

NBER WORKING PAPER SERIES

LEARNING AND FORGETTING:
THE DYNAMICS OF AIRCRAFT
PRODUCTION

C. Lanier Benkard

Working Paper 7127
<http://www.nber.org/papers/w7127>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1999

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 1999 by C. Lanier Benkard. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning and Forgetting: The Dynamics
of Aircraft Production

C. Lanier Benkard

NBER Working Paper No. 7127

May 1999

JEL No. L62, I11, L16

ABSTRACT

This paper introduces a new cost dataset for a commercial aircraft firm and uses this data to analyze the dynamics of learning in commercial aircraft production. This dataset is found to be inconsistent with the simple learning hypothesis, and particularly the prediction that a firm's unit cost must decline with its cumulative production. Instead, strong support is found for the hypothesis of *organizational forgetting*, a more general learning model where unit costs are similarly dependent on a firm's past production experience, but where that experience depreciates over time. Additionally, it is found that some, but not all, of a firm's production experience transfers from one generation of an aircraft to the next. This evidence adds to our understanding of productivity in industries with learning and thus has implications to many fields of economics.

C. Lanier Benkard

Graduate School of Business

Stanford University

518 Memorial Way

Stanford, CA 94305-5015

and NBER

lanierb@leland.stanford.edu

Learning and Forgetting: The Dynamics of Aircraft Production

C. Lanier Benkard*

Many past empirical studies have documented “learning-by-doing”, the hypothesis that unit cost decreases with cumulative production, in military aircraft production (T. P. Wright (1936), Harold Asher (1956), Armen Alchian (1963), Thomas R. Gullledge and Norman Keith Womer (1986)). Utilizing newly available production data for the Lockheed L-1011 Tri-Star, this paper studies commercial aircraft production, with an emphasis on the dynamics of production technology. Because commercial production is subject to many uncertainties not present in military production, the data presented here allows consideration of a richer set of hypotheses than was previously possible. In addition to learning, support is found for organizational forgetting (the hypothesis that the firm’s production experience depreciates over time), and incomplete spillovers of production expertise from one generation of an aircraft to the next.

This research is directed toward commercial aircraft for several reasons. The aircraft industry has many unique features — massive entry costs, dynamically increasing returns, imperfect competition, the fact that many countries consider it “strategic” — that make it important from a policy perspective, and the industry has frequently been the target of industrial policy, most notably in Europe. Dynamically increasing returns associated with learning make the industry a rare case in which infant industry protection may be theoretically valid, perhaps justifying such specialized policies.¹ A case has also been made for restrained antitrust policy based on the fact that concentrating learning is socially optimal.²

Aircraft production was also the first published case study of the learning curve, and remains among the most widely cited examples. However, despite this reputation, anecdotal evidence and discussions with industry executives have indicated that the marginal costs of producing aircraft do not always decrease over time, as would be expected if production was subject to pure learning, and a brief inspection of the data presented here (see Figure 1) confirms this fact. A goal of this paper is to reconcile these two conflicting stories and to show why they are both correct.

The evidence presented here also has important consequences. The organizational forgetting hypothesis in particular adds another degree of dynamic complexity to industries where learning is present. Hence, there are implications for many areas of industrial organization, including dynamic pricing and the evolution of the structure of an industry. These strategic outcomes in turn influence what industrial policies may be optimal, particularly in areas such as antitrust policy and strategic trade policy. While the current theoretical literature is relatively scarce with respect to the learning hypothesis,³ it is scarcer still with respect to organizational forgetting.⁴

At the macroeconomic level, organizational forgetting implies that recessions may lead to a reduction in productivity that lasts beyond the rebound in output. It also has implications for production scheduling, and leads to a particular kind of dynamic returns to scale in which greater production rates lead to greater productivity over time.

The hypothesis of incomplete spillovers implies that production costs increase generally when firms introduce new models of a product. Since competition often forces firms to produce

many variants of their products in order to secure sales (e.g., there have been over 25 different types of 747 produced in approximately 1000 units sold), there is reason to believe that the number of products in an unconstrained dynamic equilibrium may be too high from a social perspective. If that is the case, there would be further support for restrained antitrust policy, beyond that already called for due to the learning hypothesis.

Learning can also be thought of as the accumulation of human capital, which leads to an appealing interpretation of the two hypotheses considered. The organizational forgetting hypothesis suggests that organizational human capital depreciates, an assumption often maintained in other lines of the literature but thinly studied empirically. The extent to which production experience spills over across products or firms is a measure of specificity of the firm's human capital.⁵ This paper provides precise estimates of both the depreciation rate of the firm's experience and the rate at which experience spills over across aircraft models.

Finally, the data set used in this paper is quite special, uniquely suited to evaluation of the hypotheses of interest, and distinctive in the literature on several counts. It contains labor requirements data for each aircraft unit produced individually and over a period of highly variable production rates, both features that make it highly informative on the existence and extent of organizational forgetting. It is also the only data set that I know of to contain observations for several different models of a product produced simultaneously in the same plants on the same production line, making it also highly informative on the extent of learning spillovers.

I. Hypotheses on Aircraft Production

Past studies of learning in aircraft production have concentrated on military data because it was all that was previously available. However, in many ways military data is free of the dynamics that make commercial production interesting. A military producer typically gets a contract to produce a certain number of identical planes, N , over a certain number of time periods, T . Gulledge and Womer (1986) correctly characterize this as an optimal control problem in which the producer schedules production so as to minimize costs. Typically this results in production schedules that are relatively smooth over time and free from many of the uncertainties present in the marketplace. For commercial producers: orders are inconsistent due to both competition and business cycles, return on investment is not guaranteed, aggressive competition leads to the need to periodically upgrade models as well as to simultaneously produce several variants within each model generation, production rates may need to be changed quickly to fulfill sales agreements, and even within a given model variant every unit produced is customized. In essence, commercial producers experience more uncertainty, and a shorter planning horizon.

Because of the regularity of production in military programs, organizational forgetting and spillovers of production experience are less apparent. If forgetting is present, it may be very difficult to identify (e.g., data could be consistent with either a 20 percent learning rate or a 25 percent learning rate with 5 percent forgetting). And, in most cases there simply are not many model variants, so spillovers are not important.⁶ The Lockheed L-1011, on the other hand, was produced at a time when the commercial aircraft market was extremely volatile,

1970-1984. It was in fierce competition with another plane, the McDonnell-Douglas DC-10, that was very similar. As a result, there were four basic models of L-1011 in only 250 units produced and production rates varied widely.

A. *Learning*

It was long ago realized that aircraft production was characterized by strong learning effects (Wright (1936), Alchian (1963), and Asher (1956)). Wright (1936) noticed that labor, material, and overhead requirements declined with cumulative production. His estimates led to wide acceptance of the “20 percent learning curve” in the engineering literature, which has the interpretation that production input requirements reduce by 20 percent for every doubling of cumulative past production. He also introduced the learning curve specification that is now typical of the economics literature:

$$(1) \quad L_i = AE_i^\theta$$

where L_i is labor input per unit, A is a constant, E_i (“experience”) is cumulative past output, and θ is a coefficient describing learning.⁷ Alchian (1963) ran similar learning curve regressions on WWII military data for 22 plane types in four major classifications and came to the conclusion that the learning rate varied across each plane type. In a more recent study, Gullledge and Womer (1986) link the two ideas of the production function and the learning curve by allowing for various kinds of scale effects and profit maximization in an optimal control setting. They apply their model to data for several military programs and find that

it fits the data well. Learning has also been documented in many other industries, e.g. aircraft engines, machine tools, metal products, nuclear power plants, chemical processing, semiconductors, and shipbuilding.

Learning may take on many different forms depending on the particular nature of production. In more capital-intensive industries such as chemical processing and semiconductors, learning primarily results from the fine tuning of production techniques. In such industries, engineers and managers analyze current output and constantly make small changes to the process with the result that productivity gradually improves. In labor-intensive industries such as aircraft and shipbuilding, learning primarily results from workers becoming more efficient at the tasks they perform through multiple repetition. Many industries may be subject to both types of learning. The exact nature of learning may also influence the transferability and persistence of experience.

Because the unit of output is so large, commercial aircraft production is highly labor intensive and production rates are very low.⁸ Learning thus results primarily from a more experienced work force.⁹ In recognition of the importance workers' experience plays, aircraft producers are very careful to maintain stability in their manufacturing process over time.

B. Organizational Forgetting

Organizational forgetting is the hypothesis that a firm's stock of production experience depreciates over time. Since an aircraft firm's experience is embodied in its workers, it seems likely that turnover and layoffs may lead to losses of experience. The traditional learning

hypothesis does not allow for this possibility. Linda Argote, Sara L. Beckman and Dennis Epple (1990) and Argote and Epple (1990) suggest that based on rapid job turnover, highly variable production rates, and anecdotal evidence that production costs did not uniformly fall over time, the L-1011's situation was consistent with the hypothesis of organizational forgetting. Argote, Beckman and Epple (1990) also find support for the organizational forgetting hypothesis in WWII ship production.¹⁰

Experience may depreciate in times of falling production rates since typically such times are accompanied by layoffs. During subsequent increases in production, the firm is often unable to acquire the same workers that it formerly released, and must completely retrain new ones. Normal rates of employee turnover may also lead to experience depreciation during periods of constant production, particularly if turnover rates are high.

The term “forgetting” may seem somewhat inappropriate when referring to an organization, since under the explanation given above it may be that no individual is forgetting anything. Furthermore, the data in this paper is only able to distinguish “forgetting” at the organizational level without accounting for what fraction is due specifically to individuals. However, since aircraft production rates are so low (individual assemblies are worked on for as long as four weeks) it is entirely plausible that forgetting may occur both at the individual and organizational levels.¹¹

An implication of the organizational forgetting hypothesis is that recent production is more important than more distant past production in determining a firm's current efficiency. This prediction is perhaps more intuitive than that of the learning hypothesis, which treats all

production equally no matter how old. In a practical example, it is hard to imagine that Boeing's rapid production of 747s in the early 1970's is as important to current unit costs as production in the early 1990's, particularly because it is unlikely that many of Boeing's workers from that period remain with the company today.

The organizational forgetting hypothesis is especially relevant to aircraft production due to a unique contract that producers have with the International Association of Machinists and Aerospace Workers (IAM) regarding seniority called "displacement rights" or "bumping rights". As hourly employees attain seniority, they have the option of requesting an upgrade into a higher level job classification if one should become available. At the same time, the company is bound to fill positions from within the existing ranks whenever possible, even if it means retraining. Of course, if an opening is filled from within then the position vacated by the worker upgraded is subject to the same rules. Similarly, when there are layoffs, the higher seniority workers can bump down lower seniority workers, and so forth down the ranks. Whenever employment changes, there is a domino effect throughout the company that has been known to affect as many as ten positions for just one job opening. This "bumping" has the greatest impact on increases in employment since it could mean that the company must retrain its existing workers in their new positions as well as the usual training of new workers. Under "bumping", turnover rates are high, which leads to greater organizational forgetting.¹²

C. Experience Spillovers

The question of to what extent experience spills over across firms (external) or products (internal) is one of how specific the firm's production experience is. Because of their implications to trade policy, external spillovers have been a major topic in semiconductors production (e.g., Douglas A. Irwin and Peter J. Klenow (1994), William W. Nye (1996)). Internal spillovers, or spillovers across programs or plants within the firm, also have strategic implications, particularly in aircraft production due to firms' joint military and commercial production. The European Union frequently defends its huge subsidies to the Airbus Consortium on the basis that U.S. commercial firms are implicitly subsidized by their large military contracts. This argument hinges partly on the existence of internal spillovers across programs.

These are important issues that are, however, not easily analyzed with the data presented here, which covers just one aircraft program. This paper instead looks at to what extent experience spills over incompletely across different models of an aircraft within one overall program. If it is true that the skills required to build a new model of an existing aircraft do not transfer over completely from past production, then upon introducing the new model, the firm experiences a setback in learning and higher production costs for the whole aircraft program. Furthermore, the extent to which experience is specific to similar models within one program represents an upper bound on the extent to which experience is specific across different programs.

Despite there now being only two firms in the industry, the commercial aircraft industry is characterized by extreme competition. One way that firms compete is by aggressively

tailoring their product to customer's needs, leading to their providing many different models of every aircraft they produce.¹³ If learning does not spill over completely between these models, then production costs for the whole program are higher as a result. It seems likely, then, that in an unconstrained dynamic equilibrium there may be too many models from a social perspective. This argument thus offers additional support for restrained antitrust policy in this industry, beyond that already called for due to learning.¹⁴

II. Data

Underlying this research is a new dataset for L-1011 production made available to the author by Lockheed that shows how labor requirements varied across the 250 L-1011s produced between 1970 and 1984 (see Figure 1).¹⁵ The data refers to direct man hours incurred by Lockheed California in the production of each plane including detail fabrication in Burbank, Burbank assembly, Palmdale final assembly and flight test.¹⁶ It does not include major subassemblies that were subcontracted such as the complete wing with all flight controls and systems, complete empennage, S-duct, doors, etc.¹⁷ Unfortunately, at the present time very little data (annual reports, newspaper articles etc.) is available to document the cost of other inputs to production such as capital investment and materials. Due to high development costs typical of the industry and the fact that the L-1011's price rarely exceeded even marginal cost, the L-1011 was phased out in the early 1980's for a total loss of several billion dollars.

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

flight date, which is taken to be the plane's date of completion.¹⁸

With these two data sets, it was possible to compile the production schedule and sales pattern of the L-1011 across time and models (Figures 1 and 2). The first model produced, the L1011-1, comprised 65 percent of total production and was produced during the whole period. The L1011-100 and -200, which together account for just 15 percent of production, utilize the same fuselage as the -1 and are otherwise quite similar to produce, the difference being in accessories only (fuel tanks, landing gear, engines, etc.).¹⁹ The -500 on the other hand has a shorter fuselage with two fewer doors and a different cargo compartment configuration and galley, differences significant enough that it would not be unreasonable to expect an impact on production. The -500 accounts for the remaining 20 percent of L-1011s and was produced from 1978-1984.

Figure 1 graphs direct labor requirements and yearly output for the L-1011. This graph emphasizes the differences between this dataset and the military datasets used in the existing literature on learning curves. While military production is not perfectly smooth, L-1011 production rates varied sharply (4-42 planes per year) over the period 1970-1984 reflecting the chaotic state of commercial air traffic in the 1970's.

The labor requirements data (Figure 1) has a fairly classic learning curve shape up until the overall trough in man hours at unit 112, where 226 thousand direct labor hours went into production of the plane. At this low point, which occurs in the fifth year of production, the direct labor requirement is about one-sixth that of the first plane.

The most striking feature of the data is the upturn in labor hours for units produced after number 112. Much of the empirical work presented here is devoted to an explanation of this seeming departure from the learning curve.

The most obvious explanation for this feature would be some kind of supply disruption such as a strike or materials shortage, but that explanation does not hold up in this particular case. There was in fact an IAM strike during the last three months of 1977. However, as noted above, this strike took place during the tail end of a recession in commercial air travel when production rates were very low anyway, and hence had little impact on Lockheed's production schedule. The planes produced during this strike correspond to the very first part of the upturn on the graph in Figure 1 (approximately units 148-150). The second peak in labor requirements occurs at plane number 181 which was completed three years later in March 1980. Though it is entirely plausible that the strike could have caused an upturn in labor requirements, there is a great deal of evidence that production technology was not permanently influenced by the outcome of the strike.²⁰

Two other possible explanations for the increase in labor requirements are labor hoarding and scale effects. If this were a labor hoarding story then it would be expected that the peak in labor requirement would occur near the trough in output. Instead, the peak in labor requirement occurs at the second peak in output. While scale effects are investigated in more detail in the estimations that follow, it is apparent in Figure 1 that over the first half of the data there is an increasing returns relationship, while over the second half of the data the relationship is the opposite: output and labor requirements grow together.

III. A Model of Commercial Aircraft Production

A. *Production Function*

This section concentrates only on the production function itself, but makes assumptions with consequences that extend to the underlying dynamics. For an example of a dynamic model of the commercial aircraft industry that is consistent with the assumptions made here, see C. Lanier Benkard (1998).

Consistent with practice in the industry, it is assumed that firms must build a plant prior to beginning production and that from then on the firm maintains the plant but cannot change its size.²¹ That assumption transfers to a production function with fixed capital, denoted \bar{K} . The firm's single-product production function is assumed to take a Leontief form in variable inputs labor (L) and materials (M):

$$(2) \quad q = \min \left(G \left(L, E, \bar{K}, S, \epsilon \right), H \left(M, E, \bar{K}, S, \nu \right) \right)$$

where E is the firm's "experience", a state variable that determines the firm's current production technology, S is line-speed (or current production rate), ϵ and ν refer to productivity shocks to labor and materials respectively, and $G(\cdot)$ and $H(\cdot)$ are functions to be described later.

The "line-speed" variable is a measure of current output rate and measures within period scale effects. It is included primarily for comparison with previous work (Werner Z. Hirsch

(1952), Gullledge and Womer (1986)) and because its importance could not be entirely rejected in the estimations that follow. However, several cautions must be raised. First, line-speed is a choice variable and is hence endogenous. Second, the relationship between line-speed and production in a fixed capital model is not obvious. For very low production rates, it is presumed that an increase in line-speed would lead to an increase in productivity. However, since capital is fixed, it must be the case that at high production rates diminishing returns set in.²²

The Leontief form of (2) was chosen for several reasons, the most important being that it is a reasonable approximation of the industry. By far the greatest input to production next to labor is engines, which can not substitute with labor since they are fully subcontracted. In addition, the same strict adherence to a production process discussed above in the context of keeping the capital stock constant over time also applies here, limiting the scope for input substitutability. The second reason is that many other candidate specifications such as Cobb-Douglas in all inputs would place strong restrictions on the joint substitution of experience with labor and materials which seem unlikely to hold in practice. Note also that the importance of the Leontief assumption to the results is quite limited since, despite the fact that exact materials inputs data is not currently available, there are straightforward ways of testing the substitutability of labor and materials by adding proxies for materials prices to the labor requirements equation.

A unique feature of aircraft production is the large unit size of production (one plane), which leads to very low production rates and a production technology that is noticeably different (due to learning) for every unit produced. Production technology changes over time as

learning progresses. Additionally, building an aircraft is such a large undertaking that firms can adjust variable inputs at the unit level. Fortunately, the data also contains observations at the unit level. Hence, consider the efficient production frontier along the surface of $G(\cdot)$,

$$(3) \quad T(q, L, E, \bar{K}, S, \epsilon) = 0$$

Since firms adjust the variable inputs at the unit level, it is as if they are always operating along the unit isoquant ($q = 1$). Assuming that the efficient frontier takes the Cobb-Douglas form and rearranging terms gives:

$$(4) \quad \ln L_i = \ln A(\bar{K}) + \theta \ln E_i + \gamma_0 \ln S_i + \epsilon_i$$

The parameter θ above measures the rate of learning (the *learning rate* is calculated as $1 - 2^\theta$), while γ_0 measures the within period returns to production rate. The i subscript, omitted prior to equation (4), indicates the unit number of the product being produced. More specific assumptions on the E_i and ϵ_i processes follow below. ²³

Equation (4) implies a stable relationship between line-speed and labor requirement which goes against the discussion of line-speed above. Hence, in estimation more flexible functional forms were also used to allow for a more general relationship.

The timing assumptions on the production process defined above are as follows: Line-speed is set at the beginning of the period so that output equals demand for the period. At any time the firm can produce a plane according to the technology in (2). When it produces a plane, it first observes the shocks ϵ_i and ν_i and then injects the required amount of labor and materials

to produce exactly one unit. The exact technology used is determined by the firm's current experience level, the current line-speed, and productivity shocks. The timing of observing the shocks before injecting the variable inputs is a reasonable approximation of the fact that during the 7-15 months it takes to produce a plane, anything that goes wrong requires increased labor (and/or materials) to right. This increased input requirement is unforeseeable and unavoidable unless the firm were to scuttle the unit entirely in mid-production — an unlikely prospect given the size of the investment involved.

B. *Experience*

In the traditional learning model, experience is defined very simply as cumulative past output:

$$(5) \quad E_i = E_{i-1} + 1, \quad \text{with } E_1 = 1.$$

In order to incorporate the hypotheses described above, a new specification is used. The first goal is to incorporate organizational forgetting. In that regard, this paper uses the specification introduced by Argote, Beckman and Epple (1990) which has experience depreciating by a factor δ each month. Experience in a given month is then the cumulative depreciated production experience of all production that took place prior to that month:

$$(6) \quad E_t = \delta E_{t-1} + q_{t-1}, \quad \text{with } E_1 = 1.$$

There are many ways to justify this specification theoretically.²⁴ I prefer an abstract interpretation of experience in which the company can only invest in experience by producing

and this experience depreciates at a constant rate due to a variety of factors including geographic turnover, company turnover, literal forgetting, and nonreinforcement of routine (see Chapter 5 of Nelson and Winter (1982)). If these various effects occur at constant rates, the specification yields a constant overall discount factor.²⁵ It may also be desirable to allow the discount factor to vary over time. This generalization is discussed further in the estimation section of the paper.

Just as (5) is a reduced form for learning, (6) is only a reduced form representation of organizational forgetting. The forgetting hypothesis hinges upon very micro-level employment effects such as turnover, layoffs, and literal forgetting that are not accounted for separately. Instead, output rates and the timing of output are used to proxy the underlying learning/forgetting processes. Within the model, learning can only take place by producing planes, and unless production rates are maintained, experience begins to depreciate. These effects are regardless of any underlying employment patterns. In actuality, it seems likely that more than one time series of employment could lead to the same set of outputs, while simultaneously resulting in a different path for production technology. For instance, it may be within the power of the firm to raise the value of δ , presumably at some cost. While this is certainly an interesting prospect that may point toward optimal employment strategies, it is a subject for further research. The goal of this paper is simply to document the presence of organizational forgetting, and to understand the dynamics of technology in aircraft production. Additionally, the reduced form specification (7)-(9) is found to lead to extremely stable results, justifying its use.

The second addition to the model allows for experience spillovers to be incomplete across models. To test that hypothesis, the four L-1011 types were divided into two overall classifications. The -1, -100, and -200 were joined together under the heading -1 since they share the same fuselage. The -500, which has a slightly different fuselage, is kept separate under its own heading.

With these considerations, the new experience variable is defined as follows:²⁶

$$(7) \quad E_i = \begin{cases} E_{1,t} & : \text{ if } i \text{ is type } -1, -100, -200 \\ E_{500,t} & : \text{ if } i \text{ is type } -500 \end{cases}$$

$$(8) \quad E_{1,t} = \delta E_{1,t-1} + q_{1,t-1} + \lambda q_{500,t-1} \quad \text{and} \quad E_{1,1} = 1$$

$$(9) \quad E_{500,t} = \delta E_{500,t-1} + q_{500,t-1} + \lambda q_{1,t-1} \quad \text{and} \quad E_{500,1} = 1$$

$E_{1,t}$ refers to the firm's experience in producing type -1s and determines its production technology for type -1s at any given time. $E_{500,t}$ is defined similarly. Thus, the firm now has two experience state variables where it used to have one. λ is the experience spillover parameter. The case where $\delta = 1.0$ and $\lambda = 1.0$ corresponds to no organizational forgetting and complete spillovers respectively. In that case, E_i is essentially the same as (5).²⁷

Note that (7)-(9) assumes some cross equation restrictions. Several alternatives to this specification exist, none of which are supported by the theoretical arguments given previously.²⁸

One such alternative would be to allow for spillovers to be asymmetric. However, asymmetric spillovers would be contrary to the theoretical intuition that experience spills over due to workers becoming proficient at tasks common to both aircraft models. This alternative

hypothesis was nevertheless tested and the data did not support it.²⁹

C. *Estimation*

Equation (4) can be estimated using OLS if ϵ_i is *iid* and independent of E_i and S_i . However, this would not be an innocuous assumption. Past research (e.g., G. Steven Olley and Ariel Pakes (1996)) indicates that serial correlation in the unobserved portion of firm productivity is important. In that case, both E_i and S_i are correlated with ϵ_i as follows: If ϵ is not independent across observations, then as indicators of future productive efficiency, past values of ϵ are state variables for the firm. As such, they influence the firm's choices of prices and line-speeds, S_i . Since E_i is a function of past quantities, which are determined in part by past prices, which are in turn a function of past ϵ 's, it too must be correlated with ϵ_i .

As a result, it is necessary to instrument for these correlations in order to regain consistency. For the traditional learning model (5), where E_i is not a function of parameters and is thus observable, two stage least squares (2SLS) is sufficient. However, in the other two models, it is necessary to use a nonlinear estimator.

The nonlinear estimator used here is based on the standard literature for estimating GMM models with a conditional moment restriction as described in Lars P. Hansen (1982). It departs from that literature primarily in estimating the covariance matrix of the moment conditions, which is used in calculating both standard errors and an optimal weight matrix. To account for serial correlation without making any specific assumption as to its form, a nonparametric heteroskedasticity and autocorrelation consistent (HAC) estimator of the

covariance matrix suggested by Andrews (1991) is used.³⁰

D. *Instruments*

Both the 2SLS and GMM estimators require exogenous instruments. As above, the two variables being instrumented are line-speed and experience. Under the assumptions of the model, line-speed should be highly correlated with current output, while experience should be highly correlated with recent output. Hence, good instruments (\mathbf{Z}) for this model include present and lagged demand and cost shifters. To be valid instruments, the variables used must be exogenous in the sense that they are not choice variables for the firm.

Demand shifters used include various world GDP measures (world, OECD, third-world), the price of oil, and a time trend. For the GDP measures, one- and two-year lags were used, while for oil prices lags up to four years were included. The longer lags on oil prices were designed to instrument for demand for the more efficient model (type -500), the development of which significantly lagged the initial oil price shock.³¹

Cost shifters used include the world aluminum price and the U.S. manufacturing wage rate. For both cost shifters, the one- and two-year lags were included. It is assumed that in both cases, Lockheed amounts to a negligible amount of demand so that both variables are truly exogenous to Lockheed's productivity shocks. Due to high industry concentration, it would not be appropriate to make a similar assumption regarding more industry specific prices, e.g., the aircraft industry wage rate.

If the Leontief assumption in equation (2) were relaxed and instead production was assumed to be Cobb-Douglas in all inputs, then production only at the unit level would imply a labor requirements equation similar to (4) with the addition of both the wage rate and materials prices. In that case three of the instruments, the manufacturing wage rate which is highly correlated with the industry specific wage rate and the world prices of oil and aluminum which ought to proxy materials prices, may also belong as regressors. If these instruments belong as regressors but are excluded, consistency is lost if they are included in the instrument set. This discussion implies two straightforward tests of the validity of the Leontief assumption. First, the value of the GMM objective function provides a simple test of the overidentifying restrictions. This test should fail in the event that the Leontief assumption is false. Second, a perhaps more powerful test would be to include input prices in the production function as would be correct in the Cobb-Douglas case and see if that has any impact on the results.

IV. Results

A. *Traditional Learning Model*

For comparison purposes, consider first the traditional learning hypothesis specification. For the whole of this section, experience is defined as cumulative past output as in equation (5). See also Tables 1-2.

Regressions 1 and 2 show estimates for the traditional learning model in its most basic form for both the first half of the sample and the entire sample. The results for the first

112 units are very similar to those found in the literature (e.g., Aircraft: Alchian (1963); Semiconductors: Irwin and Klenow (1994), Harald Gruber (1992)). The implied learning rate is slightly higher than average at 30 percent, degree of fit is very high, and the learning parameter, θ , is estimated very precisely.³² The primary reason for running this restricted regression is to show that the first half of the sample is consistent with previous results. However, a brief visual inspection of data is enough to tell us that this model will not fit the full sample well and that expectation is born out in regression 2 (see also Figure 3). Thus, it seems that the traditional learning hypothesis explains the first half of the data well, but falls short when applied to the full sample.

The hypothesis maintained in this paper is that this discrepancy is caused by the fact that the firm's experience is not being fully retained over time, which only becomes apparent when production rates are uneven and new models are introduced. However, before testing that hypothesis specifically, I first show that the discrepancy is not easily explained if we maintain the traditional learning hypothesis.

Regression 3 adds line-speed and line-speed squared³³ to the equation in order to allow flexibly for within period returns to scale. Both terms are found to be significant, and the estimated relationship implies that returns to scale are increasing at a decreasing rate in the observed output range. While this relationship is intuitive, it nevertheless adds little explanatory power to the model as evidenced by the high *SSR*.³⁴

For comparison with some of the existing literature (e.g., Gruber(1992), Martin B. Zimmerman (1982)), regression 4 tests the hypothesis that learning takes place over calendar time

as well as from production experience. Despite improved fit and the small standard error on the time variable, the model is rejected due to the sign of the time variable, which implies technological change is negative.

I also tested several model extensions which attempt to directly explain the increase in labor requirements after unit 112. The first of these, regression 5, is a test for the presence of adjustment costs in labor and includes two variables representing the change in line-speed, one for a positive change and one for a negative one. The adjustment costs specification was tried because some contemporary newspaper articles imply that Lockheed incurred great difficulty in increasing production after 1978.³⁵ This preliminary evidence weakly supports that adjusting production upward increases labor requirement whereas adjustment downward decreases it. However, this specification also has very little explanatory power, and the result disappears when the learning model is generalized.³⁶ Another possible explanation for the upturn in labor requirements is the IAM strike of 1977, but several specifications testing this were tried with no success.³⁷

The regressions in Table 2 test the robustness of the production function specification. The cleanest regression under the traditional learning hypothesis is regression 6 and comes from adding a “scope” dummy, which equals one for all units produced after the introduction of the type -500, the period of joint production.³⁸ The coefficient on scope is significant and positive, indicating strong diseconomies of scope, and the fit of the model is also much improved. Remaining misspecification occurs primarily for units 140 onward.

Regressions 7 and 8 include input prices as regressors as would be implied if the production

function were Cobb-Douglas. The positive sign on the aluminum price is as expected, but the strong positive sign on the wage rate implies that a higher wage leads to much greater use of labor, which seems to make little sense. Furthermore, if the scope dummy is added to this regression the wage rate coefficient becomes close to zero and insignificant, suggesting that the result is spurious. To resolve this matter and especially since the aluminum price was significant in both regressions, this issue will be reconsidered in the context of the general learning model.

Note that throughout this section the significant and positive estimates of ρ_ϵ throughout Tables 1 and 2 are as much an indication of misspecification in these models as they are of serial correlation so they are not addressed until the next section.

B. *General Learning Model*

Regressions for the general learning model are shown in Table 3. The significant and positive estimates of ρ_ϵ in all regressions indicate that, while serial correlation is significantly reduced from before, it is still present. Hence, NLLS is inconsistent due to endogeneity of the right hand side variables. Additionally, while the HAC-IV estimates were found to be statistically significantly different from those of the NLLS model — the NLLS estimations were rejected — their interpretation was essentially the same and therefore the inconsistent NLLS results have been omitted. There was also little difference between the point estimates using the HAC weight matrix and the more standard serial independent weight matrix (both are consistent) in the GMM estimation. The only differences were that the HAC technique gave slightly

larger standard errors and acceptance of the χ^2 test of the overidentifying restrictions in all cases.

Regression 9 generalizes the learning model to allow for organizational forgetting ($0 \leq \delta \leq 1$), but still enforces that spillovers are complete ($\lambda = 1$). This regression shows that organizational forgetting alone is very powerful in explaining the labor requirements data. With the addition of only one parameter, the *SSR* falls from 12.9 to 2.9. In addition, the depreciation parameter, δ , is estimated extremely precisely and is significantly different from one in all cases, strongly rejecting the hypothesis of no forgetting.

The estimate of δ rises significantly when the model is extended to allow for incomplete spillovers (regression 10, see also Figure 4). This is because the introduction of the -500 occurs near the time when forgetting is causing labor requirements to increase and incomplete spillovers help to explain the increase in labor requirements. Leaving out incomplete spillovers from the model thus leads to a downward bias in the experience depreciation parameter. The addition of the spillovers parameter also improves the fit of the organizational forgetting model to the point of being nearly perfect. Most of the remaining specification error is present only for units 1-10, and is most likely due to discretization errors in the model that result from modeling experience as accumulating all at once upon completion of a unit rather than smoothly over the whole production of the unit. However, even over this range the model correctly captures the timing of two local maxima and two local minima in the data caused by early production delays.

It is particularly noteworthy that the model explains both halves of the data well. It accounts

successfully for the rise in labor costs in the latter half of production while still explaining the first half of production, which was shown above to be consistent with the simpler learning model. Moreover, the model fits both turning points in the observed data. In comparison with the diseconomies of scope regression, the organizational forgetting model captures the overall productivity trend much better from unit 140 onward. The scope regression misses the early upturn in labor requirements as well as its subsequent peak. Importantly, for approximately two years after its introduction, production of the type -500 becomes progressively *less* efficient even though the scope of production was unchanged and production rates were increasing. It is this feature of the data, along with the increasing labor requirements for the type -1 planes, that supports the organizational forgetting model and rejects diseconomies of scope. Similarly, the fact that labor requirements for the two models converge over time supports that learning spillovers are incomplete.

Despite the fact that the estimate of the depreciation parameter for this data is much higher and more precise than previous results obtained by Argote, Beckman, and Epple (1990) and Darr, Argote, and Epple (1995), the implied depreciation rate still seems quite high. The value $\delta = 0.96$ implies that 61 percent (0.96^{12}) of the firm's stock of experience existing at the beginning of a year survives to the end of the year.³⁹ This high rate of depreciation is in part justified by low aircraft production rates. Some tasks performed by workers are repeated only rarely. However, it is helpful to keep in mind that we are not measuring simply the depreciation in an individual's skills, which would be unlikely to be that high, but the depreciation in a firm's stock of a very specific kind of human capital. Thus, a high depreciation rate is also justified by the effects of turnover, remembering that the "displacement rights" contract

provision serves to multiply the effects of company level turnover. Lastly, in the past it has proven difficult to measure human capital and its depreciation, so any intuition we may have about it is likely driven by either theoretical considerations or estimates of depreciation for other kinds of capital, neither of which are likely to give us very accurate or precise estimates.

One reason for the stability of the organizational forgetting model in this dataset is that aircraft producers, and Lockheed in particular, do not make significant technological changes to the production process once production is underway. Hence, the data is not influenced very much by factors not accounted for in the model, as may occur in other industries. What is perhaps most interesting is that organizational forgetting may also explain the lack of technological change itself. In discussions with industry executives they have expressed the belief that disruptions in production, even those designed to improve efficiency, may lead to setbacks in productivity since they upset workers' routines. Companies thus choose to maintain current processes for as long as possible in order to retain production experience.⁴⁰

One consequence of allowing for depreciation in the model is that the learning rates in Table 3 (35-40 percent) are much higher than those estimated under traditional learning hypothesis. The reason for this is that learning is no longer relative to cumulative production, but is now relative to accumulated experience, which is constantly depreciating. The interpretation of the learning rate is thus not as simple as before since the current production rate (monthly output) also matters. The new learning rate implies that if experience were doubled, then labor requirements would fall by 35-40 percent. However, whether or not this reduction is even attainable depends on the current experience level and planned production rates.

In all of the regressions (except 11 – see below) the spillover parameter, λ , is estimated to be significantly different from 1.0, so the hypothesis of complete spillovers is rejected. Given that the two L-1011 models are quite similar and were produced in the same plants together, the spillover parameter estimates were surprisingly low, suggesting that aircraft production experience is very specific. The parameter estimates imply that approximately 70 percent of production tasks overlap between the two models. This number can also be assumed to be an upper bound on the amount of spillovers that would take place across two different aircraft programs and suggests that the level of cross program spillovers may be quite low. This finding is consistent with the fact that labor requirements are usually very large upon introduction of a new aircraft program.⁴¹

The results in regression 10 were also tested for robustness against all of the alternatives mentioned in the previous section with similar results (see Table 4). Regressions 11-13 show that input prices and diseconomies of scope are rejected in the general learning model, though there is clearly some collinearity between the scope dummy and the level of experience spillovers as the standard error on the spillovers parameter is quite large in regression 11. The spillovers parameter changes value somewhat in regressions 12 and 13 and is generally the least stable parameter of those estimated, but interpretation of the results is not significantly altered. The results were also tested for robustness against a variety of other alternatives. The strike and adjustment cost variables were included and found to generate small insignificant coefficients. The line-speed coefficient, γ_0 is found to have a small amount of explanatory power, but the restriction that $\gamma_0 = 0$ does not substantially change the results. When included, the squared term in line-speed (γ_1) and the time variable produced coefficients that were very

small and insignificant.

Using the parameter estimates in Table 3 it is possible to quantify the impact of the L-1011-500 on Lockheed's variable cost of production. If we assume values of $\lambda = 0.70$ and $\theta = 0.63$, then the first -500 produced required approximately 25 percent ($= 0.70^{-0.63}$) more labor than a -1 would have required. If Lockheed chose to produce both -500s and -1s in approximately equal numbers for the rest of the program, then labor requirements would continuously run approximately 11 percent ($= ((1.0 + 0.70)/2)^{-0.63}$) higher for each model than in the single model case.⁴²

These figures show that due to incomplete spillovers, the decision to bring out a new (albeit similar) model can involve a significant setback in learning, and an associated large and immediate increase in variable costs. Moreover, simultaneous production of more than one model causes variable costs to be continuously higher than they would be for the same production rate of a single model.⁴³ When research and development expenses are added in, it becomes evident that introducing new models is a costly endeavor even within an existing aircraft program.

V. Alternative Production Schedules

This section considers the implications of the general learning model to productivity through alternative production schedules. The organizational forgetting model implies a kind of dynamic increasing returns to scale that is somewhat different from that in the traditional

learning model. In the organizational forgetting model the number of units produced recently is most important in determining productivity, so that greater production rates can lead to greater average productivity. In the traditional learning model only the total number of units produced matters and hence production rates do not change average productivity for a production run of fixed length. Thus, in the general learning model we should expect to find that some bunching of production is optimal, much as producing all output in one period would be if production were subject to static increasing returns.⁴⁴ On the other hand, production scheduling is irrelevant in the traditional learning model with respect to a production run of fixed length.

Three counterfactual production schedules are considered. The first is a smooth production schedule of 19.8 units per year, the average output for the L-1011, but leaves the set of models produced unchanged. The second uses the same smooth production schedule as the first, but assumes that only one model of plane is produced. The third further assumes that production rates are doubled so that production occurs in half as many years. Figure 5 shows actual direct labor requirements versus direct labor requirements for the three alternatives. Note that per unit direct labor requirement is the inverse of labor productivity.

Of the four schedules shown, no one schedule is best for every unit produced. In particular, the actual production schedule is best for most units in the range 80-120. However, Table 5 shows that alternatives 3 and 4 both result in both higher average productivity and higher productivity in expected discounted value than what actually occurred. On the other hand, smoothing production over time would have increased labor requirements by both measures as expected.

Figure 5 also shows that in a smooth production schedule, the labor requirements curve generated by the organizational forgetting model appears similar to one generated by the traditional learning model, the only difference being that it asymptotes to a positive value. This feature of the model may explain why previous studies of military production have found support for the traditional learning model. A comparison of lines 3 and 4 in Figure 5 suggests that Alchian's (1963) conclusion that learning rates varied across different aircraft programs may result from the use of too narrow a learning model. In the more general learning model presented here, these apparently different learning rates may be explainable through either varying production rates across the different programs or varying levels of cross program experience spillovers.

None of the four schedules shown is likely to be the optimal production schedule given demand for the product. Lockheed's actual production choices may have been optimal ex ante but, even if that were true, they are unlikely to have been optimal ex post. Calculation of the ex post optimal production schedule using the learning model estimates would require further assumptions about inventory costs. I do not attempt to make this calculation here since these costs are difficult to observe and any assumptions I might choose to make would thus be somewhat arbitrary. In the absence of inventory costs, the production schedule that would minimize the EDV of cost according to the model estimated is one in which production is run essentially at capacity from start to finish. Interestingly, most military production runs are of this form. If inventory costs are large enough, then the optimal production run looks much like the actual one, which closely matched scheduled deliveries.⁴⁵

VI. Concluding Remarks

Previously, commercial aircraft and other markets with learning have been analyzed using theory that is based on smooth learning curves. The evidence presented in this paper suggests that, in actuality, the dynamics of production are much more complex. In order to better understand these industries, the strategic effects of organizational forgetting need to be considered. This is an area of theory that has, to my knowledge, not been studied.

At the same time, it seems unlikely that forgetting is important in all industries where learning takes place. This brings up the issue of what properties of the firm, and of production, are most important in causing the depreciation of experience. Commonalities between aircraft, ship, and service franchises — the only three industries where forgetting has yet been documented — include: the products are labor intensive, learning is thought to be important at the individual worker level, and all three experience relatively high turnover. These may be criteria for evaluating how important organizational forgetting is to an industry. However, that hypothesis remains to be tested.

Additionally, it seems plausible that the rate of experience depreciation is under the firm's control. For example, avoiding layoffs may be in the firm's interest even if workers are temporarily underutilized. In the event that layoffs are unavoidable, it may be in the firm's interest to make some provision for reacquiring laid off workers in future periods. Further study with some additional data on employment and turnover might point toward ways that management could maximize experience and productivity by minimizing loss of experience. A full understanding of the causes of forgetting may also lead to a better understanding of

its macroeconomic effects.

The learning model estimated here (and elsewhere) assumed a deterministic learning process, but it seems more likely that learning actually takes place stochastically, i.e., upon executing a task sometimes something is learned and sometimes not. However, a stochastic version of this model was estimated using simulation techniques and yielded almost identical results. In hindsight, this result is quite intuitive given that the unit of output actually consists of many individual tasks. If learning occurs stochastically at the individual task level, then the learning process at the unit level would have very low variance and thus resemble a deterministic process.

Finally, the results above show that, in addition to their development costs, new models can have a great impact on variable production costs across an entire aircraft program. Therefore, it seems likely that in an unconstrained dynamic equilibrium firms may produce too many models. A social planner or multiproduct monopolist would internalize the business stealing effects of product differentiation and may produce fewer models to avoid these higher costs.

Table 1: Traditional Learning Model Regressions

	$\ln A$	θ	γ_0	γ_1	Time	Adj.+	Adj.-	SSR	ρ_ϵ	$L.R.$
Basic Regressions:										
1. Units 1-112	7.90	-0.51	-	-	-	-	-	1.36	0.73	30%
	(0.06)	(0.01)							(0.04)	
2. Units 1-238	7.16	-0.29	-	-	-	-	-	15.0	0.97	18%
	(0.08)	(0.02)							(0.02)	
Line-Speed:										
3.	6.51	-0.35	0.95	-0.20	-	-	-	11.0	0.92	21%
	(0.21)	(0.02)	(0.17)	0.03					(0.03)	
Calendar Time:										
4.	6.03	-1.08	-0.04	0.004	1.16	-	-	5.4	0.56	53%
	(0.15)	(0.04)	(0.13)	(0.025)	(0.07)				(0.04)	
Adjustment Cost:										
5.	6.75	-0.31	0.67	-0.16	-	0.07	-0.04	15.8	0.43	19%
	(0.26)	(0.03)	(0.25)	(0.05)		(0.03)	(0.02)		(0.06)	
$N = 238$	$TSS = 33.7$									
NOTES: All regressions are 2SLS. Instruments (Z_i) are present and lagged demand shifters (various world GDP measures, the price of oil, and a time trend – see text) and present and lagged cost shifters (U.S. wage rate, aluminum price). $L.R.$ is the implied <i>learning rate</i> .										

Table 2: Traditional Learning Model Regressions: Input Prices and Diseconomies of Scope

	$\ln A$	θ	γ_0	γ_1	Wage	P_{AL}	P_{oil}	Scope	SSR	ρ_ϵ	<i>L.R.</i>
Diseconomies of Scope:											
6.	7.35	-0.49	0.49	-0.10	-	-	-	0.55	2.4	0.70	29%
	(0.10)	(0.01)	(0.08)	(0.02)				(0.02)		(0.04)	
Oil Price:											
7.	5.88	-0.54	1.36	-0.27	-	-	0.27	-	9.3	0.83	32%
	(0.21)	(0.03)	(0.17)	(0.03)			(0.04)			(0.04)	
Input Prices:											
8.	-15.9	-0.52	0.45	-0.09	8.68	0.50	-	-	10.0	0.81	30%
	(3.37)	(0.03)	(0.18)	(0.03)	(1.32)	(0.09)				(0.04)	
$N = 238$	$TSS = 33.7$										
NOTES: All regressions are 2SLS. Instruments (Z_i) are present and lagged demand shifters (various world GDP mea- sures, the price of oil, and a time trend – see text) and present and lagged cost shifters (U.S. wage rate, aluminum price). <i>L.R.</i> is the implied <i>learning rate</i> .											

Table 3: General Learning Model Regressions

	$\ln A$	θ	γ_0	δ	λ	SSR	GMM(p)	ρ_ϵ	$L.R.$
OF Only:									
9. [$S_N^* = 9.3$]	7.63	-0.65	0.14	0.952	-	2.9	0.60	0.51	36%
	(0.01)	(0.02)	(0.12)	(0.003)				(0.05)	
Spill-Overs:									
10. [$S_N^* = 6.9$]	7.73	-0.63	0.11	0.960	0.70	2.3	0.62	0.45	36%
	(0.01)	(0.03)	(0.17)	(0.003)	(0.07)			(0.05)	
$N = 238$	$TSS = 33.7$								
NOTES: All regressions in this table use the HAC-IV method described in the text. Instruments (Z_i) are present and lagged demand shifters (various world GDP measures, the price of oil, and a time trend – see text) and present and lagged cost shifters (U.S. wage rate, aluminum price). S_N^* is the optimal bandwidth used in estimation the GMM covariance and optimal weight matrices. L.R. is the implied <i>learning rate</i> .									

Table 4: Alternative Hypotheses

	$\ln A$	θ	γ_0	δ	λ	Wage	P_{AL}	P_{oil}	Scope	SSR	GMM(p)	ρ_ϵ	<i>L.R.</i>
Scope:													
11. [$S_N^* = 7.7$]	7.91	-0.71	0.13	0.954	0.73	-	-	-	-0.03	3.1	0.62	0.45	39%
	(0.01)	(0.02)	(0.10)	(0.004)	(0.17)				(1.92)			(0.05)	
Oil Price:													
12. [$S_N^* = 9.0$]	7.72	-0.72	0.17	0.946	0.79	-	-	-0.003	-	3.3	0.89	0.51	39%
	(0.01)	(0.02)	(0.05)	(0.002)	(0.05)			(1.13)				(0.05)	
Input Prices:													
13. [$S_N^* = 7.9$]	7.63	-0.70	0.14	0.952	0.77	0.06	0.009	-	-	2.9	0.78	0.46	38%
	(0.13)	(0.01)	(0.08)	(0.003)	(0.04)	(6.36)	(2.96)					(0.05)	
$N = 238$	$TSS = 33.7$												
NOTES: All regressions in this table use the HAC-IV method described in the text. Instruments (Z_i) are present and lagged demand shifters (various world GDP measures, the price of oil, and a time trend – see text) and present and lagged cost shifters (U.S. wage rate, aluminum price). S_N^* is the optimal bandwidth used in estimation the GMM covariance and optimal weight matrices. <i>L.R.</i> is the implied <i>learning rate</i> .													

Table 5: Alternative Production Schedules

	Unit Labor Requirement	
	Average	EDV
1. Actual	381.5	62440
2. Smooth Production (20/yr)	385.1	63570
3. One Model (20/yr)	373.7	62330
4. One Model (40/yr)	324.6	55860

Figure Titles:

Figure 1: L-1011 Production: Direct Labor Requirement and Yearly Output

Figure 2: L-1011 Production by Model

Figure 3: Traditional Learning Curve: All 238 Units (log-log)

Figure 4: Organizational Forgetting and Incomplete Spill-Overs (Regression 10)

Figure 5: Labor Requirements: Actual and Three Alternatives




Figure 1:




Figure 2:




Figure 3:




Figure 4:

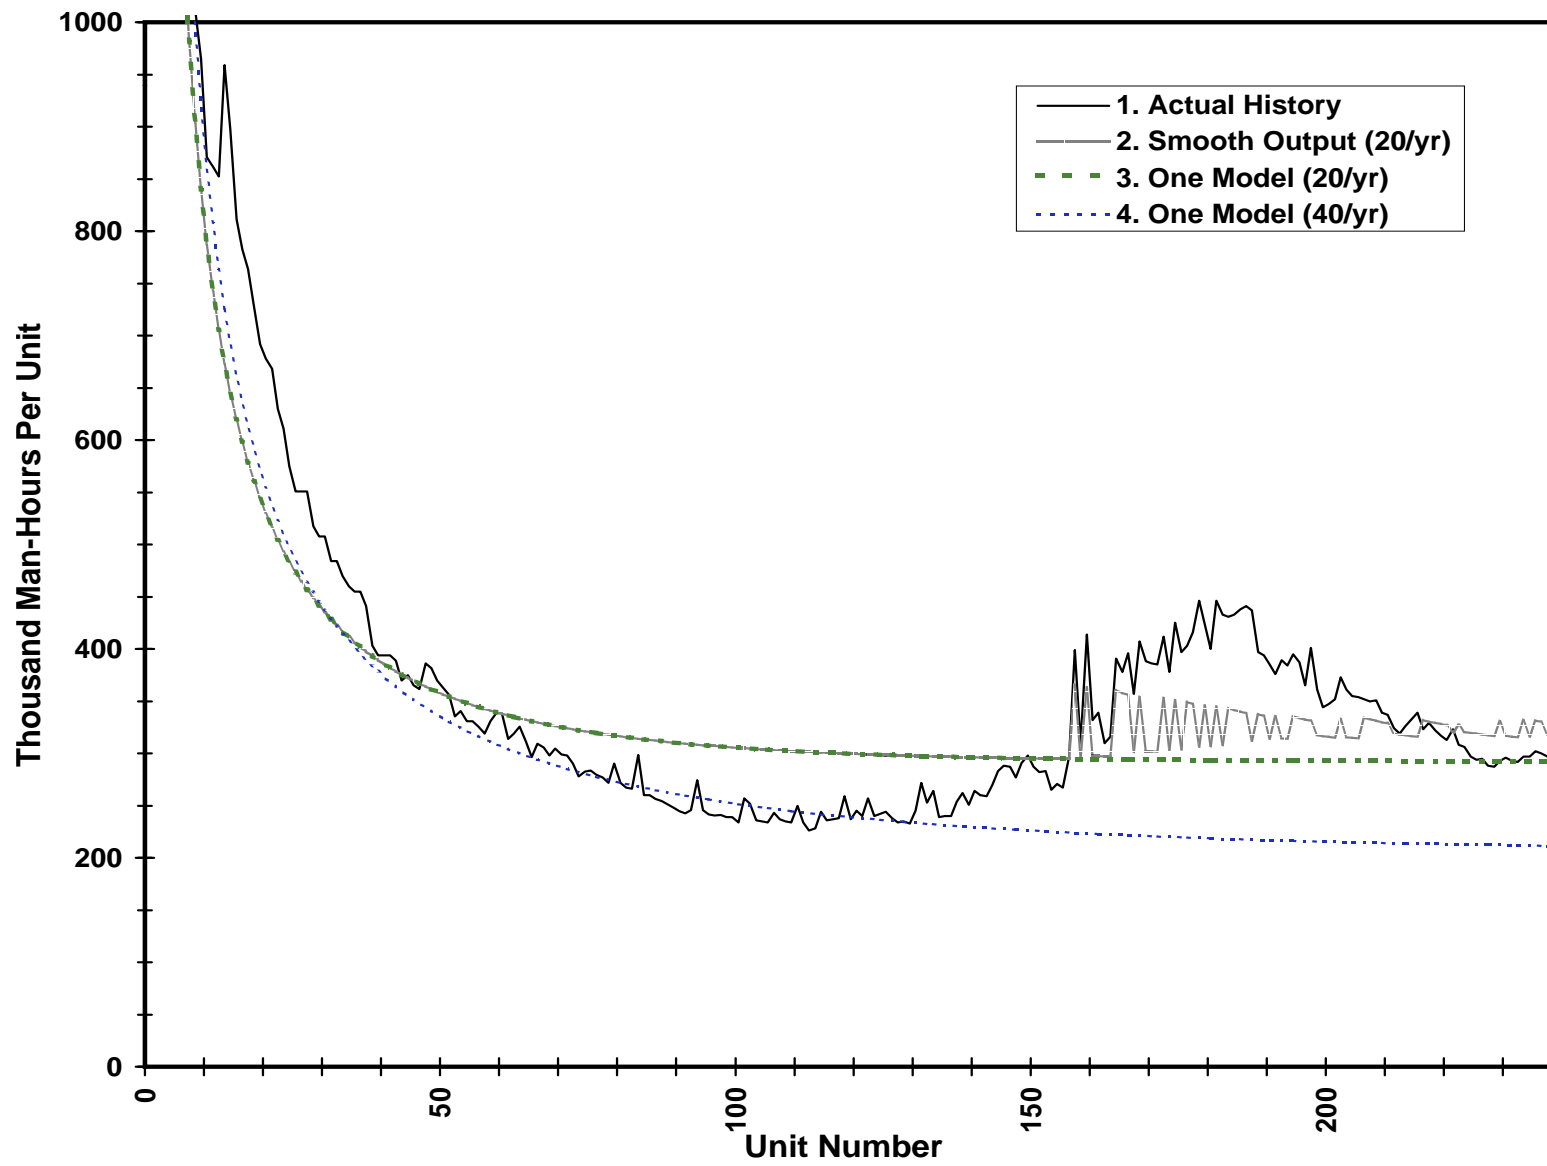


Figure 5:

Notes

*Graduate School of Business, Stanford University, Stanford, CA 94305, e-mail: lanierb@leland.stanford.edu. I would like to thank Ariel Pakes, John Rust, Steven T. Berry and Donald W. K. Andrews for their invaluable advice. I also thank Daniel Akerberg, Michael Riordan, Christopher Timmins, George Hall, two anonymous referees, and many seminar participants for helpful comments, and especially Tom Crawford at Lockheed for providing many important insights into the industry. Financial assistance from the Alfred P. Sloan Foundation is gratefully acknowledged.

¹Richard Baldwin and Paul Krugman (1988) discuss this issue in the context of the European Airbus subsidies. Partha Dasgupta and Joseph Stiglitz (1988) do not refer to aircraft specifically, but do prove the result that, in the presence of learning curves, there are circumstances in which an import ban may be welfare enhancing.

²See Dasgupta and Stiglitz (1988) for a theoretical argument and conditions under which restrained antitrust policy is preferred. Additionally, the Department of Justice chose not to block the recent Boeing–McDonnell-Douglas merger, a move that would be hard to imagine in other industries given the market shares involved.

³Luis M. B. Cabral and Michael Riordan (1994) give a good summary of the current theory on strategic interaction under learning.

⁴To my knowledge, the strategic implications of organizational forgetting have not been studied.

⁵Machines can be either very specific in their use, such as a die which can only be used

to fashion one specific part, or very transferable, such as a hand-held welding device which can be used to join any number of parts. The same can be said for workers' skills. The skill of operating a welding device is very transferable, but the intricate details of welding a particular assembly may not be.

⁶Gulledge and Womer have one example (F102A/TF102 program) in which there are two models produced but the two models are very similar and hence the authors do not attempt to account for the differences.

⁷The learning rate is calculated as $1 - 2^\theta$. A 20 percent learning rate corresponds to $\theta = -0.32$.

⁸The highest selling commercial jet airliner ever, the Boeing 737, has sold less than 3000 units in eight models over 32 years. It is more typical to sell fewer than 1500 units, especially for larger planes. Among wide-bodies, the Boeing 747 is the only plane to have sold more than 1000 units, and it has been through four model generations over 28 years.

⁹This belief has been represented to me in many conversations with industry executives. It is also supported in many current and past anecdotal accounts, e.g., *The New York Times*, November 14 1997 and *Aviation Week & Space Technology*, October 15 1979.

¹⁰Eric Darr, Linda Argote, and Dennis Epple also found support for organizational forgetting in franchises.

¹¹Richard R. Nelson and Sidney G. Winter (1982) refer to an individual's "skills" (chapter 4) and an organization's "routines" (chapter 5). The organizational forgetting hypothesis relates to their claim that organizations "...remember a routine largely by exercising it..."

much as an individual "...remembers skills by exercising them." (p. 99). See also Charles D. Bailey (1989) for an empirical study of individual rates of forgetting, and Jacob Mincer and Haim Ofek (1982) for a study of the depreciation of human capital at the individual level.

¹²The existence of such a contract provision may in itself seem contrary to the importance of learning. However, it should be noted that this contract, which originated prior to the unions, most importantly serves to provide workers with insurance against being laid off in recessions.

¹³Most aircraft come in many different configurations which usually vary in seating capacity and range. Examples of individual model tailoring are also extensive and include everything from cabin configuration, engine type, and optional fuel tanks down to prayer rugs, and custom galleys and lavatories, e.g., a recent article in *The New York Times* (November 14 1997) claimed that Boeing offers a choice of 37 different cockpit clipboards.

¹⁴Dasgupta and Stiglitz (1988) show that under certain circumstances, the learning hypothesis implies that unconstrained monopoly may be socially preferable to oligopoly.

¹⁵Complete data for all levels of production is not available for the last 12 planes produced and hence they are excluded from this study.

¹⁶The data provided by Lockheed includes a breakdown for each of these parts of the production process. However, for a majority of the planes, data for the first stage of production (accounting for about one-third of direct labor hours) is only available by lot. Hence, for the purposes of this paper only the total direct labor requirement is used.

¹⁷The level of subcontracting was consistent throughout the production run. The primary

subcontractor was Lockheed Georgia (the primary L-1011 producer was Lockheed California). I have been unable to acquire additional production data for any of the portions of production that were subcontracted.

¹⁸For reference, the time to completion for the L-1011 ranged from 7-15 months depending on the overall rate of production. Delivery usually took place soon after the flight test except in a few isolated occurrences.

¹⁹Because the differences between the -1 and -100 and -200 models were in accessories only, conversion was a simple task and was ordered by several airlines for their existing fleets upon introduction of the -100 and -200. For the purposes of this exercise, conversions are ignored.

²⁰Lockheed managers have indicated (consistently with concurrent articles in the press) that the strike was initiated by management in trying to break the “displacement rights” contract provision described previously. The strike was settled when management gave in and the settlement did not involve any significant contract changes.

²¹Though normal practice is to build a plant prior to production and to essentially leave the plant as is during the entire production run, there has on occasion been a plant change-over in mid production. Lockheed relocated production of the P-3 Naval Patrol Aircraft twice, once in mid-production with a “disastrous affect” on the program, and the second time between two separate production runs with minimal effects. In the case of the L-1011, however, production consistently took place in the same plants for the entire production run. When production rates were low, these plants were producing below capacity.

²²Since plant size is fixed, aircraft firms increase production rates by adding workers to the

existing line and increasing the line speed. Increasing returns to line-speed at low production rates are thus likely to result from worker specialization. In the extreme case, one can imagine that one worker assembling the entire plane would not be very proficient at any of the individual tasks. However, because the plant size is fixed, adding more workers would eventually lead to overcrowding and inhibit production.

²³Since we only observe production at the unit level, the returns to labor are not identified in this model and have been normalized out.

²⁴One rather mundane way is to define E_i as the (normalized) sum across tasks of the number of times a current employee has executed that task, where production of an aircraft is assumed to be broken up into a certain number of small “tasks”. Then, suppose that the company loses a constant fraction of its employees each month and that the average experience of lost employees is a constant fraction of the average experience of current employees (e.g., 1).

²⁵Note that assuming that a company invests in a regional experience pool and that geographic migration occurs at a constant rate seems more plausible than assuming that company level turnover is constant.

²⁶To be clear, the sequential (as opposed to recursive) representation of equations 7-8 is: $E_{1,t} = \sum_{s=0}^{t-1} \delta^{t-s-1} q_{1,s} + \lambda \sum_{s=0}^{t-1} \delta^{t-s-1} q_{500,s}$. The sequential representation of $E_{500,t}$ is analogous.

²⁷The only difference between (7)-(9) vs. (5) in the case of $\delta = \lambda = 1.0$ is that experience accumulates monthly instead of unit by unit.

²⁸It should be noted here that due to the very large scale of production, as is typical in the industry, L-1011 production narrowed to one line in its later stages. This provides a unique environment in which the same workers in the same plants sequentially work on different models.

²⁹Another interesting alternative to (7)-(9), the idea for which I owe to Dennis Epple, would have initial condition $E_{500,t} = \Phi E_{1,t}$ for the period in which the first -500 was produced, and then follow (7)-(9) for the remaining planes. This alternative specification would allow separate estimation of the Φ and λ parameters and would have very different implications to (7)-(9) if it were found that $\lambda \neq \Phi$. However, estimation of this specification strongly supported the hypothesis that $\lambda = \Phi$, which is consistent with the “task overlap” intuition for experience spillovers and further supports the use of (7)-(9).

³⁰In particular, Andrews’ Quadratic-Spectral (QS) kernel example is used since he shows that an estimator based upon the QS kernel is best in terms of an asymptotic truncated mean square error criterion among kernels that necessarily generate positive semidefinite estimates. The bandwidth was chosen optimally using a parametric first stage. The resulting GMM estimator is efficient under the class of IV estimators with instruments \mathbf{Z} , but does not obtain the asymptotic efficiency bound due to the nonlinearity of the experience variable. Details on all aspects of estimation are available from the author upon request.

³¹In the serial independent case the optimal instrument for the spillover parameter would be highly correlated with the ratio E_1/E_{500} , and thus it is desirable to include instruments which separately capture demand for the two different models.

³²The standard errors listed do not account for serial correlation and are hence likely to be too low. However, these results are comparable with others in the literature.

³³Because line-speed is not directly observable, in all of the estimations that follow, line-speed is measured as the annualized number of planes produced in the period beginning three months before the current month and ending three months after. Alternatives to this definition that were considered included varying the period length and using the simple annual output rate. Results were nearly identical in all cases.

³⁴Note that these regressions are not orthogonal projections and hence the use of the *SSR* as a measure of fit is only approximate.

³⁵See articles from *The Wall Street Journal* and *Aviation Week & Space Technology* listed in the references.

³⁶Several other versions of the adjustment cost test suggested by seminar participants were also tried involving various specific functional forms of the rate of change in output. The one shown contains the strongest results found.

³⁷See previous revisions of the paper for details.

³⁸This specification may not be the most intuitive test for scope economies, but it is the specification that gave by far the strongest results. Several other specifications were also tested that actually measured the scope of production continuously.

³⁹In the two previous examples in the literature, the depreciation parameter was estimated to be in the range 0.75-0.85, suggesting that only 5-15 percent of experience was retained

from year to year.

⁴⁰Despite this argument and the stability of the depreciation parameter in the estimations, several attempts were made to allow the depreciation rate to vary over time. Several seminar participants and an anonymous referee suggested allowing depreciation to be a function of output or the change in output, but this was not supported by the data. Other specifications tried included high order polynomials in time and splitting the sample into subperiods. Some support for the latter specification was found with the sample split into two at approximately the 90th unit. However, the p-value in a test of equality of the two depreciation parameters for the subsamples was only 0.014, which is surprisingly high given the precision of the estimates, and the true (unconditional on the split point) p-value of such an arbitrary search must be much higher.

⁴¹High initial labor requirements and relatively fast learning rates imply that only a small amount of the firm's production experience from other models is relevant to current production.

⁴²An alternative specification was tried that lends further support for the hypothesis that simultaneous production is costly. That specification, which is discussed in footnote 29, attempted to separate experience gained toward production of the -500 from previous production of the -1, from spillovers of experience during simultaneous production. The two spillover rates were estimated to be equal.

⁴³The extent of the increase in costs would likely vary in each instance depending on how many of the production tasks overlap. The spillover parameter estimates shown here only

apply to this particular case.

⁴⁴This is only true if the depreciation rate of experience is not a function of current output rates, as was found to be the case here. If it were true that higher production rates led to less depreciation then it may instead be optimal to smooth production.

⁴⁵In some simple experiments I have found that including a small inventory cost on the order of the interest cost of the investment in the plane is enough to make the actual production run very close to ex post optimal.

VII. References

- Alchian, Armen. "Reliability of Progress Curves in Airframe Production." Econometrica, October 1963, Vol. 31 No. 4, pp. 679-694.
- Andrews, Donald W. K. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." Econometrica, March 1991, 59(3), pp. 817-858.
- Argote, Linda and Epple, Dennis. "Learning Curves in Manufacturing." Science, February 1990, 247, pp. 920-924.
- Argote, Linda, Beckman, Sara L. and Dennis Epple. "The Persistence and Transfer of Learning in Industrial Settings." Management Science, February 1990, 36(2), pp. 140-154.
- Asher, Harold. Cost-Quantity Relationships in the Airframe Industry, Santa Monica, CA: The RAND Corporation, R-291, 1956.
- Aviation Week & Space Technology. "TriStar Production Costs Offset Lockheed Profits." October 15, 1979, p. 32.
- Bailey, Charles D. "Forgetting and the Learning Curve: A Laboratory Study." Management Science, March 1989, 35(3), pp. 340-352.
- Baldwin, Richard and Krugman, Paul. "Industrial Policy and International Competition in Wide-bodied Jet Aircraft." in Robert E. Baldwin, ed., Trade Policy and Empirical Analysis, National Bureau of Economic Research Conference Report series Chicago and London: University of Chicago Press, 1988, pp. 45-71.
- Benkard, C. Lanier. "Dynamic Equilibrium in the Commercial Aircraft Market." Mimeo, Stanford University, 1998.
- U.S. Department of Labor, Bureau of Labor Statistics. Nonfarm Payroll Statistics from the Current Employment Statistics (National). <http://stats.bls.gov/sahome.html>.
- Cabral, Luis M. B., and Riordan, Michael H. "The Learning Curve, Market Dominance, and Predatory Pricing." Econometrica, September 1994, 62(5), pp. 1115-1140.
- Darr, Eric D., Argote, Linda, and Epple, Dennis. "The Acquisition, Transfer, and Depreciation of Knowledge in Service Organizations: Productivity in Franchises." Management Science, November 1995, 41(11), pp. 1750-1762.
- Dasgupta, Partha, and Stiglitz, Joseph. "Learning-by-Doing, Market Structure and Industrial and Trade Policies." Oxford Economic Papers, June 1988, 40(2), pp. 246-268.
- Davidson, Russell and MacKinnon, James G. Estimation and Inference in Econometrics. Oxford: Oxford University Press, 1993.
- Fudenberg, Drew, and Tirole, Jean. "Learning-by-Doing and Market Performance." Bell Journal of Economics, Autumn 1983, 14(2), pp. 522-530.

- Gruber, Harald. "The Learning Curve in the Production of Semiconductor Memory Chips." Applied Economics, August 1992, 24(8), 885-894.
- Gulledge, Thomas R., and Womer, Norman Keith. The Economics of Made-to-Order Production, New York: Springer-Verlag, 1986.
- Hansen, Lars P. "Large Sample Properties of Generalized Method of Moments Estimators." Econometrica, July 1982, 50(4), pp. 1029-1054.
- Hirsch, Werner Z. "Manufacturing Progress Functions." The Review of Economics and Statistics, May 1952, 34(2), pp. 143-155.
- Hirsch, Werner Z. "Firm Progress Ratios," Econometrica, April 1956, 24(2), pp. 136-143.
- International Monetary Fund. International Financial Statistics Yearbook. Washington, DC: International Monetary Fund, 1985, 1995.
- Irwin, Douglas A., and Klenow, Peter J. "Learning-by-Doing Spillovers in the Semiconductor Industry." Journal of Political Economy, December 1994, 102(6), pp. 1201-1227.
- Mincer, Jacob and Ofek, Haim. "Interrupted Work Careers: Depreciation and Restoration of Human Capital." Journal of Human Resources, Winter 1982, 17(1), pp. 3-24.
- Nelson, Richard R., and Winter, Sidney G. An Evolutionary Theory of Economic Change. Cambridge, MA: Harvard University Press, 1982.
- The New York Times. "Stacked Up at Boeing; Plane Giant Tries to Recover From Its Snarls in Production." November 14, 1997, p. 1.
- Nye, William W. "Firm-Specific Learning-by-Doing in Semiconductor Production: Some Evidence from the 1986 Trade Agreement." Review of Industrial Organization, June 1996, 11(3), pp. 383-394.
- Olley, G. Steven, and Pakes, Ariel. "The Dynamics of Productivity in the Telecommunications Equipment Industry," Econometrica, November 1996, 64(6), pp. 1263-1298.
- Prescott, Edward C., and Visscher, Michael. "Organizational Capital." Journal of Political Economy, June 1980, 88(3), pp. 446-461.
- Rapping, Leonard. "Learning and World War II Production Functions." Review of Economics and Statistics, February 1965, 47(1), pp. 81-86.
- Roach, J. R. and Eastwood, A. B. Jet Airliner Production List Volume 2. Farnborough: The Midas Press, 1995.
- Scherer, Frederic M., and Ross, David. Industrial Market Structure and Economic Performance. 3rd ed., Boston: Houghton Mifflin, 1990.
- Tirole, Jean. The Theory of Industrial Organization, Cambridge, MA: MIT Press, 1988.

The Wall Street Journal. "Lockheed is Planning a Write-Down of \$600 Million on Its TriStar in a Complex Financial Restructuring." June 3, 1974, p. 1.

The Wall Street Journal. "Lockheed Net Soared 108% in 3rd Period but Outlook for TriStar Sales is Dismal." November 7, 1975, p. 1.

Womer, Norman Keith. "Learning Curves, Production Rate, and Program Costs." Management Science, April 1979, 25(4), pp. 312-319.

Womer, Norman Keith. "Estimating Learning Curves from Aggregate Monthly Data." Management Science, August 1984, 30(8), pp. 982-992.

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

Zimmerman, Martin B. "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power." Bell Journal of Economics, Autumn 1982, 13(2), pp. 297-310.