Learning Rate Sensitivity Model

Nichols F. Brown
System of Systems Engineering Office
Systems Engineering Division

Timothy P. Anderson
Integrated Cost and Schedule Analysis Department
Systems Engineering Division

Approved for public release. OTR 2018-00468.

Abstract

Learning curves are frequently used in cost estimating applications to model the regular and predictable reduction of per unit cost associated with a manufacturing production run. In the space system cost estimating community, learning curves have been reliably used to estimate cost of multiple-spacecraft acquisitions, usually consisting of two to five units. However, spacecraft providers have started proposing constellations of unprecedentedly large numbers, consisting of hundreds of spacecraft. Since learning rate best practices for the space industry have been based on relatively small production runs, it has become necessary to re-visit the applicability of learning curves and learning rate assumptions for space cost estimating applications with large production numbers. The Aerospace Corporation has developed a methodology to test assumptions about learning rates vis a vis proposed cost estimates. For example, if a spacecraft provider makes claims about the learning rate associated with a cost estimate for a large number of spacecraft, this methodology provides a data-driven assessment of whether or not this learning rate/cost combination is feasible or even likely. The sensitivity model further describes what learning rate would be necessary in order to meet a proposed cost estimate. While this process was developed for a space application, it is equally applicable to other manufacturing processes in which large numbers of units are to be built, and can be used to assess cost estimates relative to learning rate assumptions.

Contents

Figures

Tables

1. Introduction

Learning curves are frequently used in cost estimating applications to model the regular and predictable reduction of per unit cost associated with a manufacturing production run due to improvements over time. In the space system cost estimating community, learning curves have been used to estimate cost of multiple-spacecraft acquisitions which to date have usually consisted of two to five units. However, spacecraft providers have started proposing constellations of unprecedentedly large numbers, consisting of hundreds of spacecraft. Since learning rate best practices for the space industry have been based on relatively small production runs, it has become necessary to re-visit the applicability of learning curves and learning rate assumptions for space cost estimating applications with large production numbers. The Aerospace Corporation has developed a methodology to test assumptions about learning rates vis a vis proposed cost estimates in a high rate production environment. For example, if a spacecraft provider makes claims about the learning rate associated with a cost estimate for a large number of spacecraft, this methodology provides a data-driven assessment of whether or not this learning rate/cost combination is feasible or even likely. The sensitivity model further describes what learning rate would be necessary to meet a proposed cost estimate. While this process was developed for a space application, it is equally applicable to other high rate production manufacturing processes, and can be used to assess cost estimates relative to proposed learning rate assumptions.

2. Introduction to Learning Curves

Learning curves have enjoyed wide use as a tool to estimate recurring costs in a manufacturing process. The general idea holds that as production quantity increases, there is a predictable reduction in manufacturing cost, with direct labor learning generally assumed to be the predominant driver. Conversely, automation is generally assumed to reduce learning as it reduces touch labor and machines are assumed to repeat the same process in the same way each time. However, it is unclear if this is always true; processes and machines can always be reprogrammed or tinkered-with by humans looking to reduce costs. In the future, the possibility of machine learning may also break the maxim that more automation means less learning.

Two primary theories of learning curves are commonly used – unit theory, and cumulative average theory. The unit theory learning curve, credited to J. R. Crawford (1947), assumes that in a production process, the unit cost of some doubled quantity is equal to the unit cost of the undoubled quantity times the slope of the learning curve, where the slope is expressed as a percentage. Another way of stating this is that every time the quantity of units produced doubles, the cost to produce a unit decreases by a constant percentage.

The cumulative average theory learning curve, credited to T. P. Wright (1936), is similar, but uses cumulative average cost as its basis, rather than individual unit costs. The cumulative average theory assumes that in a production process, the cumulative average cost of some doubled quantity is equal to the cumulative average cost of the undoubled quantity times the slope of the learning curve, where again, the slope is expressed as a percentage. In other words, every time the quantity of units produced doubles, the cumulative average cost decreases by a constant percentage.

For this study, all learning rates were calculated and executed using the cumulative average (Wright) theory. Thus, whenever a learning curve percentage is mentioned, we refer to cumulative average theory learning curves.

Regardless of the theory chosen, the purpose of learning curves is to model the predictable reduction in the cost of multiple sequential units of production in a manufacturing process. The estimated cost of production is directly related to the learning rate. For example, if one were to model the production cost of 100 units using an 80% cumulative average theory learning curve, with first unit cost of $100K, then the cumulative average cost estimate, for a given quantity x, would be modeled using the following learning curve:

$$Y_x = Ax^b$$

Where

Y_x refers to the average cost of units 1-*x*;
A refers to the theoretical first unit cost (also known as *T1*);
x refers to the unit in question; and
b refers to the learning parameter

Figure 1 illustrates learning curves for a variety of different learning rates. Note that as the learning rate decreases, the learning curve gets progressively steeper.

Modeling the cost in this fashion is very important if one desires to capture the cost reduction due to learning. Failure to model learning when it exists is equivalent to assuming a 100% learning curve. That is, each doubled quantity's average cost is 100% the average cost of the undoubled quantity – in essence, no learning.

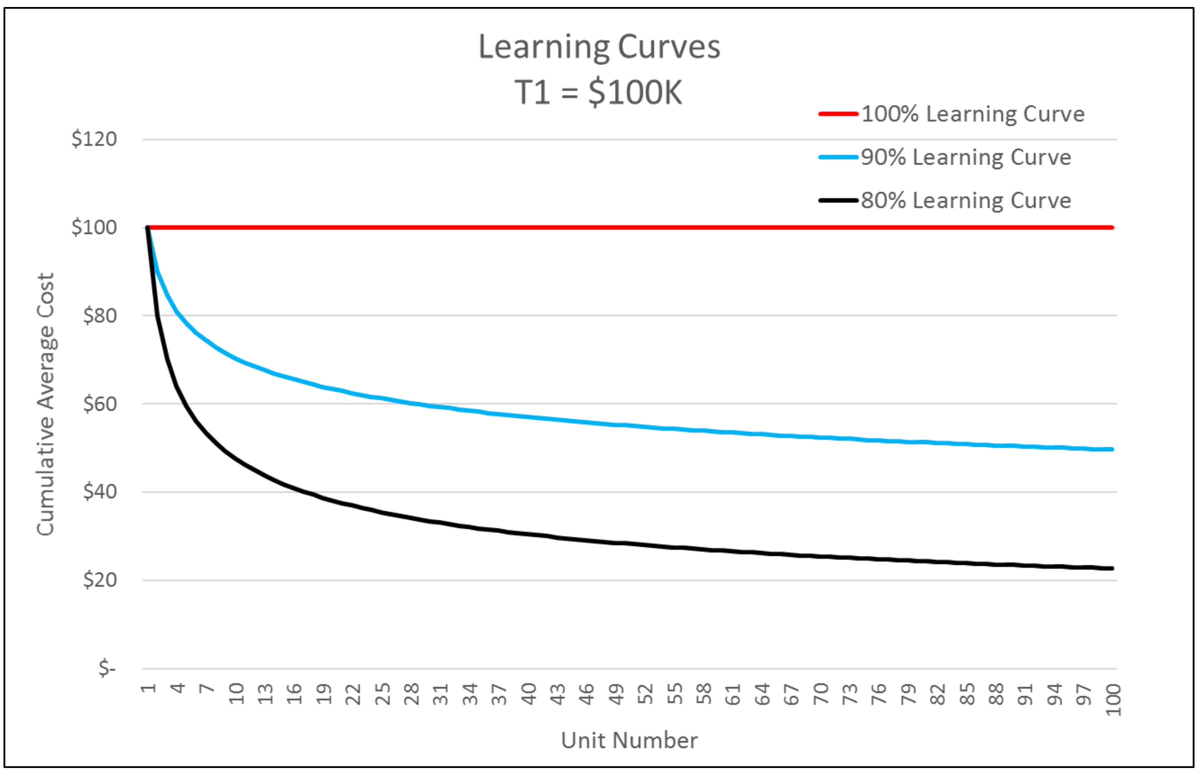

Figure 1 - Typical Learning Curves

The lesson here should be obvious. If an acquisition organization believes the manufacturing process should be subject to some degree of learning, then it is advantageous to model that learning in the cost estimate, which should indicate the potential to achieve significant cost savings.

3. Problem Description

Suppose an acquisition organization is acquiring multiple copies of some hardware, using a high rate production (HRP) process, and the development contractor claims that they expect a certain learning curve in their manufacturing process, and has provided an accompanying cost estimate. Is the learning curve assumption used in the cost estimate of a HRP manufacturing processes credible?

This research is intended to provide a way to answer that question. We provide a mechanism that says, "in order to achieve the proposed cost, the manufacturing process would have to experience the equivalent of a learning curve with slope *X*%." Moreover, we provide historical context so that we can say, "this *X*% learning slope was achieved by *Y*% of a broad collection of manufacturing programs."

Our process assumes that a cost model exists that enables computation of a first unit cost for the hardware under examination. Then, we determine the learning curve slope that would be needed to achieve the proposed cost. Next, we compare the derived learning curve slope to a collection of learning rates achieved by over 100 projects across a variety of industries to provide some context as to how difficult it would be to achieve the derived learning rate. Finally, we have developed a sensitivity analysis process that automatically varies the learning rate for production and creates a plot of the total cost versus the learning rate. This enables one to say, "given the learning rate proposed by the developer, their cost estimate should be $*X*," or similarly, "given the cost estimate proposed by the developer, their learning rate would need to be *Y*%."

As the saying goes, "this is not rocket science." Learning curves are already standard practice in the space system cost estimating community. Until recently, space system acquisitions have usually been limited to very small production quantities. Many are one-of-a-kind developments, and most space programs contain fewer than five copies (some exceptions include the GPS and Iridium constellations).

As a result, there is limited data with which to estimate what a typical learning rate should be for small quantities of spacecraft. In some cases, the answer is "zero" observed learning. This may be due to volatility in the manufacturing of small quantities, even though a longer-term trend could be achieved if more units were built. Also, there are frequently few incentives to reduce unit costs when only a handful of copies will be produced. Finally, unlike high rate production industries, the rate of production of spacecraft is often measured in years, allowing plenty of time between production units, which tends to lead to a reduction in learning. Therefore, the space system cost community typically uses a relatively flat 95% learning rate for small quantities of spacecraft.

But, the space cost community is now faced with a paradigm shift. Future space systems are being proposed with unprecedented production runs in the hundreds and thousands of units. With these kinds of numbers, learning curve assumptions can drive huge swings in cost estimates. At the same time though, the learning rate to be achieved remains one of the least known, unpredictable aspects of the process. Manufacturing will have to change from less specialized, labor intensive shops, to more specialized, and more automated manufacturing lines. But more automation means there is less recurring touch labor from the start, which generally leads to flatter learning curves, while at the same time automation implies lower production costs.

Given these changes, the space cost community still needs an independent way of evaluating the reasonableness of the proposed cost of a high rate production process. In this research, we propose estimating the equivalent learning rate assumption that would need to be accomplished to arrive at the proposed cost estimate, and then to compare the calculated learning rate to learning rates observed in other high rate manufacturing processes in order to determine whether or not the calculated learning rate is reasonable.

4. Historical Learning Rates in Manufacturing

An important contributor to this research is a paper entitled "Learning Curves in Manufacturing[1]," by Argote and Epple, which collected data on the observed cumulative average (Wright) learning rates from more than 100 different manufacturing processes across a wide range of industries. The authors found that the modal learning rate was about 81%, however there was substantial variation around this number ranging from 54% to 108%. Learning rates above 100% are possible as demonstrated by Lockheed's production of the L-1011 Tri-Star in the 1970s with an observed learning rate of 108%. Argot and Epple found reasons for the variation observed in organizational learning curves include organizational "forgetting," employee turnover, transfer of knowledge from other products and other organizations, and economies of scale. Figure 2 is a histogram representing the learning rate frequency derived from the data provided in the Argote and Epple study.

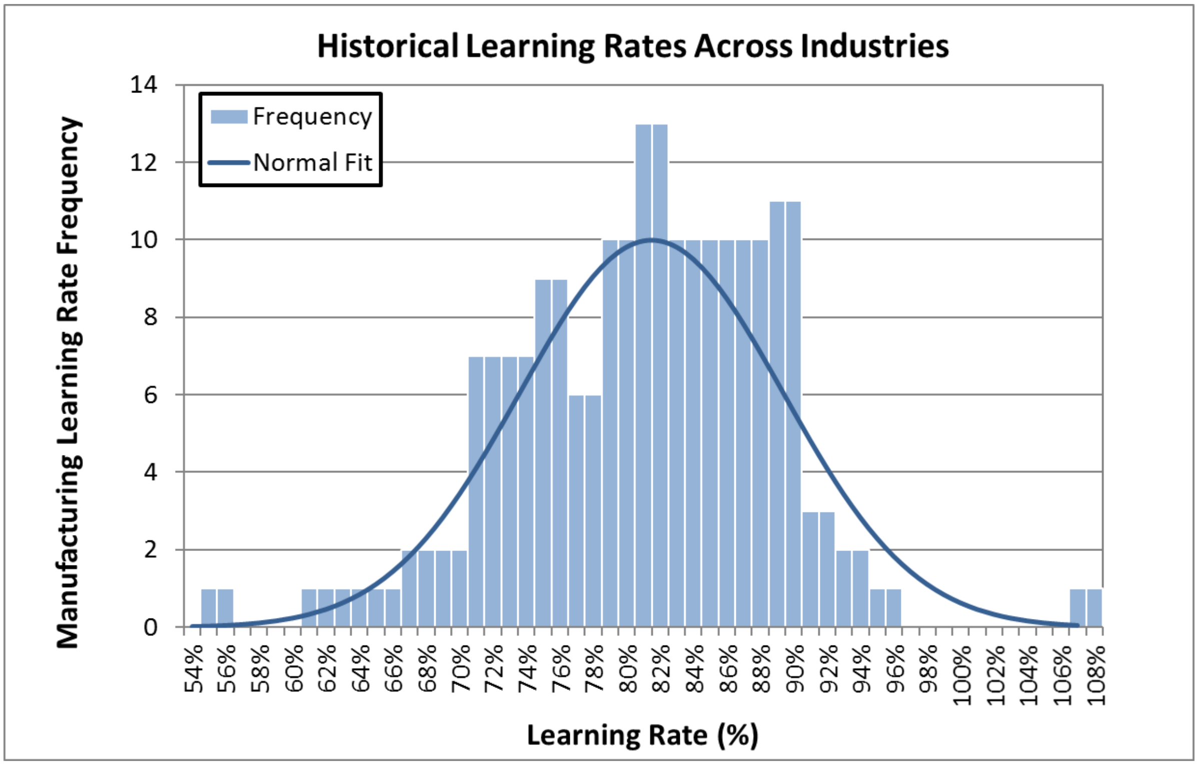

Figure 2 - Historical Learning Rates Across Industries

While production runs of hundreds of units are unprecedented in spacecraft manufacturing, we can possibly learn from other manufacturing fields that build complex products with much larger production sizes (e.g., aircraft, automobiles, shipbuilding, etc.). The objective of this research is to start with traditional cost estimating methods to arrive at an estimated first unit cost, then to determine the learning rate that would be necessary to match a total proposed cost. Upon discovery of the equivalent learning rate, we then compare that rate to learning rates achieved in other manufacturing fields, with the goal of estimating the probability, or likelihood, of the program achieving such a learning rate based on historical data. This result can be used to provide a sense of the reasonableness of the cost of the high production rate system being proposed.

[1] Linda Argote and Dennis Epple, "Learning Curves in Manufacturing," *Science*, vol. 247, no. 4945, 1990, pp. 920-924.

5. The Learning Rate Sensitivity Model

As mentioned previously, in large production runs, the learning rate plays a major role as a cost driver. At the same time, it is one of the least known, unpredictable aspects of the cost estimate. Guidelines for use in typical space system cost estimates have been developed. The most relevant set of guidelines is found in "Historical Cost Improvement Curves for Selected Satellites[2]," by Meisl and Morales, in which they propose broad-based cumulative average theory (Wright) learning rates of 95% for 1-10 units, 90% for 11-50 units, and 85% for 50 or more units, as well as specific learning rates for individual subsystems. While useful for estimating cost, setting the learning rate to some static value ignores a major driver of cost risk due to the volatility of cost, especially in large quantities, resulting from small changes in learning rate assumptions.

Acquisition decision makers need information as to how sensitive a production cost estimate is to the assumed learning rate. The Learning Rate Sensitivity Model is constructed with the goal of providing that information.

The Learning Rate Sensitivity Model, as implemented at Aerospace, is embedded in our Concept Design Center (CDC) cost model. The CDC is an Aerospace-developed resource that allows for real-time, integrated design of space system concepts. The CDC utilizes subject matter experts from major subsystems, and their engineering models, to solve complex design problems. As concept designs are completed, their technical characteristics are inserted into the CDC cost model, which then develops a rapid parametric cost estimate of the design. One of the many variables in the CDC cost model is the assumed learning rate for the spacecraft hardware development. These learning rate assumptions are fixed in accordance with spacecraft estimating guidelines.

The Learning Rate Sensitivity Model was developed by Nichols Brown using Visual Basic for Applications and was recently integrated into the CDC cost model to enable sensitivity analysis of the CDC cost estimate as a function of the learning rate assumptions. In essence, the sensitivity model computes individual CDC cost estimates using learning rates that span the range of learning rates found in industry, from 54% to 108%, in one percent steps, resulting in a table of cost estimates vs. learning rate assumptions. The model then identifies the implied learning rate of the original cost estimate using a lookup function from the table of cost estimates. Moreover, the model computes the cumulative probability of achieving such an implied learning rate based on the industry data shown in Figure 2. Finally, the model plots the cost estimates from the cost table to illustrate the corresponding cost estimates that would arise across the range of observed industry learning rates.

5.1 Example

Consider the following example of a CDC cost estimate for a generic commercial satellite program containing a large number of units: 500 spacecraft inserted into orbit via 25 launches, with a high production rate. The cost estimate table is provided in Figure 3.

[2] Peter Meisl and Lana Morales, "Historical Cost Improvement Curves for Selected Satellites: Final Report," *Management Consulting and Research, Inc.*, TR-9338/029-1, 1994.

Mean Parametric Cost Estimate for the Commercial Class D Program					
	Sat+Grnd Dev	Sat+Lnch Prod	Total	Sat T1	Sat Ta for 500
Total Cost (FY18$M)	$216	$8,501	$8,717	$56	$13
Total Cost (FY18$K)	$216,012	$8,501,137	$8,717,149	$55,826	$13,002
SPACE SEGMENT (FY18$K)	$216,012	$6,177,692	$6,393,704	$53,048	$12,355
Payloads	$34,314	$3,221,773	$3,256,087	$27,666	$6,444
Communication System	$34,314	$3,221,773	$3,256,087	$27,666	$6,444
Bus	$13,901	$1,915,873	$1,929,775	$16,452	$3,832
Propulsion	$505	$133,038	$133,543	$1,142	$266
ADCS	$2,157	$283,243	$285,400	$2,432	$566
TT&C	$1,157	$134,778	$135,935	$1,157	$270
C&DH	$2,445	$284,723	$287,168	$2,445	$569
Thermal	$476	$51,185	$51,661	$440	$102
Power	$3,471	$632,300	$635,771	$5,430	$1,265
Structure	$3,689	$396,606	$400,296	$3,406	$793
Flight Software	$138,552		$138,552		
Integration, Assembly & Test	$10,300	$438,961	$449,261	$3,769	$878
Program Level	$18,945	$601,085	$620,030	$5,162	$1,202
LAUNCH SEGMENT (FY18$K)		$2,323,445	$2,323,445	$2,777	$647

Figure 3 - Cost Estimate of a High Production Rate Program

The total cost is estimated at $8,717M (FY18) when applying the learning rate guidance from Meisl and Morales, in this case an 85% cumulative average theory learning rate for spacecraft production greater than 50 units. (Note no learning is assumed for launch vehicles.) Upon completion of the initial estimate, the Learning Rate Sensitivity Model is activated, which cycles the cost estimate through all of the learning rates experienced in industry (54% to 108%), and the following table of cost estimates vs. learning rates is produced.

Table 1 - Cost Estimates vs. Learning Rates

Learning Rate Assumption	Mean Total Cost (FY18$M)	Learning Rate Assumption	Mean Total Cost (FY18$M)	Learning Rate Assumption	Mean Total Cost (FY18$M)
54%	$2,285	72%	$3,347	90%	$14,437
55%	$2,298	73%	$3,518	91%	$16,048
56%	$2,314	74%	$3,714	92%	$17,859
57%	$2,331	75%	$3,936	93%	$19,893
58%	$2,352	76%	$4,190	94%	$22,177
59%	$2,376	77%	$4,478	95%	$24,739
60%	$2,404	78%	$4,806	96%	$27,610
61%	$2,436	79%	$5,178	97%	$30,827
62%	$2,473	80%	$5,601	98%	$34,428
63%	$2,516	81%	$6,080	99%	$38,456
64%	$2,566	82%	$6,622	100%	$42,958
65%	$2,623	83%	$7,237	101%	$47,988
66%	$2,689	84%	$7,932	102%	$53,601
67%	$2,765	85%	$8,717	103%	$59,863
68%	$2,852	86%	$9,604	104%	$66,843
69%	$2,952	87%	$10,605	105%	$74,619
70%	$3,067	88%	$11,734	106%	$83,274
71%	$3,198	89%	$13,005	107%	$92,902
				108%	$103,605

The table shows what the cost estimate would be for a given cumulative average (Wright) learning rate, for example an 85% learning curve assumption would lead to a mean cost estimate of $8,717M (FY18).

Now we overlay the cost estimates vs. learning rates with the industry learning rate data in Figure 5. This is a combination of the cost estimate vs. learning rate from Table 1 (red curve) superimposed, for historical context, with the observed learning rates from Figure 2 (blue histogram).

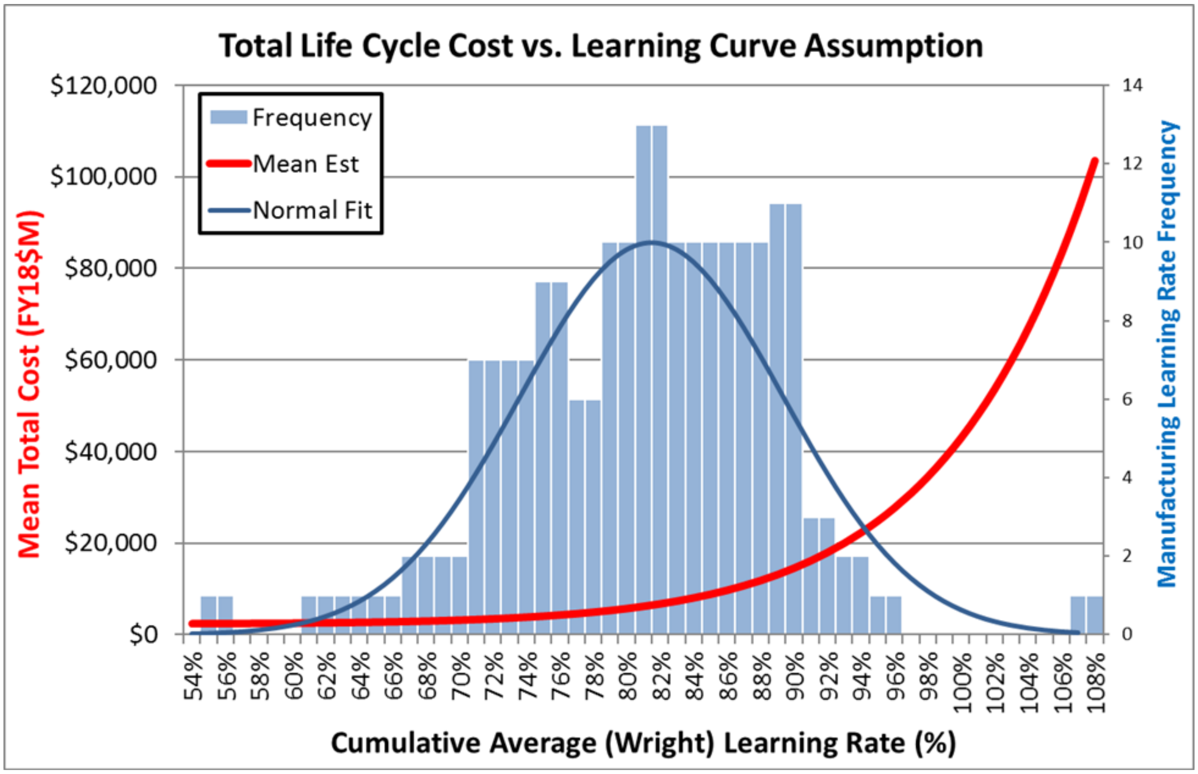

Figure 4 - Total Cost vs. Learning Curve Assumption

The cost estimate vs. learning rate (red curve) illustrates graphically the sensitivity analysis of the cost estimate as a function of the assumed learning rate. The most obvious use of this curve is to provide the decision maker with an assessment of the sensitivity of the cost estimate to the learning rate. That is, if the decision maker has his mind set on estimating the future cost of this system using a specific learning rate (e.g. near 100% due to its automated processes), one can easily display that this assumption should lead to a much higher cost than what is currently estimated.

However, suppose now that the spacecraft developer has a proposed cost for this high production rate system of, say, $3,000M (compared to the buyer's estimate of $8,717M), while claiming an 80% cumulative average (Wright) learning rate. This can be shown to be an optimistic estimate by the proposer. An 80% learning rate implies a cost of $5,601M, while a $3,000M proposed cost implies a learning rate of about 69% to 70%. The decision maker should come away from this thinking that the developer will need to experience a very aggressive learning rate (70%) in order to deliver at $3,000M, or that the developer will have to have a very low, optimistic first unit cost for the total production cost be only $3,000M while assuming an 80% cumulative average (Wright) learning rate.

6. Conclusion

The Learning Rate Sensitivity Model described herein is presented as one of the tools used by The Aerospace Corporation to assess the reasonableness of proposed cost estimates. The model is useful in evaluating the credibility of one or more cost estimates which might have substantially different equivalent learning rates, providing a basis for assessing the reasonableness of learning assumptions used in the competing estimates. It can also be used to estimate the sensitivity of cost estimates to learning assumptions in large quantity acquisitions when little is provided except a stated total production cost.

One weakness of this study, and an area for further research, is that the historical data shown in Figures 2 and 4 are predominantly large hardware and labor intensive systems. These types of systems may be comparable to traditional spacecraft manufacturing methods. But, as spacecraft sizes tend toward microsats and cubesats, with emphasis on electronics design rather than hardware design, the traditional learning curve theories described herein may not adequately apply.

7. Authors

Nichols F. Brown

Nichols (Nick) Brown is an Engineering Specialist in the System of Systems Engineering Office at The Aerospace Corporation, and a professional aerospace engineer and cost analyst with over 17 years' experience. He is co-developer of Aerospace's Small Satellite and Concept Design Center cost models. He has led cost, schedule, and programmatic risk assessments for NASA programs ranging from Small Explorer missions to large joint-agency and flagship programs on behalf of NASA Headquarters' Science Mission Directorate.

Timothy P. Anderson

Timothy Anderson is Director of the Integrated Cost and Schedule Analysis Department at The Aerospace Corporation, and a professional cost analyst with over 22 years' experience, primarily in Department of Defense (DoD) weapon systems and national security space acquisition. He has also served as the Aerospace Director of NASA Program Assessments, managing programmatic assessments for NASA Science Mission Directorate projects. Tim is a frequent presenter of topics related to cost estimating and cost uncertainty analysis.