AFIT-CE-18-008

LEARNING CURVE ANALYSIS IN DEPARTMENT OF DEFENSE ACQUISITION PROGRAMS

Technical Report prepared for the Naval Postgraduate School Acquisition Research Program

John J. Elshaw, PhD

Adedeji B. Badiru, PhD

Sharif F. Harris, First Lieutenant, USAF

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

Approved for public release; distribution is unlimited.
Prepared for the Naval Postgraduate School, Monterey, CA 93943.

The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).

Abstract

Learning curves are used to describe and estimate the cost performance of a serial production process. There are numerous models and methods; however, it is not precisely known which model and method is preferred for a particular situation. The primary objective of this research is to compare performance of the more common learning curve models. The research goals are to improve understanding of the systemic cost drivers of a production process, to clarify the relationship of these drivers to cost, and to present modeling methods. We use qualitative analysis combined with statistical regression modeling to assess fit. The research identified that preference for one function or another depended upon the shape of the data and how well a model formulation could be made to fit that shape. This was reliant upon the model's basic shape and the available parameters to alter its appearance. The typical learning curve model assumes that cost is a function of time but commonly omits factors such as production process resources changes (capital and labor) and the impact it has on cost. A learning curve model that includes the effects of resource changes would likely provide higher estimative utility given that the model establishes a systemic relationship to the underlying production process.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

THIS PAGE LEFT INTENTIONALLY BLANK

Introduction

On Friday, February 3, 2017, the U.S. Department of Defense (DoD) and Lockheed Martin Corporation signed an $8.2 billion contract for the next 90 F-35 fighter jets. Lockheed CEO Marillyn Hewson stated that per unit prices have decreased 60% since the program's initiation and that "this demonstrates a learning curve as efficient as any achieved on any modern tactical aircraft" (Weisgerber, 2017). In this particular instance, the phrase "learning curve" refers to the often observed decrease in unit cost of a serially produced item as additional units of that item are made. This is particularly prevalent during the earliest segment of production for a new product, which was the case for the F-35 program. Typically, this effect is attributed to increased familiarity with the product and its production method. However, management decisions to alter the production process through capital investment and increased employment has often been an underappreciated source of product cost reduction. Understandably, much attention is paid to learning curves, as they are used to describe and estimate the unit cost of a production process. This research seeks to further develop the understanding and use of learning curves within the DoD cost-estimating community.

The DoD has historically underestimated the development and production costs of Major Defense Acquisition Programs (MDAPs). A 2006 RAND Corporation review of 68 MDAPs spanning from 1968 to 2003 identified that average quantity adjusted cost growth from Milestone B estimate to project completion was 46% and from Milestone C estimate to completion was 16% (Arena, Leonard, Murray, & Younossi, 2006, pp. 21–22). Approximately 16.83% of the historical error was attributed to cost estimates; however, previous research did not address any particular estimating methodology or technique most responsible for estimative error, making it difficult to develop generalizations or targeted assessments of specific methods. Learning curves are a common component of cost estimates, and they are often applied to the production component of the acquisition process. Improving the estimating and modeling methodologies thereof would benefit overall cost estimate accuracy. Minor variations to the learning curve modeled estimate can have significant ramifications on total estimated procurement cost (Department of the Air Force, 2007, p. 3). Given the relative impact of learning curve estimates to the overall estimate, the objective of this research was to investigate the efficacy of the more common existing modeling techniques based upon the power function, as well as to compare the relative performance of those models to a sigmoid function. The specific questions this research aims to address are as follows:

1. How well do power function based models perform relative to empirical data?
2. How well does a sigmoid function perform relative to the same data as above?
3. Do the common variants of the power function based models compare to one another as well as to the sigmoid function, and if not, why?
4. How does contemporary learning curve modeling methodology impact and explain the cost behavior observed in the data?

Background

The principal objective of individuals and organizations is to maximize the value of their resources through effective management. Generally, businesses seek to maximize profitability and the productivity of employees and fixed investment. Governments desire the maximum public good of tax expenditures and investments. Managers depend heavily upon budgets and plans to assist resource management and decision-making. Budgets and strategic plans are programs for future action and are estimate driven. The utility of either depends upon the accuracy and reliability of those estimates and assumptions from which they are formulated. Consequently, the success of a manager is particularly reliant on the overall quality of those estimates provided. Those past and ongoing DoD major weapons procurement programs exemplify this relationship.

An MDAP is a research and development effort expected to require in excess of $480 million or a procurement exceeding $2.79 billion in fiscal year (FY) 2014 constant dollars (DoD, 2015, p. 44). The current annual Government Accountability Office (GAO) DoD procurement assessment released March 31, 2016, indicates 79 active MDAPs with a total value of 1.44 trillion in FY16 constant dollars (GAO, 2016, p. 8). Milestone B is the point in the procurement process at which research and development is mostly complete and a functioning prototype has been prepared. From here the product design is finalized, production methods chosen, and production begins. Milestone C is the point at which the design and production method are finalized. From this point, mass production begins and the item is operationally fielded (DoD, 2005, pp. 16–30).

The presently ongoing F-35 Joint Strike Fighter (JSF) procurement program is a noteworthy example of cost underestimation. A June 14, 2012, GAO report indicates that in 2012, the total program cost was estimated at $161 million per aircraft compared to the 2001 program initial baseline estimate of $81 million per aircraft (GAO, 2012, p. 5). This is an approximate 99% increase in an 11-year span, and additional increases have since occurred. Cost estimates that have a 46% post–Milestone B average increase for MDAPs, as well as the near doubling of estimated cost in the case of the F-35, are of limited utility to a decision-maker.

The specific reasons for cost estimate deviation relative to the actual expenditures are numerous and system peculiar. Frequently, the cost deviation observed is attributable to influences well beyond the reasonable control and foresight of a cost estimator. A 2008 RAND Corporation study of 35 MDAPs indicates that cost estimation error accounted for approximately 16.83% of the cost growth observed in those programs (Bolten, Leonard, Arena, Younossi, & Sollinger, 2008, p. 27). A similar 2004 study by David McNicol of 138 weapon systems procurements notes cost deviation as a result of mistakes (unrealistic estimates or poor management) in 70% of the systems reviewed. The average estimative deviation was -20% to 30% (Arena et al., 2006, p. 8). McNicol suggests that services showed a tendency toward optimistic estimates (Arena et al., 2006, p. 15). Nevertheless, estimative errors have been shown to be at least partially responsible for total cost deviation, and improving the estimating methods and techniques most frequently used by the DoD is an important component of the effort to enhance estimate accuracy. The primary research focus is the more common serial production process cost estimating techniques used within the DoD. Typically, the final output of a weapon system procurement program, be it aircraft, vehicles, or ships, are manufactured through a serial production process. During production, unit costs are frequently observed to decline with incremental production. The general explanation for this phenomenon is that increased familiarity with production tasks enables the reduction of the time and cost to produce additional units (Department of the Air Force, 2007, p. 4). Given sufficient time, a production process stabilizes and individual unit costs generally remain constant going forward. This phenomenon is modeled into cost estimates and is usually referred to as the learning curve.

Learning curves are developed using a statistical modeling technique called least squares regression. The objective is to fit a mathematical function to a data set in question by manipulating the function's controlling constants. The specific task is to minimize the total error (sum of squares of differences between mathematical functional output and the corresponding data points) in the model (McClave, Benson, & Sincich, 2014, p. 606). The coefficient of determination (R^2) is a measure of a model's explanatory power and fit quality relative to the data set (McClave et al., 2014, pp. 634–636). Generally, data is sourced from the production process being analyzed or an analogous program. The model is then used to describe the data as well as estimate future expected unit costs for the production process.

The power function is the mathematical function most regularly used by cost estimators to model the anticipated effects of learning in a production process. A power function is a mathematical function of the form, shown in Equation 1.

$$y = x^a \qquad \text{Equation (1)}$$

where

y = function output

x = independent variable

a = exponential constant

Power function-based learning models assign a negative value to the constant *a*, which produces a convex curve, as shown in Figure 1.

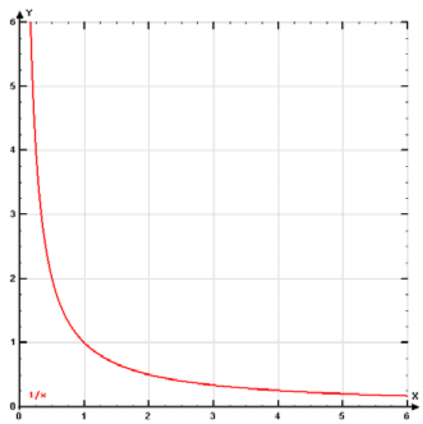

Figure 1. Power Function (Basic Form Example)

The shortfall of the power function, however, is that it does not exhibit behavior which could reasonably be interpreted as a long-term steady state. As the independent variable approaches positive infinity, the function output nears zero but does not visually stabilize horizontally. However, both logical intuition and empirical observation have shown that a production process, with a constant set of capital and labor, will stabilize given sufficient time. The extent of improvements identified over the lifetime of a specific production process are limited. Furthermore, the cost of additional or continual production process change would eventually exceed the benefits, once what is reasonably construed as the most efficient arrangement is identified. Given that the modeled phenomenon eventually stabilizes with time, whereas the power function does not, a divergence will emerge which widens with time, making any estimate based thereupon increasingly inaccurate with each successive iteration.

A possible alternative might be a mathematical function that initially decreases with increasing independent values, then stabilizes horizontally as the independent values approach infinity. A function of this nature would better approximate the typical expected learning behavior. A sigmoid function, or an s-curve, exhibits this behavior. A sigmoid function is one of the basic forms, shown in Equation 2.

$$y = \frac{1}{1 + e^x} \qquad \text{Equation (2)}$$

where

y = function output

e = natural logarithm (constant)

x = independent variable

A graph of the curve produced by this function is shown in Figure 2.

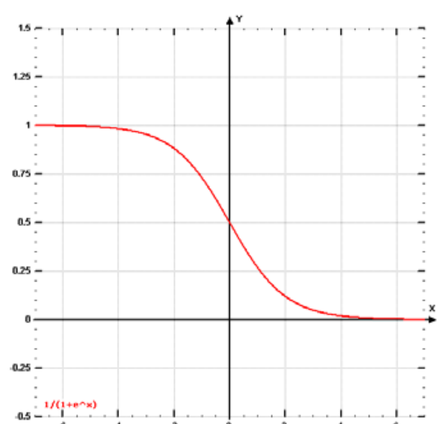

Figure 2. Sigmoid Function (Basic Form Example)

A sigmoid function unlike a power function begins and ends in what could be reasonably interpreted as a horizontal steady state phase. Additionally, a sigmoid function provides a practitioner increased control over the behavior of its shape, and potentially offers higher precision when modeling the learning effect compared to a power function. The increased control could possibly raise implementation difficulty compared to a power function, but the added value of comparative accuracy improvement will likely compensate. A sigmoid function is presumably preferred to a power function for modeling the cost behavior of a serial production process, and it is expected to have a higher coefficient of determination (R^2) value than that of a power function when modeled to the cost data for such a process.

THIS PAGE LEFT INTENTIONALLY BLANK

Theoretical Context

Learning is the accumulation of knowledge, understanding, or skill resulting from study, instruction, or direct experience. A learning curve is the two-dimensional graphical representation of the relationship between a quantitative measure of task performance and time. Performance is measured along the vertical axis and time on the horizontal axis. In finance and economics, performance is often a cost measure, such as the number of labor hours required to manufacture a product. The horizontal axis (time) is either a continuous amount when engaged in the activity of interest measured from a specific starting point, or discrete trials representing a specific production task or particular unit of production. The performance measure is not necessarily a direct function of time as implied by the graph, but is often influenced by numerous different factors. Typically, learning curves are used to assist with understanding and estimating the effect of those various influencing factors on performance. The analytical task is understanding and reasonably well segregating the influencing factors from each other to independently determine its effect on performance. The goal of learning curve analysis is to understand the impact of learning (accumulation of relevant knowledge and experience) on task performance. The implicit assumption is that relevant learning will accumulate from the initiation of an unfamiliar or foreign task and will be positively applied to enhance measured performance.

Analytical Context and Framework

The objective of learning curve research and analysis is to identify the factors and attributes most influential to performance changes. In the fields of finance and economics, analysts are generally concerned with the impact of learning on resource productivity. The underlying objective is to manage and direct the learning process toward a desired outcome. Learning curve analysis within the discipline of finance or economics is generally applied to the analytical framework of a production process. A production process is the structural arrangement of resources (labor and capital) to transform some input(s) to some desired output(s) (Anupindi, Chopra, Deshmukh, VanMieghem, & Zemel, 2012, pp. 3–6). Labor is the human workforce staffed in a production process and can be divided into direct and indirect labor. Direct labor is defined as those persons whose efforts can be traced to any one particular unit of output. Indirect labor includes the various support functions such as maintenance, engineering, and management. Capital is the firm's fixed infrastructure, machinery, and equipment required for production. A firm or organization can fully encompass a single production process, or there may be multiple production processes within a firm. A financial analyst or a cost estimator is tasked to understand and reasonably project the impact of learning on the production process.

Associated with the production process are certain performance measures that are essential to its effective management. Throughput is the total units of output produced in a standard amount of time (days, weeks, months, etc.; Anupindi et al., 2012, pp. 55–58). Productive resources (labor and capital) likewise have an associated economic cost that is measured per the standard unit of time. The ratio of economic cost for a specific resource, resource group, or the entire production process to corresponding throughput is economic cost per unit of output. Throughput is managed by changing the quantity, quality, or type of productive resources. The various combinations of those productive resources and the resulting throughput constitutes the production function for the process in question and is often expressed as a mathematical function such as Cobb-Douglas (Baye, 2010, pp. 156–165). This interpretation generally assumes that productivity levels of resources employed by a production process are relatively constant. Frequently, however, this is not true. Resource productivity levels often fluctuate depending upon the circumstances, and performance is affected by numerous factors. Learning has great potential to affect resource productivity. Learning is of great interest to process managers as it presents an opportunity to improve resource productivity without requiring additional investment. Distinguishing the changes in resource productivity consequent to learning effects and resource changes is critical to the correct analysis of a production process.

Mathematical modeling is used extensively in performing learning curve analysis. A mathematical model is the representation in mathematical terms of the relevant features and the behavior of some real-world phenomena (Bender, 1978, p. 1). Within the context of an economic production process, a learning curve is mathematically modeled such that the dependent variable (output) is the per unit production cost and the independent variable is either continuous time starting from absolute reference or the discrete count of production iterations. The most frequently used mathematical function for this purpose is the power function, shown in its basic form in Equation 1 and corresponding graph in Figure 1.

Although unit cost is not a direct function of time, it is modeled as such due to the implicit relationship with the learning effect. Learning accumulates with time and its positive application is believed to reduce production costs. Learning itself is a function of time resulting from the exposure, experience, and increasing knowledge of a specific production process or technique. The critical assumption with this methodology is that learning is the only significant influence on unit cost. If other factors significantly affect cost, it compromises the learning analysis and another modeling methodology would be more appropriate. Essentially the modeling process devolves into curve fitting—identifying a mathematical function that has the best fit to a set of data points. Curve fitting is useful for describing and estimating behavior in a data set, but the function which best fits the curve might not necessarily have a systemic relationship with the data it describes. A model that includes the significant variables that govern the dependent variable (unit cost) is preferred because it encapsulates the systemic relationship of the production process and would better explain and predict outcomes.

Development of Learning Theory

Theodore Paul Wright is credited with being the first to apply learning curve analysis to an economic production process with his research *Factors Affecting the Cost of Airplanes*, published in February 1936. Wright conducted his research within the Curtiss-Wright Corporation. Curtiss-Wright is a large aerospace corporation founded in 1929 and at the time of publishing was the largest nationally of that type. Wright's overall research objective was to improve the understanding of those significant factors most influential to aircraft production cost. Wright specifically wanted to develop a heuristic to describe and estimate aircraft unit production cost. Wright listed numerous important categorical areas including tooling, specification changes, aircraft size, and batch production quantity. The central finding of Wright's research was a cost quantity relationship shown in Equation 3. This formulation has since become the basis of contemporary learning curve modeling.

$$y = ax^b \qquad \text{Equation (3)}$$

where

y = average cost per unit of output

a = estimated cost to produce a single unit

x = total number of units to produce

b = ln(1 – cost reduction% per doubling)/ln(2)

Wright (1936) described this formula as a cost quantity curve that is used to compare the cost of a completed airplane in different quantities (p. 125). A cost estimator would input the total number of aircraft for a single production run, the expected cost of a single unit, and the expected percentage cost decrease per quantity doubling. The result is the expected average per aircraft unit cost. Wright indicates that he began developing the cost quantity relationship in 1922. The curve initially began from two or three data points of cost quantity pairings for identical aircraft production and was later supplemented with additional data when it became available (Wright, 1936, p. 122). Wright did not necessarily intend that his cost quantity relationship be used as a function for pricing individual units as either an independent variable of time or incremental unit production. Wright does not provide explicit detail of the internal specifics of production methods or production scale. Nevertheless, individual unit cost within a production run of a specific quantity will differ from those of another production run of differing quantity because the production process design would differ among batch sizes. Interpreting Wright's cost quantity relationship as a continuous function ignores the scaling limitations of a production process. Even within the scope of the cost quantity relationship the limitations to scale a production process are evident with increasing quantities. Wright acknowledges this long-term possibility stating that percentage rate reduction of cost gradually declines with increasing production totals (Wright, 1936, pp. 125–126). Notwithstanding its limitations, Wright's curve has since been transformed and used as a continuous function for measuring individual unit cost. The formulation shown in Equation 3 has become the most common method for measuring cost decreases thought to be the consequence of learning. Several adaptations have since been made based upon the original formulation and have since been applied in numerous applications.

The most typical simplification of the learning curve is that increased cumulative output results in decreased average unit cost. However, the lingering question is why that may be the case. The most commonly offered explanation is that with each additional unit of production, task familiarity and knowledge of the production process increases and thus allows for the identification of cost-reducing improvements. This brings to focus the key difference between how Wright presented the learning curve and how it is generally used. From the perspective of a cost quantity relationship as determined by Wright, the unit cost differential from one output quantity to another is predominantly a function of scale. That the selection of a specific production quantity would also imply the selection of a certain production process design, and the key difference being the process throughput and per unit average total cost. The alternative interpretation does not necessarily consider effects from the production process scale but implies that marginal cost reductions from one unit to the next is a predominant function of learning. Wright does not describe the particulars of the production process prevalent in his analysis, but suggests the process would change depending on desired production quantity (Wright, 1936, pp. 124–126). Wright provides three general explanations for the behavior of the cost quantity curve. Wright reasons that labor learning, economies of scale, and resource selection impact the curve (Wright, 1936, p. 124). Wright goes on to say that economies of scale is the principal factor of the three.

Learning Based Improvement

Labor learning is the most frequently offered reason for the expected behavior of learning curves and one of three reasons mentioned by Wright when developing the cost quantity curve. Wright specifically states that "improvement in proficiency of a workman with practice and particularly if time in motion studies are made, is well known" (Wright, 1936, p. 124). Often, a large part of the cost reduction in a production process is believed to originate from improvements in the direct labor force. This thinking was confirmed to an extent by the empirical study conducted by Dr. Nicholas Baloff, in which he found a gradual decrease and an eventual flattening of the learning curve in a capital-intensive production process, but found no such decrease or flattening in the labor-intensive processes studied (Yelle, 1979, p. 310). The common understating was that a human paced production process will likely show more cost improvement than one paced by machines given the human capacity to learn and improve (Hirschmann, 1964). Additionally, the results of human task performance studies ranging from the earliest such as those of Ebbinghaus (1913) and Bryan & Harter (1897) to modern studies are supportive. The general belief is that individual performance improvement in the context of a production process will aggregate and accrue to the firm enhancing overall productivity. However, other comprehensive studies of the firm offer a different narrative. Dr. Kazuhiro Mishina conducted an analysis of Boeing B-17 bomber production during World War II. The study focal point was Boeing Plant Number Two in Seattle, WA. With respect to labor learning, Mishina identified that direct labor learning did not play as significant a role as once thought in improving overall production productivity. He found that workers generally were not particularly skilled or experienced in aircraft production and turnover was high, disallowing for much in the way of learning (Mishina, 1999, pp. 162–163). The overwhelming majority of the productivity gains observed in the plant was a function of managerial improvements to the overall production process and not improvements of direct labor proficiency (Mishina, 1999, p. 164). Mishina stated that management's resource employment and organization decisions rather than gains in proficiency of the resources themselves accounted for the overall success of the plant. Similar findings were noted in a study of a truck assembly plant conducted by Dr. Dennis Epple. The objective was to compare the learning gains of two separate production shifts which at one point had operated as a single shift. He found that both shifts were equally productive and that gains in learning are embedded in the organizational structure as well as its technology. He also identified that the second shift had a reduced improvement rate compared to the first because of lowered managerial and industrial engineering oversight (Epple, Argote, & Murphy, 1996, pp. 84–85). This research emphasizes the importance of management and structural arrangements within the firm as it pertains to its overall productivity. Those gains from direct labor learning are not as significant a contributor to organizational gains when compared to the resource employment and production process design choices.

Economies of Scale

The second explanation offered by Wright for the underlying behavior of the cost quantity relationship is economies of scale. Economies of scale exists when average total costs decline as productive output (throughput) is increased (Baye, 2010, p. 185). Logically, this is congruent with the cost quantity interpretation of Wright's report, that unit cost is predominantly a function of the production process design or subsequent changes thereto. The cost quantity relationship implies that for higher expected total production quantities, a production manager would choose to produce on the most optimal point on the long run average total cost curve that does not exceed the total amount and remains within the boundaries of productive resource constraints.

Wright indicated that per unit capital and labor costs will decline as a function of increasing scale. Wright specifically states that increased scale allows for the economical substitution of labor for capital, reduced setup costs for both labor and capital, and a more efficient spread of indirect (overhead) cost (Wright, 1936, pp. 124–126). Mishina's research supports this finding. His B-17 production process analysis indicates that the cost reductive learning which occurred was primarily the result of management decisions and not increased direct laborer proficiency. Mishina indicates that B-17 direct hour labor requirements declined from approximately 71 worker-years to 8 worker-years from 1941 to 1945 (Mishina, 1999, pp. 150–151). Additionally, he identified that the predominant reason for the observed improvement, particularly from the early phases of the program, resulted from increasing production scale (Mishina, 1999, pp. 175–176). Furthermore, the majority of aircraft labor cost reduction occurred during the operations scale-up phase occurring from approximately May 1940 until December 1942 (Mishina, 1999, p. 159). During this period, Boeing emplaced 91% of the total fixed investment in terms of Plant Number Two floor space, established the Tooling Department (managed 70,000 dies and jigs), and increased the total direct labor force from 9,972 employees in August 1941 to 17,000 in February 1942 (Mishina, 1999, pp. 157, 162).

In a comparative analysis between Boeing Seattle, WA, and Ford Willow Run, MI (B-24 Liberator), Mishina notes the respective differences between their learning curves. Mishina specifically indicates that the Willow Run learning curve stabilized more rapidly than Boeing Seattle. He attributes this to the process design difference. Ford used a more capital-intensive design, while Boeing had a labor-intensive focus. Mishina described the Ford Willow Run learning curve as a period of necessary adjustment before the process achieved its total designed potential (Mishina, 1999,167–168). This corroborates previous research by Baloff, who identified that the learning curves of capital-intensive production processes plateaued in 75% of the cases studied where the labor-intensive processes did not (Yelle, 1979, pp. 310–311). It is possible that Baloff, like Mishina, was observing the higher relative flexibility of a labor-intensive process over a capital-intensive one. Research from Dr. Peter Thompson identified similar findings regarding investment and production scale increase. Thompson analyzed World War II Liberty Ship production with a focus on seven of the largest producing shipyards at the time. Thompson found a positive relationship between capital investment and labor productivity levels. Additionally, there is also a positive relationship between periods of increasing capital investment, productivity growth, and increasing throughput (Thompson, 2001, pp. 121–122). Thompson concludes that traditional learning curve analysis suffers from an omitted variable bias. Specifically, with respect to Liberty Ships he found that omitting the capital investment component overstates the contribution of learning to productivity increases (Thompson, 2001, p. 132).

Given the aforementioned empirical studies, the possibility exists that the typical visualization of decreasing unit costs concurrent with increasing production quantities is not necessarily a function of incremental learning, but instead represents the modification of the production process as it moves to a more efficient location on the average cost curve. The traditional method of modeling learning analysis does not develop a systemic relationship to those other major underlying contributing factors, which potentially results in misunderstanding the behavior of a production process.

Resource Selection

The third and final component is resource selection. The nature and design of the production process alters labor force expertise requirements. Generally, the resource cost of production workers is measured with time (hours typically) and not absolute economic cost. This is appropriate for long-term comparative purposes as it discounts for wage rate fluctuations and inflation. However, it does not show the cost impact of selecting labor of differing skill classes. Wright mentions the association of the production process design, economies of scale, and resource selection. Specifically, he notes that increased production scale, supplementing or supplanting labor with capital, and procedure standardization will reduce the need for highly skilled or specialized labor reducing costs (Wright, 1936, p. 124). Mishina states that construction of the first B-17s in 1935 was done primarily with highly skilled workers using hand tools (Mishina, 1999, p. 157). This highly contrasts with the later scaled-up production operations when Boeing employed large numbers of workers who on balance had little if any manufacturing experience or knowledge (Mishina, 1999, p. 163). In this situation, an adequate pool of experienced craftsman was not available. However, the production process transformation from low volume to standardized high volume did not require such skilled labor, and consequently wage rates declined, lowering average unit labor cost.

THIS PAGE LEFT INTENTIONALLY BLANK

Design/Model Analysis

The objective of the methodology is the development of an unbiased quantitative and qualitative analysis of serial production process cost performance data. The purpose is evolving an improved understanding of the effect of learning on cost performance, and to improve the estimating methods thereof. Probabilistic mathematical modeling is the central method of analysis to develop the results. Field data was collected from several different serial production processes. Final results are used to further develop insight and appreciation for the learning effect in a production process.

Tools and Equipment

Spreadsheet modeling and nonlinear optimization were the principal means by which the analysis was accomplished. Microsoft Excel 2016 installed on a Microsoft Windows 7 desktop PC was the platform used. The Solver add-in software from Frontline Systems for Excel was used to accomplish the nonlinear optimization calculations. The pairing of Excel and Solver allows for simple development of probabilistic mathematical models as well as accompanying data visualizations such as scatterplots, box plots, line graphs, and histograms. The Real Statistics Resource Pack add-in for Excel was also used to automate the calculation of certain statistical measures. High functionality and high relative ease of use compared to similar alternatives prompted these decisions.

Field Data

The field data selected for analysis in this research was sourced from varied serial production processes. The data was collected by the process owning firm or organization. Data collection methods and systems were automated and manual. The data are assumed highly reliable as it is used by the firm for its internal process management decisions and externally for billing and compliance. Each data set varied, but at minimum each had a trial number and cost performance measure. Each unit of output is produced in sequence and assigned an ordinal number starting with the initially completed unit and continuing until the process is terminated. This sequence or trial number is often augmented with the date and time of completion. The measure of cost performance in all the data analyzed was the labor time required for task completion. Other quantitative and qualitative measures, such as production breaks or product variants, were available but differed from one data set to the next. Data recording generally started with production process initialization particular to a specific output.

Procedure

Probabilistic mathematical modeling is the central analytical technique employed in this process. A probabilistic model or function is one that has a deterministic aspect as well as a random error component (McClave et al., 2014, p. 603). The functions used to model the data are the basis of the deterministic component. The random error is found though the regression analysis techniques. The regression process is designed to identify those functional parameters which minimize the error or total deviation between function output and the corresponding observation (McClave et al., 2014, p. 607). The cost data, collected from the various production processes, is being regressed with the experimental functions to assess their suitability for modeling and estimating cost performance. Data is visualized with a scatterplot to help develop a rudimentary behavioral assessment and for the consideration of certain qualitative factors. Finally, a review of regression quality and the validation of certain key assumptions is performed.

The first task in executing the analytical method was to review the field data and to standardize it for insertion into the modeling template. Specifically, the relevant time values and the cost performance measures were reviewed to ensure conformity with the spreadsheet standard. Following a review of the data, it was then placed into the modeling template for analysis. A segment of example data is shown in Figure 3.

No.	*Group*	*Item*	*Date*	*Hours*
1	EMD	91-A4001	08/08/97	33614
2	EMD	91-A4002	07/25/98	22667
3	EMD	91-A4003	02/24/00	28592
4	EMD	91-A4004	10/31/00	25159

Figure 3. Spreadsheet Modeling Template (Example Data)

Once the data was placed into the modeling template the regression process commenced. The principal method for accomplishing this was though the complimentary Solver Excel add-in software. Solver was used to minimize the total sum of squares of the errors (SSE) for each experimental function. The error term or residuals are the difference between the observed value (hours) and the value calculated by the model, which corresponds to the same independent value (item number). The errors are squared and summed to calculate SSE. The principal basis of regression analysis is identification of those parameters that minimizes the SSE for a mathematical function (McClave et al., 2014, p. 607). Solver was utilized to minimize SSE by adjusting the corresponding parameters using nonlinear optimization. Figure 4 presents an example of the Solver interface.

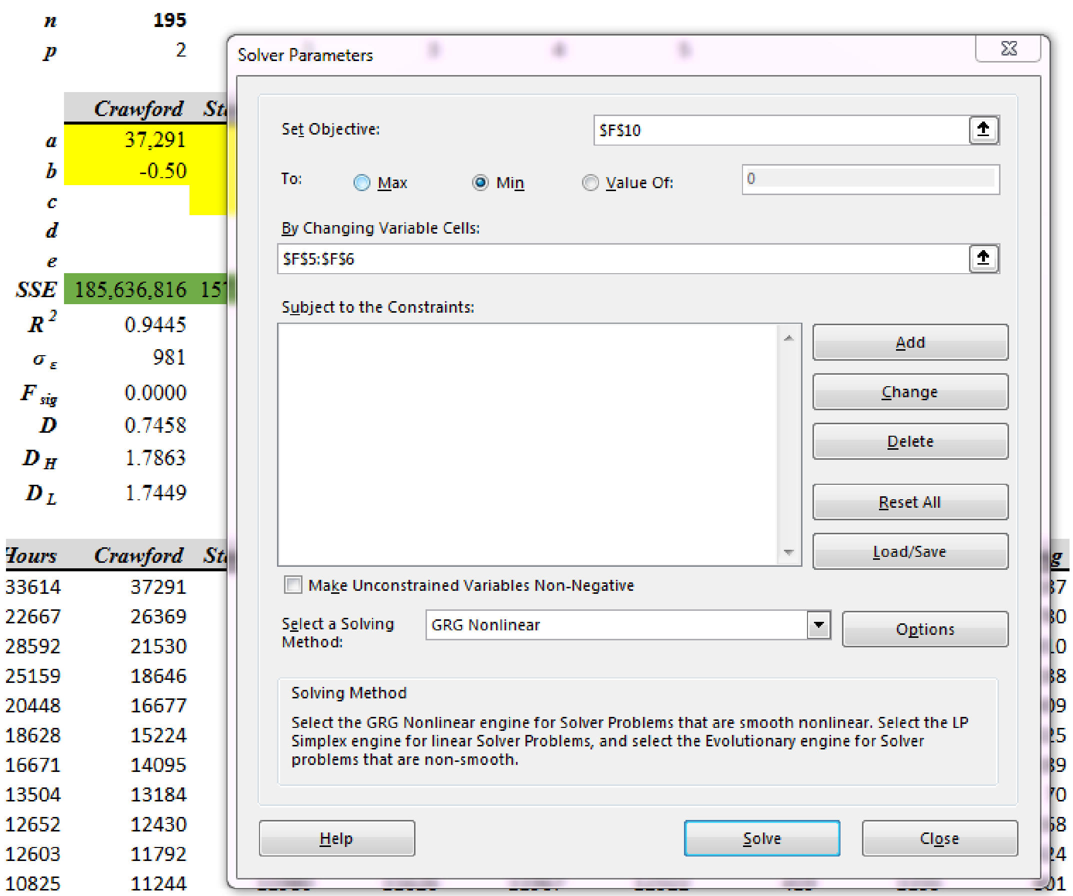

Figure 4. Spreadsheet Modeling Template (Solver Interface)

In total, five univariate mathematical models were tested for each data set. Four of the five functions are variants of the base form power function first shown in Equation 1 and graphed in Figure 1. The fifth function is the sigmoid function (s-curve) displayed in Equation 2 and graphed in Figure 2. The cost quantity relationship originally proposed by Wright is the basis of the power function based models. Several variants have since been created with the objective of developing a more robust model. Four of them were tested as part of this research. The first is the reinterpretation of the Wright cost quantity curve to a unit cost curve generally attributed to James R. Crawford (Department of the Air Force, 2007, p. 4). Rather than the independent variable being the total production quantity and the output variable the average total cost to produce that amount, the unit curve interpretation states that the independent variable is an ordinal production unit and the output variable is the cost of that particular unit. The unit theory is shown in Equation 4.

$$y = ax^b \qquad \text{Equation (4)}$$

where

y = function output (cost)

x = independent variable (unit number)

a = constant (first unit cost)

b = exponential constant (learning)

The second is the Stanford-B model, which adds an additional constant that allows for the lateral shifting of the curve along the horizontal axis (Badiru, 1992, p. 178). This added variable allows for increased ability to fit the curve to any given set of data compared to the basic unit curve (see Equation 5).

$$y = a(x + b)^c \qquad \text{Equation (5)}$$

where

y = function output (cost)

x = independent variable (unit number)

a = constant (first unit cost)

b = lateral shift (prior units of experience)

c = exponential constant (learning)

The third is the DeJong learning formula. This variant introduced the incompressibility factor (M) concept to the unit theory curve. Incompressibility is the percentage amount (ranging from 0 to 1) of the production process that is machine automated (Badiru, 1992, pp. 178–179). A negative exponent basic form power curve approaches zero on the vertical axis as the independent values approach positive infinity. Within the context of a learning curve, this implicitly means that costs will effectively approach zero with increased unit production. The incompressibility factor essentially places a floor beneath the learning curve and in the context of the DeJong model represents the unchanging nature of machines beneath infinitely compressible unit labor costs. Mathematically this allows the unit theory curve to be shifted vertically, potentially allowing for improved fit to a data set compared to the unit curve. The mathematically transformed DeJong model is shown in Equation 6.

$$y = a + bx^c \qquad \text{Equation (6)}$$

where

y = function output (cost)

x = independent variable (unit number)

a = constant (capital operation cost)

b = constant (first unit labor cost)

c = exponential constant (learning)

C = a +b (first unit total cost)

M = a ÷ C (incompressibility factor)

The fourth and final variation of the power curve that was tested as a part of this research is the combination of the Stanford-B and the DeJong model attributed to Gardner W. Carr (Badiru, 1992, p. 178). Mathematically this variant provides full control over the positioning of the curve in two-dimensional space, maximizing the possibility for fitting to a data set (see Equation 7).

$$y = a + b(x + c)^d \qquad \text{Equation (7)}$$

where

y = function output (cost)

x = independent variable (unit number)

a = constant (capital operating cost)

b = constant (first unit labor cost)

c = constant (prior units of experience)

d = exponential constant (learning)

C = a +b (first unit total cost)

M = a ÷ C (incompressibility factor)

Lastly, the sigmoid function, or s-curve, was tested. The s-curve is most similar to the Carr model of the power curve family. The shape of the curve can be directly controlled, and it can be placed anywhere in two-dimensional space as well (see Equation 8).

$$y = a + \frac{b}{(1 + e^{c(x+d)})^e} \qquad \text{Equation (8)}$$

where

y = function output (cost)

x = independent variable (unit number)

a = constant (lowest value)

b = constant (upper most in excess of lowest value)

c = constant (curve saturation point)

d = constant (lateral shift)

e = constant (curve symmetry control)

The quantitative measures and qualitative indicators of fit quality are reviewed for each function following the identification of its SSE minimizing parameters. Regression summary statistics are reviewed to assess quantitative fit quality. Specifically, the *F* test for overall model validity is reviewed, adjusted coefficient of determination (R^2), and the standard error of the regression for each function. The *F* test is used to determine if the regression is statistically significant to an σ level of 0.05. Adjusted R^2 is a measure of the functions explanatory power relative to the data set (McClave et al., 2014, pp. 634–636). The standard error of the regression is the standard deviation of the error (McClave et al., 2014, pp. 619–620). Valid R^2 values range from zero to one with one being a perfect explanatory model. The standard error of regression value can range from zero to positive infinity with zero being perfect predictability. The sample output is shown in Figure 5.

	Crawford	*Stanford-B*	*DeJong*	*Carr*	*Sigmoid*
a	37,291	47,720	-440	2,722	3,870
b	-0.50	1	37,241	245,850	36,119
c		-0.56	-0.47	5	1.90
d				-1.19	1.74
e					0.06
SSE	183,926,140	157,149,378	183,158,206	111,662,669	225,136,399
R^2	0.9450	0.9528	0.9450	0.9663	0.9317
σ_ε	976	905	977	765	1,089
F_{sig}	0.0000	0.0000	0.0000	0.0000	0.0000
D	0.7499	0.8686	0.7393	1.3078	0.7033
D_H	1.7863	1.7969	1.7969	1.8076	1.8184
D_L	1.7449	1.7345	1.7345	1.7239	1.7133

Figure 5. Spreadsheet Modeling Template (Summary Statistics)

One of the objectives of this research was to determine if there was any trend within the data sets with regard to one model being statistically favored over the other. In addition to the adjusted R^2 values, a pairwise analysis of the variance (ANOVA) test was performed to determine which of the models was preferred to a statistical significance level of 0.05. This test is useful to determine if one particular model is preferred to another as adjusted R^2 values alone may not necessarily be indicative (Motulsky & Ransnas, 1987, p. 371). Sample output from this test is shown in Figure 6.

	Crawford	*Stanford-B*	*DeJong*	*Carr*	*Sigmoid*
Crawford	0.5000	0.8620	0.5115	0.9997	0.0813
Stanford-B	0.1380	0.5000	0.1448	0.9907	0.0067
DeJong	0.4885	0.8552	0.5000	0.9997	0.0773
Carr	0.0003	0.0093	0.0003	0.5000	0.0000
Sigmoid	0.9187	0.9933	0.9227	1.0000	0.5000

Figure 6. Spreadsheet Modeling Template (Comparative ANOVA Test)

The primary qualitative indicators of fit quality assessed are the scatterplot with function overlay and the corresponding residuals scatter plot. Although statistical tests are required for validating the results of a model, visual inspection of the data can provide additional insight to the behavior of the data. Sample output is shown in Figure 7.

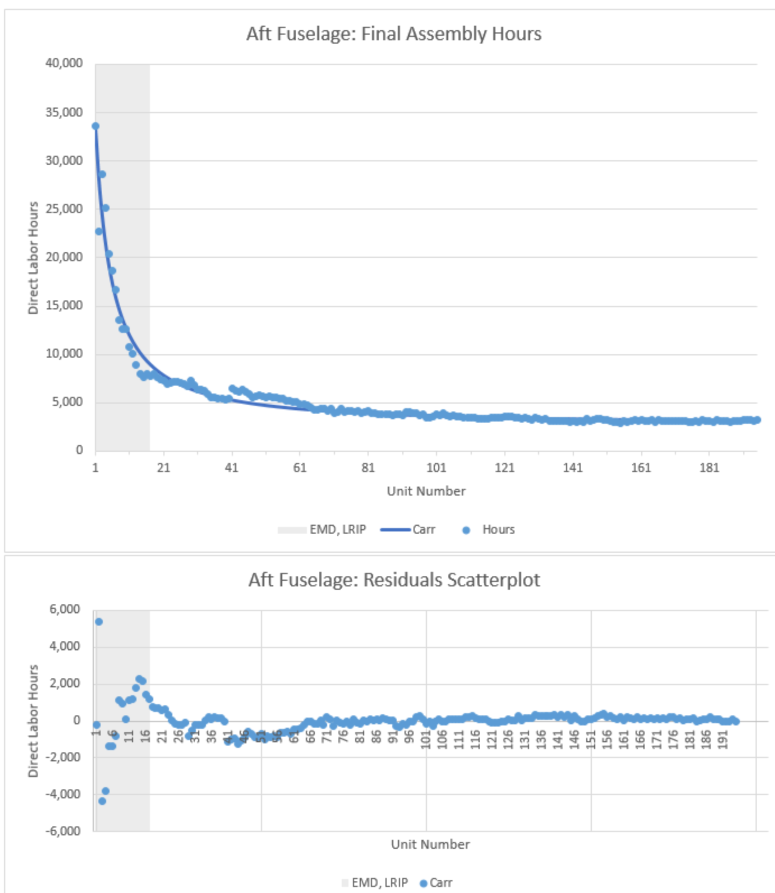

Figure 7. Spreadsheet Modeling Template (Data and Error Scatter Plots)

Analysis of Operational Data

In the context of an economic production process, learning curve analysis is used to enhance the understanding of production cost behavior over time, as well as estimating future behavior. If a mathematical learning curve function is formulated that can establish a systemic relationship with production process cost behavior, then it could be reasonably utilized to describe and estimate future cost behavior. However, if a systemic relationship cannot be established, then the mathematical function could be used only as a fitted curve, and applied on discretionary basis for description and limited extrapolation. The allure of learning curves is that production process cost behavior can be explained with only one controlling independent variable. The reality however is that numerous significant factors influence production costs, which disallows robust single independent variable modeling. Learning analysis generally has been a best fit curve analysis process, which explains the creation of the numerous competing models and methods over time. The overall research objective is to develop actionable insight using available cost data to enhance production process cost estimating methodologies.

Sequential production cost data was analyzed from five production activities with the methods previously described. Each of the five mathematical functions summarized in Table 1 were fit to the production process data. The results were reviewed to assess quality of fit as well as the relative fit quality between the different functions. Relevant supplementary information is included wherever appropriate to enhance the analysis.

Table 1. Summary of Models

Common Name	**Form**	**Type**
Unit Cost Model (Crawford)	ax^b	Power
Stanford-B Model	$a(x+b)^c$	Power
DeJong's Learning Formula	$a+b(x)^c$	Power
Carr 1946 (“S-curve”)	$a+b(x+c)^d$	Power
Sigmoid Curve (S-Curve)	$a+\frac{b}{(1+e^{c(x+d)})^e}$	Exponential

Air Force Advanced Tactical Fighter (ATF) Program

The Advanced Tactical Fighter (AFT) is a single-seat, twin-engine, all-weather, stealth, air superiority fighter aircraft exclusively operated by the U.S. Air Force (USAF). In total 195 ATFs were produced. The first ATF was delivered to the USAF on April 9, 1997, and the final on April 24, 2012. The initial nine aircraft were produced within the Engineering and Manufacturing Development (EMD) phase of the acquisition process. The remaining 186 aircraft were final production models produced thereafter. Corporation A was the prime contractor and Corporation B was a prime partner. Manufacturing responsibilities for significant aircraft components and systems was segmented between the two. Corporation B was responsible for the wings, aft fuselage, avionics integration, 70% of mission software, training systems, life support, and protection systems (Boeing, 2014). Corporation A was responsible for program management, forward and center fuselage, control surfaces and stabilizers, and critical avionics systems (Boeing, 2014). Fabrication of certain components and final assembly for the aft fuselage and wings was completed at the Corporation B Integrated Defense Systems industrial center located in Seattle, WA (Waurzyniak, 2005). The production process was subdivided into numerous activities and sub-activities which fed the respective final assembly tasks for the aft fuselage and wings. Completed units were delivered to Corporation A in Marietta, GA, for integration and final assembly (Waurzyniak, 2005).

A learning curve analysis was performed for the requisite number of direct labor hours to complete those significant activities for a single aft fuselage or wing paring. The following five major activities were analyzed: (1) aft fuselage final assembly, (2) aft fuselage feeder line, (3) wing final assembly, (4) wing spar fabrication, and (5) wing skins composite fabrication. Both aft fuselage and wing final assembly involved the integration of all requisite component parts for completion. The aft feeder line performed pre-assembly work for certain components and assemblies that were later integrated with the aft fuselage (Boeing, 2004). Wing spars are primary structural members that run lengthwise the wings. The wing skins are the outermost surface of the wing assembly. The data was sourced from DoD Cost Assessment and Program Evaluation (CAPE) and collected by Corporation B. The fields within the data set are the group, aircraft number, and direct labor hours. Group is a descriptive field indicating the production phase or lot to which the aircraft belongs. Each aircraft is assigned a unique identification number of the form FY-A4xxx, where FY is the contract fiscal year and the trailing three digits is a unique number. The data was supplemented with information from F-16.net. F-16.net is a community open source data repository for military aircraft and was used to identify the aircraft delivery dates.

Regression analysis was performed on the five data sets utilizing the spreadsheet modeling template and methodology previously described. Regression analysis for the five mathematical functions displayed in Table 1 was conducted for each data set for 25 in total. All of the regressions for the tested functions were statically significant to an α level of 0.05 per the overall ANOVA test. The approximate range of adjusted R^2 values for all regressions is from 0.81 to 0.98. This information is summarized in Table 2.

Table 2. Regression Summary Statistics

		Crawford	**Stanford-B**	**DeJong**	**Carr**	**Sigmoid**
Aft Fuselage	R^2	0.9445	0.9528	0.9450	0.9663	0.9317
	σ_ε	981	905	977	765	1,089
Aft Feeder Line	R^2	0.9237	0.9627	0.9392	0.9768	0.9678
	σ_ε	137	96	123	76	89
Wings	R^2	0.9316	0.9495	0.9341	0.9705	0.9561
	σ_ε	1,925	1,653	1,890	1,264	1,542
Wing Spar	R^2	0.8089	0.8554	0.8387	0.8560	0.8116
	σ_ε	122	106	112	106	121
Wing Skins	R^2	0.8852	0.8848	0.8846	0.8851	0.8350
	σ_ε	105	105	105	105	126

For each of the five activities analyzed, the power function based learning curve variant credited to G. W. Carr was either the superior choice or among the preferred alternatives. This was determined with comparative ANOVA testing between each of the alternatives to a statistical significance level of 0.05. For two of the five data sets this statistical test is indeterminate between the alternatives with the highest overall fit quality. A summary of this information is presented in Table 3.

Table 3. Model Fit Quality—Preferred Alternative(s)

Aft Fuselage	**Aft Feeder Line**	**Wings**	**Wing Spar**	**Wing Skins**
Carr	Carr	Carr	Stanford-B	Crawford
			DeJong	Stanford-B
			Carr	DeJong
				Carr

The initial hypothesis from the commencement of this research was that a sigmoid function would outperform a power function for modeling the data of production process. The results from the five processes analyzed as part of this research does not support this position. In each of the five production activity data sets one of the power function based variants was the preferred alternative. The predominant reason for this outcome is that the power function can generate curves that resemble a capital letter "L," where the vertical and horizontal components of the L meet in a gradual curve as opposed to a sharp point. This shape is generally well suited to approximate the near vertical cost decline for early units in the EMD and LRIP phase as well as the production units. However, the power function does not perform well in situations where the process stabilizes, and incidentally was the original motivation for this research. Four of the five production processes did eventually show signs of stabilization and it was most prevalent in the wing skins fabrication task. A scatter plot with plots for both the Crawford and Sigmoid functions are shown in Figure 8, as well as the Crawford curve error term scatter plot.

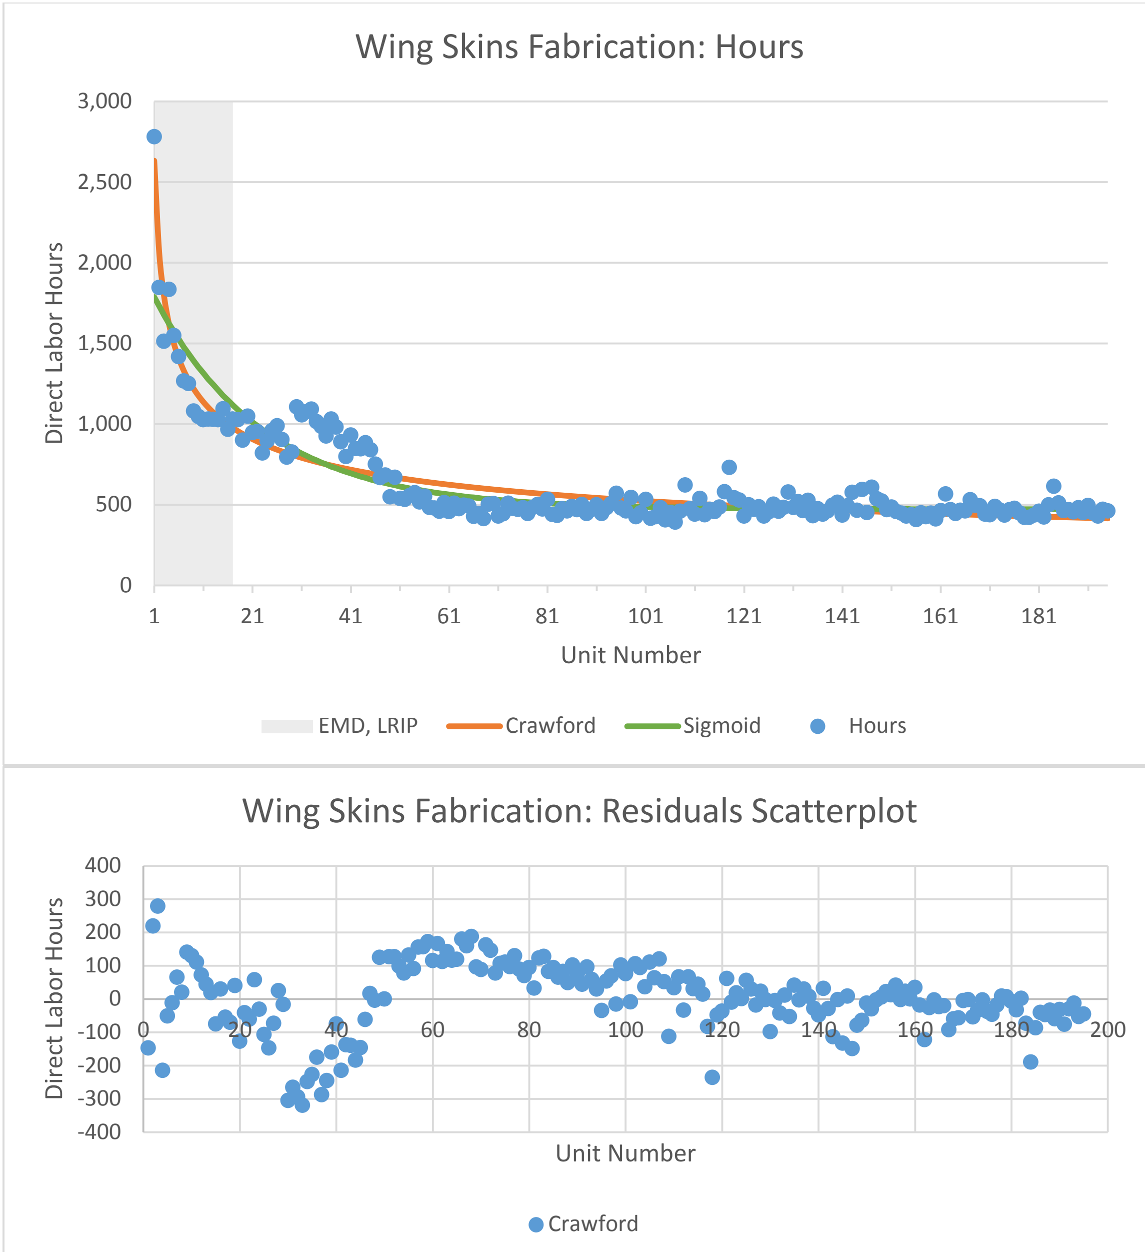

Figure 8. Wing Skins Fabrication Hours (Crawford and Sigmoid)

The upper panel of Figure 8 shows that the power function is more adept than the exponential Sigmoid function for approximating both the near vertical early phase decline and the horizontal component of the data set. However, the curve does not do well by comparison during the horizontal component as the Sigmoid function. The process stabilizes nearby unit number 60 and from that unit onward, the Crawford curve residuals (lower panel of Figure 8) show an approximate downward sloping line that crosses the horizontal axis at the approximate center point of this range. This pattern exemplifies the original hypothesis that early units are overestimated and later units underestimated when a power function is fit to cost data. Contrastingly, the Sigmoid function handles this phase well but cannot be manipulated to produce both a near vertical drop and a curved saturation in the manner of a power curve. The Sigmoid residuals scatter plot shown in Figure 9 displays the closer approximation of the steady state from the 60th unit onward to program completion.

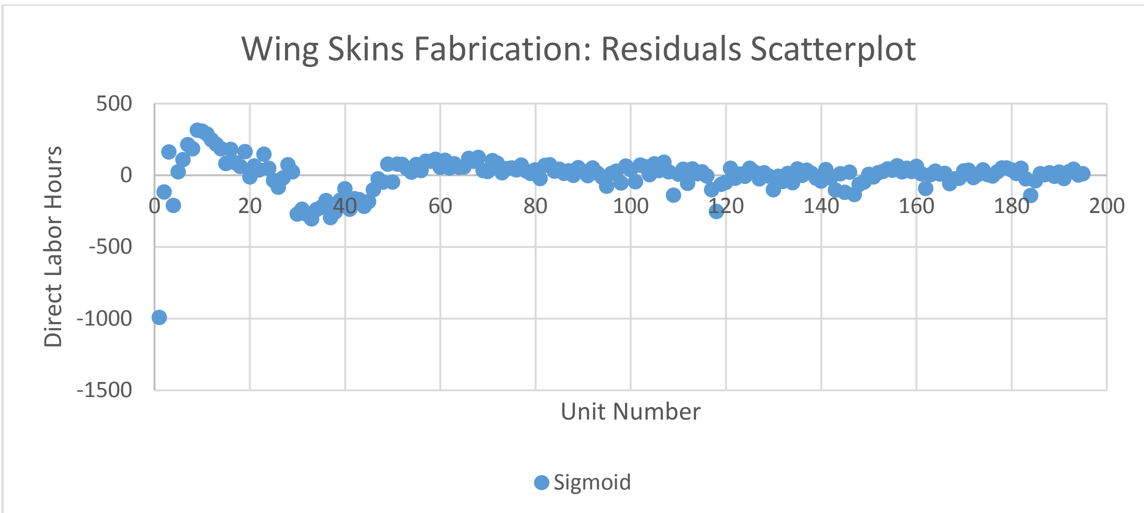

Figure 9. Wing Skins Fabrication Hours (Sigmoid Residuals)

When a production process stabilizes, it can be reasonably approximated with descriptive statistics. Using the wing skins fabrication process as an example, from the 60th unit until process termination the average cost was 477 labor hours and the standard deviation was approximately 50 hours. Figure 10 displays this relationship with a residuals scatter plot for the average (477 labor hours) starting at the 60th unit until the last.

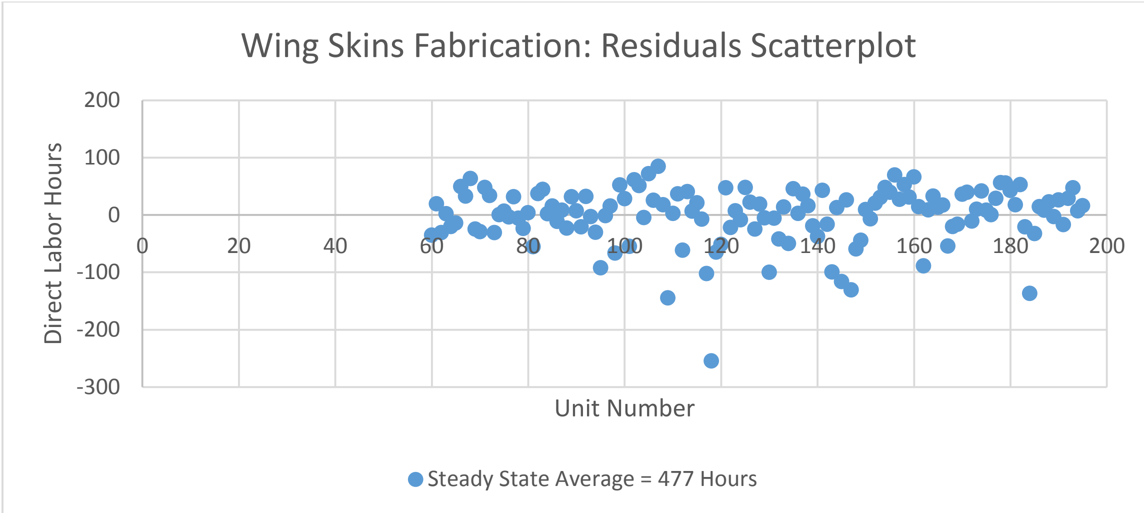

Figure 10. Wing Skins Fabrication (Steady State Residuals)

There were several outlying points (two standard deviations from the mean), and had they been absent, the average unit cost would have been 468 hours with a standard deviation of approximately 32 hours. The wing skins are made of carbon fiber and were fabricated by a programmable tape (carbon fiber) laying machine. When a replicable and satisfactory process was identified, it was maintained for the program's duration (Cantwell, 2007, p. 12). Assuming that the process was unaltered once a steady state was achieved, the mean and standard deviation would have been a reasonable cost estimate for this segment of the program. The basic shape of a mathematical function and its ability to be manipulated to fit a data set is what ultimately determines its suitability or preference to an alternative.

Fitting a mathematical function to a data set helps to explain what happened, but it does not tell you why it happened. For example, the steady state portion of the wing skins production process could have been well explained with its mean and standard deviation. However, if management were to fundamentally alter the process it would invalidate the relationship, and only after having collected adequate information could the new trend be substantiated. Cost analysts frequently use the production process data of one program to estimate another. Often the focus of the estimate is the attributes of the item in question. The production process itself, however, has just as much if not more influence over cost. Additionally, the learning curve (production process change) from one program may not necessary help describe another. Production process change generally is not the result of happenstance, but deliberate management decisions. Added focus to understanding the production process can be beneficial to developing improved cost estimates.

Learning curve analysis is commonly performed, as was done in this research thus far, using the unit number or production increment number as the independent variable. A production process, however, is generally paced by time. Production rates and production cycle time often fluctuate. Unit based production process analysis could potentially result in misunderstanding process behavior depending on the time disparity between each unit. Time sequenced learning curve analysis, on the other hand, presents a production process in its most unbiased form. It shows production just as it occurred and allows an analyst to better understand the behavior of a production process. To illustrate this concept, the time sequenced learning curve for the aft fuselage final assembly process is shown in Figure 11.

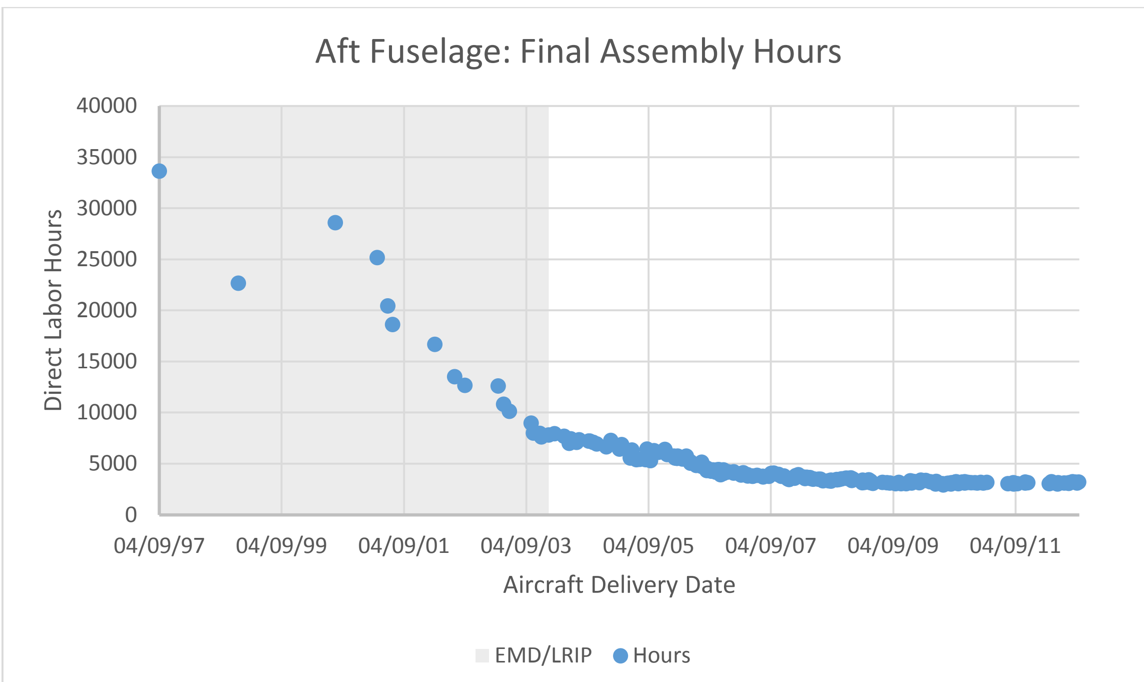

Figure 11. Aft Fuselage Final Assembly Hours—Time Sequenced

Figure 11 displays the labor hour content for each unit of production plotted along the horizontal axis by its respective aircraft final delivery date. Completion dates for each major Corporation B responsible component is unknown, but the aircraft final delivery date is a reasonable proxy. The segmented production system between Corporation A, Corporation B, and the various other parties had to be reasonably well synchronized, otherwise inventory shortages and excesses would have built throughout the system. This presentation method reveals important observations that the ordinal measured alternative (Figure 8) would not have shown. First, it clearly distinguishes the difference between the EMD/LRIP segment and full rate production as indicated by the frequency of observations. Second, it provides a better appreciation of the relative amount of time spent in each segment. Approximately 9% of total units were produced in EMD/LRIP but it consumed 42% of total project time. Finally, it shows that the labor cost decline was not as dramatic as shown in Figure 8, but occurred over a longer period of time. An ordinal scale compressed the EMD/LRIP phase relative to the program balance, altering its meaning, and accounts for the stark difference in appearance. Time phasing also facilitates throughput analysis. Figure 12 shows trailing 12-month production for the quarterly period indicated.

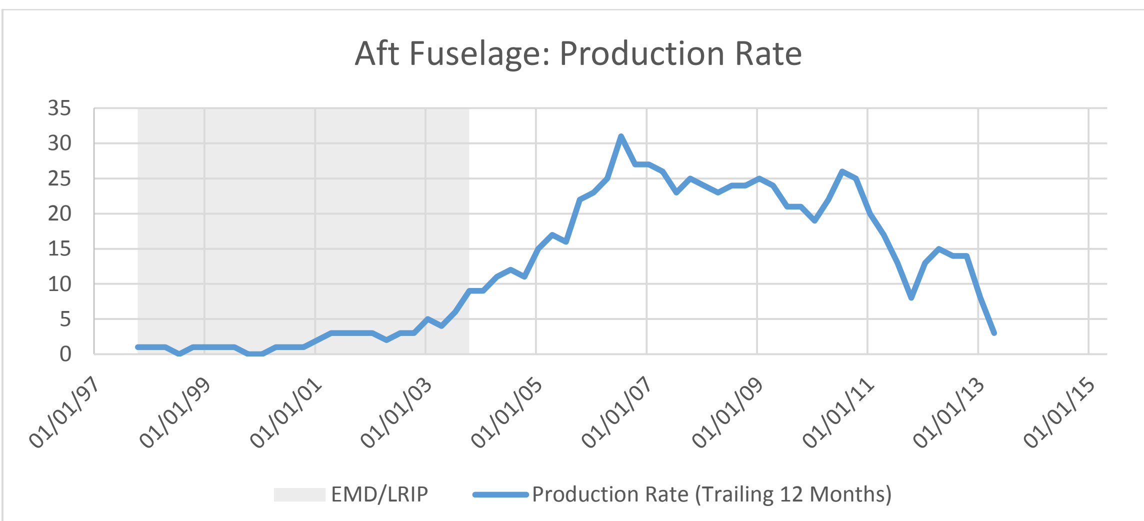

Figure 12. Aft Fuselage Production Rate (Trailing 12 Months)

The production rate increase would imply a corresponding increase in productive capacity. Likewise, the combination of decreasing unit production costs and an increasing productive capacity (implied by the rate and price change) would support an economies of scale argument—that Corporation B transitioned from one position on their average total cost curve to another. The “learning by doing” argument would explain that direct labor workers collectively improved their task performance, resulting in an approximate 91% reduction of the labor hour content for aft fuselage final assembly from 33,614 hours for the initial unit to a mean of 3,134 for the final 65. Although this very well may have been the case, additional information may suggest otherwise. An August 8, 2006, press release states the total man hour requirement to build aft fuselage units had decreased by 89% since the first delivery in October 1996 (Cantwell, 2006). The press release attributes this reduction to the lean production principles, industrial design, and capital investment. The article goes on to attribute reductions to a late 2003 transition from massive fixed assembly jigs to smaller flexible cart tooling, time savings from electron-beam welding which reduced the need for traditional fasteners by 75%, and an automated laser-guided machine to drill holes for the remaining fasteners (Cantwell, 2006). The implications are that labor productivity improvement probably had more to do with the production process design and capital investment (trading labor hours for more advanced capital) than the collective improvement of direct laborers. Figure 13 displays a conceptual graphic summarizing the relationship between long run average total cost, productive capacity, and unit cost.

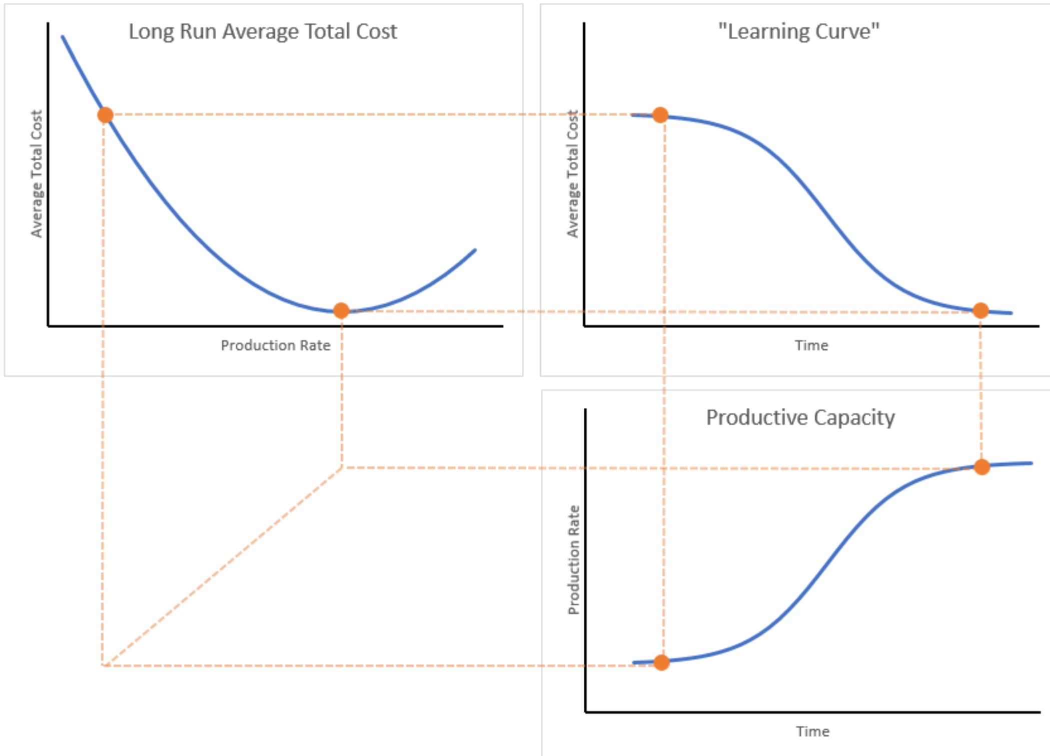

Figure 13. Relationship of Average Total Cost, Production Cost, and Production Rate

Figure 13 expresses the underlying managerial economics theory behind the ideas presented regarding Corporation B's operation. The long run average total cost curve depicts the economies of scale concept—declining long run average costs with increased throughput (Baye, 2010, p. 185). When the ATF program began, Corporation B likely used a relatively higher labor proportion production process than what they had during full rate production. Only after selecting the full rate production process design were resources committed (capital and additional employees). With sufficient time for capital emplacement and locating the employees, the production process shifted from the initial low rate point to the higher rate production point and likewise benefitted from the labor cost reduction that accompanied that transition. The missing component to the analysis is the number and cost of machine hours in both the early EMD/LRIP phase and later in full rate production. The change in capital operating expense and the amount of capital investment would help complete the analysis. From this perspective, learning curve analysis is not particularly distinct from a manufacturing cost estimate. Learning curve analysis as presently practiced is concerned with how the production process changes over time. Manufacturing and learning curve analysis attempt to answer four basic questions: what does the production process look like now, what will it look like in the future, how long will it take to get there, and how much resource growth (capital and labor) is required for the transition?

To better describe a time sequenced production process, curve fitting analysis can be applied. A time sequenced presentation of cost data has the benefit of depicting events as they actually occurred in real time. Figure 14 displays labor hours for aft fuselage final assembly fit with the sigmoid function.

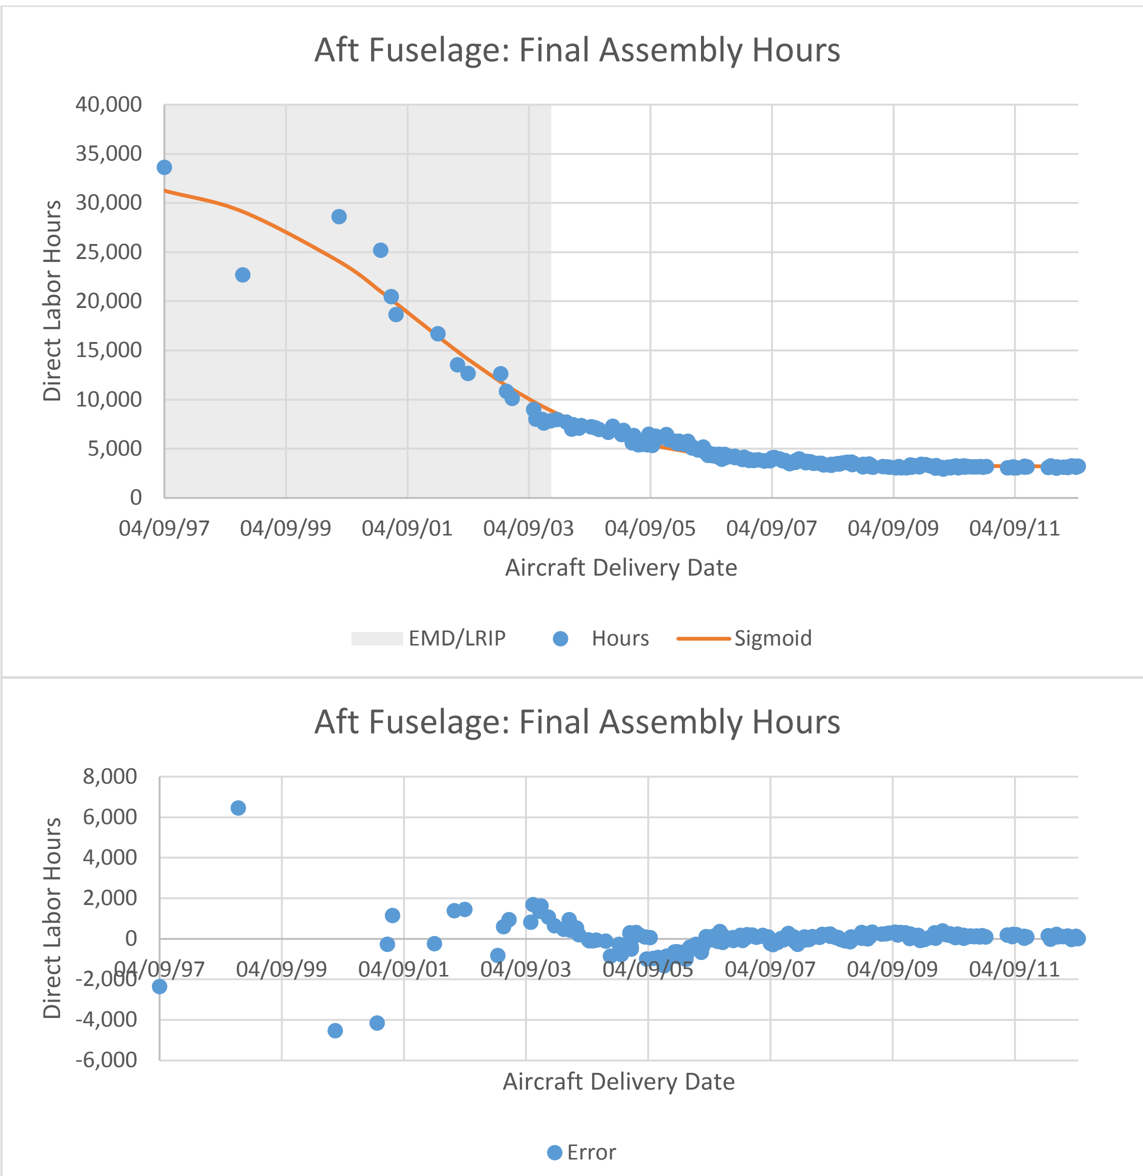

Figure 14. Aft Fuselage Final Assembly Hours—Time Sequenced (Sigmoid)

For time-sequenced production data, the sigmoid function was the preferred alternative. As before with the unit ordered analysis, this is purely a case of how well any particular mathematical function can fit data based upon its shape, and does not necessarily imply a systemic relationship to the production process. In situations which the time between cost observations is approximately equal, then unit based (ordinal) and time based sequencing of the independent horizontal axis would constitute the same analysis. Table 4 shows the regression summary statistics and pairwise ANOVA results. Given the shape of the data, none of the power curve based variants were appropriate for time sequenced curve fitting. The pairwise ANOVA test was run for the time sequenced sigmoid regression compared to the other three combinations. For this particular data set the ordinal sequenced Carr curve is preferred to the time sequenced sigmoid function, but not to a statistically significant level.

Table 4. Aft Fuselage Assembly—Time Sequenced Regression Summary Statistics

	Ordinal Sequenced		**Time Sequenced**	
	Carr	**Sigmoid**	**Carr**	**Sigmoid**
R^2	0.9663	0.9317	0.8906	0.9629
σ_ε	765	1,089	1,378	803
Paired ANOVA	0.7483	0.0000	0.0000	-

THIS PAGE LEFT INTENTIONALLY BLANK

Results and Discussion

In the context of a serial production process, the term *learning curve* refers to the observed decline of individual unit cost with incremental production. This effect is most commonly attributed to increased knowledge and experience producing the product and its chosen production method. Learning curves are of particular interest to managers, and learning curve data is often modeled mathematically to create extrapolative estimates for an existing ongoing program or to build an analogous estimate for a planned program. The initial motivation for this research was to determine if one particular mathematical model among those most commonly used for this task was more consistently preferred to others for modeling learning curve data. However, as the research progressed it became apparent that how well any particular mathematical function could fit to the cost data of a production process was not nearly as important as identifying underlying systemic drivers for the production process in question. Although a well-fit mathematical function can be useful for describing data and creating extrapolative estimates, this utility depends on the underlying systemic factors that drive the observed cost behavior to remain constant. Any deviation in effective behavior of the most relevant systemic factors would invalidate the relationship. Additionally, using a mathematical function that was well fit to the learning curve data of one particular production process to build an analogous estimate would be limited without relevant quantitative data regarding the production process drivers. The cost to produce an item is just as much a function of production process as the item itself.

General learning curve theory and analysis was initially developed in the realm of human psychology. Researchers were developing techniques to measure and understand human learning, and how learning affected memory and sensory-motor task performance. Early learning performance studies were focused on individual performance, but the basic principles have since been applied to entire organizations such as industrial production processes. Performance is typically modeled as a dependent variable of time. The implicit assumption is that learning is accumulated and actively applied to task performance over time, and that the variable of interest is the individual or the organization being observed. This relationship reasonably allows an analyst to model performance with a mathematical function with a single independent variable. However, if the particulars of task execution changed during the observation period, then performance is not just a function of learning but also that of the nature and timing of the change in task execution. In the context of a production process, if the design of the process itself changed in terms of how production was carried out, then the impact of those changes as well as worker proficiency influence the performance of that process. The more commonly used learning curve models used to analyze production processes do not consider the performance effects of alterations made thereto. Including quantitative input regarding the relative productive capabilities of the production process into the model perhaps could improve the analysis and modeling.

Conclusion

The primary research objective was to compare performance and create a preference ordering for the most commonly used mathematical learning curve models for developing probabilistic models of serial production process data. The result is improved appreciation of the systemic cost drivers of a production process, their relationship to cost, and present modeling methods. To that end, four of the more common variants of the power function and the sigmoid function, an exponential function, were fit to the data of five production activities of the Air Force Advanced Tactical Fighter (ATF) program. No definitive trend in terms of universal favorability of one curve over another was identified. Generally, the preference of one function or another was situational dependent and influenced by certain dynamics of the analytical process. The most influential factors are the visual form of the data sets when plotted, the visual form of the function to be fit and how well it can adhere to the data, and the sequencing choice of the horizontal axis—either with unit completion ordering or the time of completion. Unit of completion sequencing of the horizontal axis can alter the appearance of the learning curve, particularly if significant amounts of time dispersion exist between data points. In the five production activities analyzed, unit costs decreased from program initiation to its maturity, paralleling the transition from low rate early phase production to high rate production. Four of the five activities stabilized as the program matured. Early phase production rates were lower than full rate production, and ordinal data sequencing visually distorted the rate of unit cost decline. When the data sets are ordinally sequenced, generally the power curve models were preferred to the Sigmoid function. The opposite is true when the data is sequenced by the time of unit completion. In both cases the desirability of one function or another is totally a function of the data's visual appearance and the function of choice's ability to conform to it. Time sequencing, however, presents the production process as it actually occurred and allows for the visual assessment of production rate, production breaks, and parallel unit production. Lastly, the curve fitting process calculates the function parameters which minimize the total variance between the fit function and data set. In most cases the calculated constants do not adhere to the constraints set within the scope of each model, and frequently produced values that could not be reasonably explained particularly in the case of the power curve models.

How well a mathematical function can be fit to a data with the regression process depends on the visual shape of the data with respect to its sequencing along the horizontal axis, and the basic shape and adaptability of the math function that is being targeted to fit that data. Fit quality does not translate to a systemic relationship with the underlying cost drivers of the production process. The primary drivers of a production process are relative resource levels (labor and capital), the process arrangement and design, and learning both of the direct labor force as well as the management and support engineers. Curve fitting is useful for describing and predictive extrapolation but it does not necessarily help develop a broad systemic understating across multiple production processes. A fit curve describes the behavior of the process but does not necessarily correspond to the underlying process drivers which to an extent limits its predictive use. A fit curve would not anticipate major changes to a production process. Typically, management alters the production processes, which ultimately alters the process results but these changes do not feed to a fitted model. Specifically, the transitioning of a production process from an initial lower rate to its final planned full rate arrangement, which entails the addition and significant rearrangement of resources (labor and capital), accounts for a significant portion of production cost decline commonly observed in learning curve analysis. Production process analysis would likely improve with more focus on understanding the production process. Predictive learning curve analysis commonly entails identifying production process behavior over time using the curves of past similar efforts. The comparative focus is the production item, but often little attention is given to the attributes of the production system itself. Understanding the production process in its present state, the process as it will be implemented in the future, and the amount of time it will take to transition is ultimately what a cost analyst needs to understand.

Works Cited

Anupindi, R., Chopra, S., Deshmukh, S., Van Mieghem, J., & Zemel, E. (2012). *Managing business process flows: Principles of operations management* (3rd ed.). Upper Saddle River, NJ: Pearson Education.

Arena, M. V., Leonard, R. S., Murray, S. E., & Younossi, O. (2006). *Historical cost growth of completed weapon system programs* (F49642-01-C-0003). Santa Monica, CA: RAND.

Badiru, A. B. (1992). Computational survey of univariate and multivariate learning curve models. *IEEE Transactions on Engineering Management, 39*(2), 176–188.

Baye, M. R. (2010). *Managerial economics and business strategy* (7th ed.). New York, NY: McGraw-Hill/Irwin.

Bender, E. A. (1978). *An introduction to mathematical modeling*. Mineola, NY: Dover.

Boeing. (2004). Lean on me. Retrieved from http://www.boeing.com/news/frontiers/archive/2004/november/cover.html

Boeing. (2014). Historical snapshot: F-22 Raptor. Retrieved from http://www.boeing.com/history/products/f-22-raptor.page

Bolten, J. G., Leonard, R. S., Arena, M. V., Younossi, O., & Sollinger, J. M. (2008). *Sources of weapon system cost growth: Analysis of 35 major defense acquisition programs* (FA7014-06-C-0001). Santa Monica CA: RAND.

Bryan, W. L., & Harter, N. (1897). Studies in the physiology and psychology of the telegraphic language. *Psychological Review, 4*, 27–53.

Cantwell, D. (2016). Boeing starts production of aft fuselage for 100th F-22 Raptor. Retrieved from http://boeing.mediaroom.com/2006-08-08-Boeing-Starts-Production-of-Aft-Fuselage-for-100th-F-22-Raptor

Cantwell, D. (2007). Serving up Raptor wings, Seattle style. *Boeing Frontiers, 5*(8), 12.

Department of Defense (DoD). (2015). *Operation of the Defense Acquisition System* (DODI 5000.02). Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics (OUSD[AT&L]).

Department of the Air Force. (2007). The cost improvement curve. In *Air Force Cost Analysis Handbook* (pp. 8.1-8.85). Washington, DC: Air Force Cost Analysis Agency.

Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. New York, NY: Columbia University.

Epple, D., Argote, L., & Murphy, K. (1996). An empirical investigation of the microstructure of knowledge acquisition and transfer through learning by doing. *Operations Research, 44*(1), 77–86.

Government Accountability Office (GAO). (2012). *Joint strike fighter: DOD actions needed to further enhance restructuring and address affordability risks* (GAO-12-437). Washington, DC: Author.

Government Accountability Office (GAO). (2016). *Defense acquisitions: Assessments of selected weapon programs* (GAO-16-329SP). Washington, DC: Author.

Hirschmann, W. B. (1964). Profit from the learning curve. *Harvard Business Review*. Retrieved from https://hbr.org/1964/01/profit-from-the-learning-curve

McClave, J. T., Benson, G. P., & Sincich, T. (2014). *Statistics for business and economics* (12th ed.). Boston, MA: Pearson Education.

Mishina, K. (1999). Learning by new experiences: Revisiting the Flying Fortress learning curve. In N. R. Lamoreaux (Ed.), *Learning by doing in markets, firms, and countries* (pp. 145–184). Chicago, IL: University of Chicago Press.

Motulsky, H. J., & Ransnas, L. A. (1987). Fitting curves to data using nonlinear regression: A practical and non-mathematical review. *The FASEB Journal, 1*(5), 365–374.

Thompson, P. (2001). How much did the Liberty shipbuilders learn? New evidence for an old case study. *Journal of Political Economy, 109*(1), 103–137.

Waurzyniak, P. (2005). Lean fighter: How Boeing’s lean manufacturing team cut costs, cycle times on the F/A-22 Raptor stealth fighter program. *Manufacturing Engineering*. Retrieved from https://sme.org/MEMagazine/Article.aspx?id=46310&taxid=1411

Weisgerber, M. (2017). The price of an F-35 was already falling. Can Trump drive it lower? *Defense One*. Retrieved from http://www.defenseone.com/business/2017/01/f-35s-price-has-been-falling-can-trump-lower-it-even-more/134919/

Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, *3*, 122–128.

Yelle, L. E. (1979). The learning curve: Historical review and comprehensive survey. *Decision Sciences*, *10,* 302–328.

ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net