0.21	0.23
Freighter	2.45	0.24	0.26
No. of engines	-0.30	0.53	0.46
Price/100	-2.40	0.21	0.30
Last year dummy	-0.90	0.37	0.38
Trend	0.25	0.43	0.58
$\lambda$	0.77	0.18	0.18

Specification includes model dummies.

In the model, airlines' utility functions are,

$$u_{ijt} = x_{jt}\beta - \alpha p_{jt} + \xi_{jt} + \zeta_{igt} + (1 - \lambda)\varepsilon_{ijt} \quad (7)$$

where  $x_{jt}$  are observed characteristics of product  $j$  in period  $t$ ,  $\xi_{jt}$  is an unobserved characteristic of  $j$ ,  $\zeta_{igt}$  and  $\varepsilon_{ijt}$  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

$$\xi_{jt} = \ln(s_{jt}) - \ln(s_0) - x_{jt}\beta + \alpha p_{jt} - \lambda \ln(s_{jt/g}) \quad (8)$$

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 ( $\lambda$ ) 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 conditional moment restriction:

$$E[\xi_{jt} | Z_{jt}, \theta_0] = 0 \quad (9)$$

for a set of instruments  $Z_{jt}$ , consistent estimates of the parameter vector  $\theta$  are obtained through GMM using an optimal weight matrix.

Our instrumenting approach is now fairly standard in the literature. Instruments used include the observable plane characteristics, the hourly wage in manufacturing, the price of aluminium and the number of years a model has been on the market (to proxy learning). Note that under the assumption that  $\xi$  is first-order Markov, the number of years on the market is not correlated with  $\xi$ . However, if there is selection on  $\xi$  due to planes with low values of  $\xi$  exiting, then the moment condition would in general be violated. This problem is fundamental to the approach and is not specific to this paper. Because there is so little exit in the historical data (two data points out of 98), we feel that the selection bias should not be very severe for our data.

**4.2.1. Demand system estimates.** The demand system was estimated on demand data for the period 1975–1994. A total of eight models are observed over the estimation period, leading to 98 model-year observations. Parameter estimates are shown in Table 2. Robust standard errors were calculated because we felt that it was possible that the unobserved product characteristics were serially correlated and/or correlated with each other.<sup>20</sup>

20. The robust standard errors column refers to heteroscedasticity and autocorrelation consistent standard errors using the method of Andrews (1991).

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. Because of the high correlation among many aircraft characteristics, we only included a small number in the final specification. Only two of the parameters, the within-group correlation of utilities ( $\lambda$ ) and the price coefficient ( $\alpha$ ), are important to the dynamic model since the remaining parameters are aggregated into the two product quality states.

The within-group correlation of utilities is estimated to be high, at 0.77. This result indicates that new wide-bodies substitute much more highly with each other than they do with other aircraft. This is consistent with our expectations and was the reason for using the nested logit model. It is also important in driving the results of the paper as it implies high cross price elasticities between products. The high correlation estimate is driven by the following fact: while the market shares of individual aircraft vary substantially over time, the new wide-bodied aircraft fleet as a percentage of the total fleet in use has been very stable. Thus, the changes in the shares of individual planes are largely at the expense of other planes in the market, which implies that planes are highly substitutable. The standard error of 0.18 is on the large side, but we can easily reject the standard logit model ( $\lambda = 0$ ), and the point estimate is robust across specifications.

The coefficient on price ( $\alpha$ ) is estimated precisely and is significant and negative. Taken together, the two parameters lead to own price elasticities that average in the 4–10 range, with high cross price elasticities. These are higher elasticities than those typically found in other differentiated products industries, but we feel that they are consistent with the fact that aircraft are highly substitutable in practice and consistent with anecdotal accounts of the industry (*e.g.* Newhouse, 1982).

**4.2.2. Demand states and transition matrices.** For the purposes of the dynamic model, the plane characteristics used in estimation of the demand parameters are translated into the plane type and quality states as follows. First, three values for the plane type state ( $\mu$ ) were chosen to reflect typical small, medium and large aircraft. For those three types we simply chose the most representative aircraft of that type (L-1011, A-330, 747) and assigned to  $\mu$  the values of those planes’ observed characteristics. The three aircraft type states are listed in Table 3.

The estimated values of the quality state ( $\xi$ ) were then discretized to four points. Ideally we would have used more than this number, but to do so would have been computationally infeasible.<sup>21</sup> We then estimated the first-order stochastic process for the quality state nonparametrically using cell means.

It was also necessary to choose a discretization for the aggregate demand data. 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 was first de-trended to reflect 1994 values, and then discretized to three states. The removal of the trend makes all the state variables of the model finite and stationary,<sup>22</sup> 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 estimated nonparametrically using cell means and market size data for the complete history of the commercial jet aircraft industry, 1956–1994.

21. An earlier version of the paper used only three points, with similar results.

22. An alternative suggested by several seminar participants would be to allow some growth in the market that would eventually cease. We 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.



TABLE 3  
*Demand and other parameters*

Parameter	Explanation	Value
$\lambda$	Group corr. parameter	0.77 (0.18)
$\alpha$	Price coefficient	-0.024 (0.002)
$\mu$	Discrete plane types (small, medium, large)	{-2.6, -2.2, -1.6}
$P(\mu^e)$	Entry type distribution (small, medium, large)	(0.50 0.38 0.12)
$\xi$	Discrete plane qualities	{-0.90, -0.40, 0.11, 0.61}
$\Delta\xi$	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	(10,339 10,929 11,519)
$\Delta M$	Transition matrix for market size	$\begin{pmatrix} 0.895 & 0.143 & 0.000 \\ 0.105 & 0.786 & 0.200 \\ 0.000 & 0.071 & 0.800 \end{pmatrix}$
$\beta$	Firm's discount factor	0.925
$(\Phi^l, \Phi^h)$	Range of scrap values	(\$300m, \$700m)

4.3. Other parameters

Table 3 also lists values for the entry process. Due to computational limitations and a lack of data on entry occasions, in the empirical implementation of the dynamic model, we simplified the entry process from that outlined in Section 3. Instead of allowing entrants to choose which type of plane to enter with, entrants receive a random draw on what type of plane they can enter with as well as a random draw on development cost, conditional on that type. The entry decision is endogenous conditional on the draws received. The historical empirical distribution of plane entry types is used as the potential entry type distribution, listed in the table as  $P(\mu^e)$ . In the model, planes always enter at the second highest quality level because that corresponded to the observed entry value for seven of the eight planes in the sample.

The firm's discount factor,  $\beta$ , has been found by the past literature to be difficult to estimate. Thus, it was set to 0.925, which corresponds to a standard annual interest rate. Since the discount factor is an important dynamic parameter, in Section 9 we also solved the model for several alternative values of  $\beta$ .

The scrap value of a production facility was also difficult to measure. However, Lockheed reported significant detail on set-up 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 on whether the firm was going to continue producing other kinds of aircraft, the scrap value distribution was chosen to be close to \$500 million.

5. 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




FIGURE 3

Introductory P/MC ratios for a small aircraft entrant with three equal rivals

would have been relaxed to the point that it was no longer binding. However, computational limitations precluded that. The primary consequence of this restriction is that when the restriction binds the model predictions are slightly less competitive than they otherwise would be, with fewer firms and higher mark-ups. The model thus contains a total of 13 state variables with approximately seven million points in the state space.

### 5.1. Equilibrium pricing policies

Figure 3 graphs the equilibrium price–cost ratios for a newly introduced small 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 well below static marginal cost. 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: (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. (Cases (2) and (3) can be seen in Figure 3.) The strongest of the three effects is the learning curve. The model sometimes predicts high mark-ups in states where there are many high-quality competitors, if the competitors also have high cost, but always predicts low mark-ups 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. Thus, in general the introductory price–cost ratios predicted by the model match the industry qualitatively. The L-1011 is the only precisely measured data




FIGURE 4

Predicted vs. actual price/cost ratio for L-1011: 1972–1985

point for comparison. 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. Based on these results, our first finding is that the theoretical model is capable of rationalizing the extreme introductory discounts that are observed in the data.

We now compare the equilibrium pricing policies predicted by the model with those observed for the L-1011 over its entire product lifespan. 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.<sup>23</sup>

Perhaps surprisingly, the price–cost ratios predicted by the model are quite similar to those observed for the L-1011. The two series are very similar in overall shape and year-to-year variation, and the model predicts negative mark-ups for most of the period. The most notable discrepancy (see also Table 4 for a comparison of prices) between the two series occurs in the period immediately after the L-1011's introduction (1973–1975), where the model predicts higher mark-ups and prices than those observed. This overprediction suggests that the nested logit demand system may not be fully capturing the high degree of substitution between the L-1011 and the DC-10. For example, Newhouse (1982) suggests that these two planes were essentially perfect substitutes and thus undertook very fierce price competition. However, an alternative explanation is that during this period Lockheed was pricing at a level that was slightly lower

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, we used only the states of the own firm and the three strongest competitors. For the purposes of determining the actual historical state variables, the aggregate demand state was set to its closest discretized value.

TABLE 4  
*Predicted L-1011 prices and observed price range*

Year	Predicted average	Observed		
		Average modal	Min	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.

than optimal as suggested by the model. Mark-ups 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 mark-ups are not observable for other aircraft, comparisons of predicted vs. actual prices for other aircraft are similar.

We take away two conclusions from these results. First, we believe that these comparisons provide support for the model. The seemingly low level of aircraft prices relative to costs has previously been somewhat of a puzzle, particularly for such a concentrated industry. While the theoretical learning models have matched properties of the industry data qualitatively, past efforts at computing prices implied by parametrized versions of these models have not been successful at replicating observed prices (*e.g.* Baldwin and Krugman, 1988). Our results show that the generally low level of prices can be explained by a richer dynamic oligopoly model than was previously used. Another difference between this paper and the previous literature that may be important in matching industry pricing is that the equilibrium analysed in this paper is subgame perfect.

Our second conclusion is that Lockheed's pricing behaviour can be approximately rationalized by the model. The model suggests that Lockheed continually was in a position where there was substantial probability of future profits, which never materialized. This effect, and a strong incentive to move down its learning curve, justified Lockheed's aggressive pricing behaviour.

The model also predicts that the variance in price is much lower than the variance in cost, with mark-ups being inversely related to cost, as discussed in Section 2. The predicted price/cost ratio in Figure 4 looks similar to an upside-down graph of L-1011 costs. We present further evidence of this feature of the model in Section 6.

Note that our approach of comparing predicted with actual prices is a very rigorous test of the model. The model contains no direct information about prices beyond the parameter estimates in the supply and demand systems, and, because supply and demand were estimated separately, no information about mark-ups in the data was used in estimation. The prices and mark-ups in the model equilibrium are thus generated by the equilibrium structure of the model alone. 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 guarantee that prices will be predicted well.

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	0.55	0.75
Medium	0.23	0.21	0.23
Large	0.21	0.24	0.02

### 5.2. Industry dynamics

We now compare the model's predictions about industry dynamics with those observed in the data. Two simulations were used to calculate these statistics. The first simulation ("I.C. Simulation") shows statistics from 10,000 independent 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 plane. The second ("Invariant Distribution") shows statistics collected from a single 1,000,000 period simulation. 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. We do not believe that the historical period of 26 periods has been long enough to exhaust the memory of the process. Thus, we 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 that 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. In all cases the number of simulation draws was large enough that standard errors are negligible.

Tables 5 and 6 show actual (1969–1994) and simulated statistics for the wide-body market. 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, so the model does not make predictions about firm level concentration ratios for multi-product firms.

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

TABLE 6  
*Invariant distribution of plane values and lifetimes*

	(Invariant distribution only)			
	Median	Mean	S.D.	% positive
Distribution of plane values (\$):	1160	921	3943	60%
Distribution of plane lifetimes (years):	22	31	29	—

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.

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 well 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. However, this feature of the invariant distribution takes many periods to show up. Even in simulations as long as 200 periods, there are a high proportion of large planes due to the initial condition.

Note that these results are in part driven by the fact that the potential entry type distribution was parametrized to exactly match the observed distribution of entry types. 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 reflects the relative profitability of each product type as well as the potential entry type distribution. In fact, the simulated distribution of plane types is quite different from the potential entry type distribution (see Table 3), and the simulated distribution reflects the observed data closely, while the potential entry type distribution does not.

Unfortunately, there is no historical data available on observed plane values. Most companies do not give any public accounting of development costs and, furthermore, the majority of wide-bodied aircraft are still being produced today. However, the distribution of values generated by the model has several features that qualitatively match the industry. The model simulations show that the development of new planes is associated with a great deal of risk, which is also known to be true for the industry. Some planes lose a great deal of money, while others are very successful. The model also predicts that a majority (60%) of planes in the market are profitable. At least two (Newhouse (1982), Seitz and Steele (1985)) authors have claimed differently, but there is evidence that more products have been profitable recently than were in the past. The predicted median value of \$1160 million is within a reasonable range, but programme-level data is highly guarded and 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 appear to have been very profitable.

Product lifetimes in the invariant distribution are left skewed, with a median of 22 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. 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,<sup>24</sup>

24. 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 produced.




FIGURE 5  
Twenty-year simulation: prices

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 1 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.

## 6. REPRESENTATIVE TWENTY-YEAR SIMULATION

This section uses a typical 20 period industry simulation to display several important features of the model simulations. Figures 5–8 describe a typical 20-year 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 the model generates prices that generally do not reflect costs. An inspection of Figure 6 shows that costs generally follow the standard learning curve shape (despite the presence of forgetting in the cost equation). The first three entrants have slightly elevated initial prices due to the fact that they have few rival firms, and more importantly, no rivals that have reached the bottom of their learning curve, but once the industry reaches maturity in approximately period five, prices for all products are essentially flat, with slight year to year variations. In that sense, these graphs look much like the data in Figure 1, except that several of the firms are profitable.

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




FIGURE 6  
Twenty-year simulation: cost curves




FIGURE 7  
Twenty-year simulation: units produced

(see Figure 8). In that sense, this firm is reminiscent of the L-1011 and particularly Figure 1. 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 optimally, this firm receives






FIGURE 8

Twenty-year simulation: realized discounted cash flow of the firm

bad market realizations that cause it to make large losses totalling about \$6 billion dollars. Also similarly to the L-1011, with the benefit of hindsight, the firm's actions may appear to an outside observer to have been pathological, 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 (40%) these profits are not large enough to make the programme an overall success. See, for example, firm 5 in the simulation. Furthermore, 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.

## 7. INDUSTRY PERFORMANCE: ALTERNATIVE MARKET STRUCTURES

In this section, we compare the base market structure (MPE) with two alternatives: a multi-product monopolist (M) and a multi-product 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. Because there is only one agent, for each of the two alternatives there is a unique value function and associated policy function and the solution algorithm is a contraction mapping. Tables with statistics for the invariant distribution under each of the three market structures are available in a supplementary section.

According to the simulations, the MPE is quite efficient from a social perspective. On average, the SP increases total surplus by just 10% (\$17 billion) over the MPE. However, consumers are a great deal better off and producers a great deal worse off with the SP.

The monopolist, on the other hand, provides much lower social welfare than either the SP or the MPE, at great expense to consumers.<sup>25</sup>

While the SP does have the lowest production costs on average, surprisingly the welfare improvements from the SP are not driven primarily by lower marginal costs through learning as suggested by the theoretical models (*e.g.* Dasgupta and Stiglitz, 1988). Given the market size and the number of aircraft produced, in all three cases learning efficiencies are largely exhausted at the chosen output levels. Instead, welfare gains under the SP result primarily from more standard sources. The SP 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 competitive firms achieve this higher output primarily through a greater number of products, which results in the least efficient production (due to learning) of the three, confirming the findings of the theory literature. However, this cost effect is so small in practice that it is swamped by the output effect.

This is an interesting finding as it suggests that the welfare issues concentrated on by the theory literature are not the first-order effects. It suggests that, rather than worrying about productive efficiency, aircraft industry policy should instead attempt to ensure that the industry maintains high enough current output per product.

The second area of welfare savings under the SP results from concentrating output among just 2.4 firms on average, as compared with 3.8 in the competitive case, which leads to approximately a 30% reduction in new product investment. In the competitive case there is excess investment in development of new planes and wasted investment in learning to produce these products efficiently.

## 8. POLICY EXPERIMENT: RESTRICTING CONCENTRATION

In a recent article, *The New York Times* referred to Boeing as “essentially a government-sanctioned monopoly”.<sup>26</sup> 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, we 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 pay-offs 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.<sup>27</sup>

25. Note that it is well known that logit-based models tend to imply large amounts of welfare from each product. However, since our interest is in welfare comparisons across cases, and any tendency to overestimate welfare should be comparable across cases, the comparisons should be valid. Furthermore, the high price elasticities in our model should mitigate this tendency somewhat.

26. *New York Times*, “Aircraft Giant Tries to Recover From Hard Landing” (14 November, 1997).

27. In cases of outright monopoly, we assumed that the firm would be a regulated monopolist that would make zero current profit. In equilibrium such states were reached only 0.01% of the time so this assumption has no effect on the results.

TABLE 7  
*Statistics from 10,000 industry simulations under alternative policies*

Maximum concentration:	100%	60%	51%
Concentration ratios:	(Invariant distribution)		
1-Firm/plane	0.396	0.392	0.385
S.D.	0.102	0.094	0.081
2-Firm/plane	0.692	0.690	0.688
S.D.	0.109	0.107	0.103
Consumer surplus:			
Mean	135,373	134,917	133,895
S.D.	7040	7268	7488
Producer surplus:			
Mean	42,335	42,306	42,320
S.D.	3769	3776	3785
Total surplus:			
Mean	177,708	177,223	176,215
S.D.	10,441	10,645	10,832

Table 7 lists 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 10,000 simulations of 200 periods each.

The results show that the impact of the two alternative policies on the predicted distribution of concentration ratios is only slight. The effect of the concentration restriction policy seems to be largely limited to those states in which the policy actually binds, without having too great an impact on other states. Table 7 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 8 shows the distribution of welfare gains/losses across the 10,000 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 0.9%, or \$1.5 billion in present value terms.

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 policy has a negative effect on consumer surplus by reducing the production of the efficient dominant firm. There is also an opposing positive effect because 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 8 suggests that once all effects are accounted for, both alternative policies lead to welfare losses overall.

Due to the richness of the model, there are a wide distribution of results from the policy. While welfare losses occur in over 90% of the simulations, in the best outcome of the 10,000 simulations the 60% policy improves total welfare by approximately 12% (\$19 billion). The reason for this outcome is straightforward. There are some sequences of quality draws 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 receive good ones. In such sequences, a government

TABLE 8  
*Distribution of harm under alternative policies*

Maximum concentration:	60%	51%
Consumer surplus harm:		
Mean	457	1478
S.D.	1132	1998
% Positive	95%	94%
Producer surplus harm:		
Mean	28.5	14.6
S.D.	543	916
% Positive	31%	33%
Total surplus harm:		
Mean	485	1493
S.D.	1479	2580
% Positive	94%	93%
Mean consumer harm:	0.4%	1.1%
Mean producer harm:	0.1%	0.0%
Mean total harm:	0.3%	0.9%

policy that hurts this dominant firm in early periods and helps smaller ones is welfare enhancing because it helps to reduce the costs of the future dominant firms. Of course there is no way that a government could possibly foresee this occurrence. Moreover, there also exist sequences where the opposite occurs. In one case, the 60% restriction results in a 14% (\$22 billion) welfare loss. The conclusion that a concentration restriction would be welfare reducing thus holds only in expectation. As noted in Table 8, 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.

According to the surplus figures in Table 8, producers as a whole are hurt very little by the policy. 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 the airlines and the dominant firm to oppose the policy and weaker competing firms to support it.

Finally, note that in evaluating this policy we have implicitly kept firms' investment in product quality fixed because we have not changed the Markov processes for quality. 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.

## 9. ROBUSTNESS

In order to evaluate the robustness of the results, Table 9 shows simulations under several alternative parametrizations of the model. In addition to the base case, we solved the model with lower and higher discount factors, and lower and higher entry cost distributions. These are both important dynamic parameters for which there was little available data.

TABLE 9  
Invariant distribution under alternative parametrizations

Model:	Base		$x_1^l = B\$2.0$		$x_1^l = B\$3.0$
	case	$\beta = 0.90$	$\beta = 0.95$	$x_1^h = B\$3.0$	$x_1^h = B\$4.0$
Concentration ratios:					
1-Plane	0.40	0.44	0.39	0.39	0.40
S.D.	0.10	0.12	0.10	0.10	0.11
2-Plane	0.69	0.76	0.68	0.69	0.70
S.D.	0.11	0.12	0.10	0.10	0.11
Market size:					
No. of firms	3.8	3.3	3.9	3.8	3.7
S.D.	0.4	0.6	0.4	0.4	0.5
Avg. quantity per period:					
Small	168	171	115	142	192
Medium	60	38	92	81	33
Large	4	4	34	12	3
Avg. price per unit:					
Small	66.2	67.9	65.5	65.9	66.5
Medium	81.4	83.9	80.6	81.1	81.9
Large	104.6	107.3	103.4	104.1	105.5
Avg. MC per unit:					
Small	Lowest = 47.3	50.6	50.2	50.6	50.5
Medium	Lowest = 61.6	65.1	64.8	65.3	65.2
Large	Lowest = 83.1	86.9	86.2	87.3	86.7
Avg. (price/MC):					
	1.24	1.28	1.21	1.22	1.25
(Avg. price)/(Avg. TC):					
	1.19	1.24	1.15	1.18	1.19
New product investment per period:					
Mean	391	313	463	347	427
S.D.	74	64	88	66	81

The results do not change qualitatively across any of the five parametrizations. However, the table shows that the results are sensitive to the discount factor in one respect. When the discount factor is higher, there is more entry in general and particularly more entry of the larger aircraft types. The reason for this is that entry in the aircraft industry always involves up-front losses with offsetting profits that occur starting several years into the future. Higher discount factors help to make entry in general more profitable by increasing the weight placed on the future. For larger plane types the effect appears stronger in the simulations because the larger plane types were on average closer to the margin of profitability.

Another feature of the discount factor runs is that the higher the discount factor, the more competitive the market becomes, with lower mark-ups and prices. This happens both because there is more entry, and because with higher discount factors the future benefit of faster learning in the present has greater weight, which leads to greater present production and lower mark-ups. In our opinion, the most reasonable values of the annual discount factor are the  $\beta = 0.925$  and 0.95 cases. The only predictions of the model that change in the  $\beta = 0.95$  case vs. the base case are that the model predicts a higher percentage of large planes and slightly lower mark-ups.

The two parametrizations with lower and higher entry costs are very similar to the base case runs, with the expected effects of slightly more and slightly less entry, respectively.

## 10. CONCLUSIONS

This paper represents a first attempt at modelling an industry empirically using a dynamic oligopoly model with fully specified dynamics. We believe that this is an important problem for the I.O. literature. Such an approach is critical to understanding strategic interaction and optimal policies in many industries.

Despite several simplifying assumptions, we find that the dynamic model predicts many unusual aspects of the historical data well, particularly the persistent 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. This evidence provides tentative support for the use of this class of dynamic models in empirical work more generally as these features of the industry could not be replicated with a static model.

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. And yet, even in this case many simplifying assumptions were necessary and still the computational burden of the problem was found to be quite high. In order to apply the dynamic oligopoly framework to larger problems in the future, it may be necessary to find faster computational algorithms. However, there may also be 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 may include modelling 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 analysing various policies by simply re-solving the model with alternative institutions in place. We have evaluated three alternative market structures, with results suggesting that the single-product MPE is on average quite efficient from a social perspective. We also evaluated a policy that 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 that 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 modelling 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.

*Acknowledgements.* I would like to thank Ariel Pakes, Steven Berry, Donald Andrews and John Rust for their invaluable advice. I also thank Daniel Akerberg, Patrick Bajari, Tom Crawford, Dennis Epple, Gautam Gowrisankaran, Phillip Leslie, Michael Riordan, the editor, three anonymous referees and many seminar participants for their helpful comments. Financial assistance was provided by the Alfred P. Sloan Foundation.

## 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 EPPLE, D. (1990), “Learning Curves in Manufacturing”, *Science*, **247**, 924.  
 ASHER, H. (1956) *Cost–Quantity Relationships in the Airframe Industry*, R-291 (Santa Monica, CA: The RAND Corporation).

- BALDWIN, R. and KRUGMAN, P. (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. (1996), "A Simple Extension of Logit to Multiple Purchases and Application to Commercial Airframes" (Mimeo, Stanford University).
- BENKARD, C. L. (2000), "Learning and Forgetting: The Dynamics of Commercial Aircraft Production", *American Economic Review*, **90** (4), 1034–1054.
- BERRY, S. (1994), "Estimating Discrete-Choice Models of Product Differentiation", *RAND Journal of Economics*, **25** (2), 242–262.
- BERRY, S. and PAKES, A. (1993), "Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis", *American Economic Review*, **83** (2), 247–252.
- CABRAL, L. M. B. and RIORDAN, M. (1994), "The Learning Curve, Market Dominance, and Predatory Pricing", *Econometrica*, **62** (5), 1115–1140.
- DASGUPTA, P. and STIGLITZ, J. (1988), "Learning-by-Doing, Market Structure and Industrial and Trade Policies", *Oxford Economic Papers*, **40**, 246–268.
- ERICSON, R. and PAKES, A. (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work", *Review of Economic Studies*, **62** (1), 53–83.
- FERSHTMAN, C. and PAKES, A. (2000), "A Dynamic Oligopoly with Collusion and Price Wars", *RAND Journal of Economics*, **31** (2), 207–236.
- FUDENBURG, D. and TIROLE, J. (1983), "Learning-by-Doing and Market Performance", *Bell Journal of Economics*, **14**, 522–530.
- GOWRISANKARAN, G. (1999), "A Dynamic Model of Endogenous Horizontal Mergers", *RAND Journal of Economics*, **30** (1), 56–83.
- GOWRISANKARAN, G. and TOWN, R. (1997), "Dynamic Equilibrium in the Hospital Industry", *Journal of Economics and Management Strategy*, **6** (1), 45–74.
- GULLEDGE, T. R. and WOMER, N. K. (1986) *The Economics of Made-to-Order Production* (New York: Springer).
- HENDEL, I. (1999), "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns", *Review of Economic Studies*, **66** (2), 423–446.
- JUDD, K. L. (1996), "Cournot Versus Bertrand: A Dynamic Resolution" (Working Paper, Hoover Institution).
- KLEPPER, G. (1990), "Entry into the Market for Large Transport Aircraft", *European Economic Review*, **34**, 775–803.
- MASKIN, E. and TIROLE, J. (1988a), "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs", *Econometrica*, **56** (3), 549–569.
- MASKIN, E. and TIROLE, J. (1988b), "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles", *Econometrica*, **56** (3), 571–599.
- NEVEN, D. and SEABRIGHT, P. (1995), "European Industrial Policy: The Airbus Case", *Economic Policy*, 313–358.
- NEWHOUSE, J. (1982) *The Sporty Game* (New York: Alfred A. Knopf Inc.).
- New York Times* (1997) "Aircraft Giant Tries to Recover From Hard Landing" (14 November).
- PAKES, A. and MCGUIRE, P. (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 MCGUIRE, P. (2001), "Stochastic Approximation for Dynamic Models: Markov Perfect Equilibrium and the 'Curse' of Dimensionality", *Econometrica*, **69** (5), 1261–1281.
- REINHARDT, U. E. (1973), "Break-Even Analysis for Lockheed's Tri Star: An Application of Financial Theory", *Journal of Finance*, **28** (4), 821–838.
- RESETAR, S., ROGERS, J. and HESS, R. (1991) *Advanced Airframe Structural Materials: A Primer and Cost Estimating Methodology*, Prepared for the United States Air Force, R-4016-AF (Santa Monica, CA: The RAND Corporation).
- SEITZ, F. and STEELE, L. (1985) *The Competitive Status of the U.S. Civil Aviation Manufacturing Industry* (Washington, DC: National Academy Press).
- SPENCE, A. M. (1981), "The Learning Curve and Competition", *Bell Journal of Economics*, **12**, 49–70.
- TYSON, L. D. (1992) *Who's Bashing Whom? Trade Conflict in High-Technology Industries* (Washington, DC: Institute for International Economics).
- Wall Street Journal* (1999), "Flight of Fancy: Airbus Prepares to 'Bet the Company' as It Builds a Huge New Jet", (November) A1.
- WRIGHT, T. P. (1936), "Factors Affecting the Cost of Airplanes", *Journal of the Aeronautical Sciences*, **3**, 122–128.

Copyright of Review of Economic Studies is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.