

LEARNING-CURVE

Estimation of Production Costs and Labor Hours Using a Free Excel Add-In

BY CHARLES D. BAILEY, PH.D.



Despite the prominence of learning curves in a few industries such as aerospace, the technical complexities of implementation have hampered their widespread use in other areas. All types of service businesses, construction companies, and light manufacturers can benefit from these tools. They are appropriate for projecting costs, labor times, and quality measures wherever improvement occurs

through experience. This article reviews their use through a variety of examples, which are easy to solve using an Excel® add-in developed by the author and available free on the Internet.

ORGANIZATIONS of all sizes and descriptions can benefit from using learning curves—service businesses, small businesses, local governments, and light manufacturers. But one obstacle to widespread use has been the tediousness of the content of available books and resources—users must manipulate logarithms or mysterious tables.¹ I have developed some software that is easy to use (available free from my website) to handle the tedious calculations.

Large corporations have industrial engineers who are familiar with learning curve techniques. While the software introduced here is entirely appropriate for their use, these experienced users need no examples to demonstrate applicability. The examples below emphasize smaller businesses and non-traditional settings.

BACKGROUND

It is no surprise that as a person or team repeats an activity they become faster at it. Learning occurs through practice, and thus experienced workers should produce more output of a product or service per hour. Cost also should decrease—the cost of the 10th unit should be less than the cost of the first and so on (*quality* measures also should follow a learning curve).

The surprising insight is that the relationship between practice and improvement is rather predictable. By knowing certain data for each of the

first few units, one can make useful predictions about the production of future units (cost, time, or quality measures).

A couple of points on learning curves are worth noting. The first is that for a typical production activity, learning occurs fastest during the early stages and then levels off. Although improvement theoretically continues indefinitely, it eventually becomes almost imperceptible—a phenomenon called “plateauing.” Thus learning curves are most useful during early phases of production. The second point is that the curve can apply to an individual worker or a team of workers. Furthermore, it may apply even to the entire pool of costs associated with an activity because learning occurs not only in labor skills but also in operating systems, paperwork systems, and so on.

Learning curves can answer many types of questions: Service Businesses

- ◆ A new clerical department of an insurance agency has been processing policy applications for one month. *Will it be able to keep up with anticipated growth?*
- ◆ A beginning fast-food cook requires 60 minutes to prepare her first 20 hamburgers, 45 minutes for the second 20, and 36 minutes for the third 20. *What should her production rate be after 24 hours of experience?*

- ◆ A new check-encoding clerk in a bank required 60 minutes to encode his first 500 checks, 50 minutes for the second 500, and 45 minutes for the third 500. *When will this employee be able to produce at the standard rate of 1,000 per hour?*

Production Businesses

- ◆ An electrical contractor has wired two identical homes. *How long should the same team require to wire the third identical home? The 10th?* (Or, as the general contractor, how many labor hours should you expect the electrical contractor to include in a bid?)
- ◆ A custom boat builder has produced only a single prototype of a new sailboat. But the producer knows, from accumulated past experience, the learning curve rate for similar boats. *What are the projected labor requirements for the second, third, and so on?*

There is no one “learning curve,” although one model seems to be the most popular. A curve is simply an attempt to describe observed learning behavior by a mathematical equation. (I won’t get involved in the math here. For those interested, the manual accompanying the free software covers that topic.)

The software package, called Foresee[®], provides four kinds of learning curves, all of which have exactly the same purpose: to forecast future times (or dollars or quality measures) for each unit of production. Two of the curves use the traditional, well-known logarithmic form. The other two use a form introduced by Bailey and McIntyre² that fit better

and did a better job predicting the mechanical assembly tasks used in their experiments.

Using either form, you can fit the curve to *average* data or to *marginal (unit)* data. After forecasting either the average or the unit data, conversion to unit times, totals, and so forth is easy. This software produces unit time estimates *from whatever form of curve you use*. Confused? You’re not alone. Although I have just stated the distinction between the curves, it is hard to grasp. Academic articles have addressed the topic, but I have yet to find an accounting textbook that distinguishes clearly between the curve forms.

Let’s take a look at one learning-curve relationship, the traditional log-linear marginal (unit) curve. It assumes that every time the total number of units produced *doubles*, the marginal cost *decreases by a constant proportion* (called the learning curve percentage). A numerical example appears in Table 1, and Figure 1 shows a graph of the data. These tabular and graphical presentations demonstrate the relationship, but equations are necessary to generate forecasts about units other than the “doubling points.”

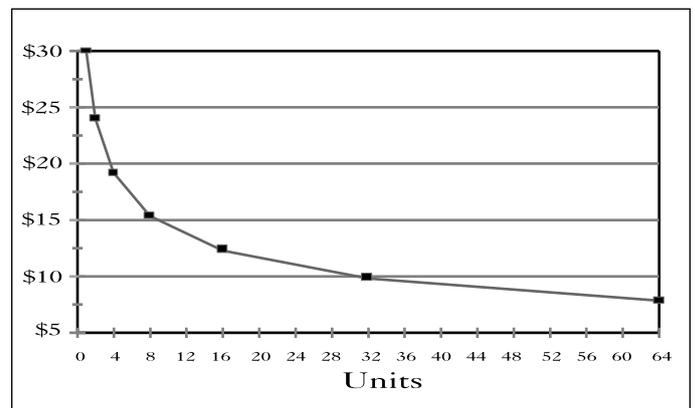
USING THE SOFTWARE

Foresee[®] is an Excel[®] add-in, written using the spreadsheet’s Visual Basic for Applications (VBA) facility. A manual accompanying the software explains the mechanics of entering data and navigating the features of the application. Except for a brief explanation following the first question, the examples below bypass these mechanical steps to show solutions to typical problems.

Table 1: Traditional Log-Linear Marginal (Unit) Learning Curve

| Cumulative Production (units) | Marginal Cost | Calculation, assuming an 80% curve |
|-------------------------------|---------------|------------------------------------|
| 1 | \$30.00 | |
| 2 | 24.00 | .80 x 30 = 24 |
| 4 | 19.20 | .80 x 24 = 19.20 |
| 8 | 15.36 | .80 x 19.20 = 15.36 |
| 16 | 12.29 | .80 x 15.36 = 12.29 |
| 32 | 9.83 | .80 x 12.29 = 9.83 |
| 64 | 7.86 | .80 x 9.83 = 7.86 |

Figure 1: Traditional Log-Linear Marginal (Unit) Learning Curve



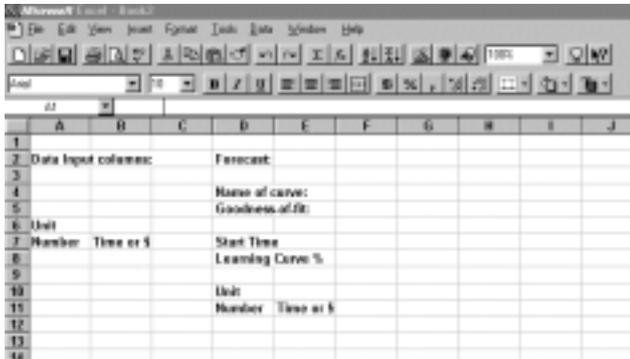


Figure 2: Basic Worksheet in Foresee[®]

Foresee[®] can perform learning-curve projections using either of two starting points.

- ◆ An assumed rate of improvement. This is appropriate for industries, such as aerospace, that have considerable experience with learning curves so that they “know” the learning rate. It also is appropriate for government contracts in which a certain rate is assumed.
- ◆ Raw data from actual experience. This is appropriate for tasks on which one has little historical experience, such as the nontraditional applications emphasized here.

After navigating the introductory menu in Foresee[®], you will reach a worksheet like the one shown in Figure 2. The columns at the left of Figure 2,

| | A | B | C | D | E | F | G |
|----|---------------------|------------|---|------------------|----------------|-----------------------|---|
| 1 | | | | | | | |
| 2 | Data Input columns: | | | Forecast: | | | |
| 3 | | | | | | | |
| 4 | | | | Name of curve: | | Marginal (Unit) Curve | |
| 5 | | | | Goodness of fit: | | 88.85% | |
| 6 | Unit | | | | | | |
| 7 | Number | Time or \$ | | Start Time | | 32.13 | |
| 8 | 1 | 29.8 | | Learning Curve % | | 67.40% | |
| 9 | 2 | 25 | | | | | |
| 10 | 3 | 17.5 | | Unit | | | |
| 11 | | | | Number | Time or \$ | | |
| 12 | | | | 5 | 12.86 | | |
| 13 | | | | 6 | 11.59 | | |
| 14 | | | | 7 | 10.62 | | |
| 15 | | | | 8 | 9.84 | | |
| 16 | | | | 9 | 9.20 | | |
| 17 | | | | 10 | 8.67 | | |
| 18 | | | | 11 | 8.21 | | |
| 19 | | | | 12 | 7.81 | | |
| 20 | | | | 13 | 7.46 | | |
| 21 | | | | 14 | 7.16 | | |
| 22 | | | | 15 | 6.88 | | |
| 23 | | | | 16 | 6.63 | | |
| 24 | | | | 17 | 6.41 | | |
| 25 | | | | 18 | 6.20 | | |
| 26 | | | | 19 | 6.01 | | |
| 27 | | | | 20 | 5.84 | | |
| 28 | | | | Total | 217.09 Minutes | | |

Figure 3: Solution to Question #1

starting with cells A8 and B8, are for entry of unit-cost data from past experience. (Alternatively, you can assume a learning rate by using an “Advanced Forecasting” option.³) Forecasts will appear in the columns at the right, starting with D12 and E12. The remaining columns are available for supplemental calculations, which can use all the regular capabilities of Excel.

FORMULATING QUESTIONS

Any question someone might ask can be addressed by using the estimated individual unit (marginal) times (or dollars or quality measures). A future version of the program might provide additional columns with total times and average times, or it might include dialogue boxes to address a variety of questions directly. But in the name of simplicity, the current version provides unit estimates only.

As noted above, the program can base its answers on two sources of information: (1) data reflecting past performance or (2) the initial time and learning rate (percentage) that you can enter directly. Here are several examples that illustrate approaches to formulating and answering typical questions:

Question #1: A worker has performed a mechanical assembly task four times. What is the projected total time for the first 20 units assembled, based on the actual times observed for the first four assemblies?

Figure 3 shows the answer I derived using actual data, in minutes, from a laboratory experiment I conducted. The steps to achieve the output in Figure 3 were as follows:

1. Enter the data in the left-hand columns.
2. Select “Tools; Learning Curve App; Fit Learning Curve.”

This Fit Learning Curve option fits a curve to the data on the current worksheet. If desired, the user can check the “Specify Curve Type” box to select the curve form. Otherwise, the program will choose the curve that best fits the data.⁴

The New option opens a new worksheet.

Advanced Forecast allows forecasting with a known learning rate.

3. Enter the numbers 5 through 20 in the forecast column.
4. Select “Tools; Learning Curve App; Forecast.”
5. Sum the *actual* times 1 through 4 and the *forecasted* times 5 through 20. (I could have produced forecasts for times 1 through 4, but that would

make little sense when I have the actual times.)

Question #2: How fast will John (worker from Question #1) be after he has completed 100 units?

An appropriate measure of ability after completing n units would be the projected time for unit $n+1$, in this case unit #101. To address this question, simply enter the unit number, 101, in cell D12 and perform step 4 from Question #1. The resulting estimate of 2.32 minutes will appear in cell E12.

| Data Input columns: | | Forecast: | |
|---------------------|------|--------------------------------------|--------|
| | | Name of curve: Marginal (Unit) Curve | |
| | | Goodness-of-fit: 99.48% | |
| Unit | | | |
| Number | Time | Start Time | 60.50 |
| 1 | 60 | Learning Curve % | 72.72% |
| 2 | 45 | | |
| 3 | 36 | Unit | |
| | | Number | Time |
| | | 4 | 31.99 |
| | | 5 | 28.88 |
| | | 6 | 26.55 |
| | | ... | ... |
| | | ... | ... |
| | | 123 | 6.63 |
| | | 124 | 6.60 |
| | | 125 | 6.58 |
| | | 126 | 6.55 |
| | | 127 | 6.53 |
| | | Minutes | Hours |
| | | 172.99 | 2.88 |
| | | 201.87 | 3.36 |
| | | 228.42 | 3.81 |
| | | ... | ... |
| | | ... | ... |
| | | 1431.94 | 23.87 |
| | | 1438.54 | 23.98 |
| | | 1445.12 | 24.09 |
| | | 1451.67 | 24.19 |
| | | 1458.20 | 24.30 |

Figure 4: Solution to Question #3

Question #3: A beginning fast-food cook, working constantly, required 60 minutes to produce her first 20 hamburgers, 45 minutes for the second 20, and 36 minutes to produce her third batch of 20. What is her predicted rate after 24 hours of experience at the task?

The solution lies in forecasting into the future and summing the unit times until they equal 24 hours. The result using hypothetical data appears in Figure 4. I have created the two columns at the far right, labeled “Minutes” and “Hours,” to show cumulative times. When the total reaches 24 hours, the estimated marginal time is 6.58 min./batch, or about 9.1 batches/hr. This is the worker’s projected speed after that much practice. A physical constraint, such as cooking time or equipment limitations, might prevent the achievement of a 9.1 batches/hr. speed.

Question #4A: A newly hired check-encoding clerk in a bank required 66 minutes to encode his first 500 checks, 59.4 minutes for the second

500, and 51.6 minutes for the third 500. After how much time will this clerk reach the “standard” rate of three batches per hour (20 min./batch)?

The solution appears in Figure 5. Projecting unit times well into the future shows that the clerk will require about 264 batches to reach the standard 20 min./batch. We could, of course, also sum up the batch times to estimate how long this will take. We can monitor his progress against these projections—and refit the curve to revise our projections as more data become available.

Question #4B: Assume that, as in #4A, checks are processed in batches of 500. We know the total number of checks the clerk has processed since starting the job, but we have recorded performance times only for batches 1, 3, and 5. Batches 1 and 3 are as given in Question #4A, but batch 5 turns out to be 43.5 minutes, not the amount forecasted above.

The “gap” in the data does not matter, assuming a unit curve,⁵ and the solution is like #4A. Cells A8 to B10 would contain the following:

| | |
|---|------|
| 1 | 66.0 |
| 2 | 59.4 |
| 5 | 43.5 |

These data lead to a forecast of 104 batches needed to reach the standard rate versus the 264 batches forecasted in #4A. This difference emphasizes the wide margin of error when relying on only three data points (see the discussion in Question #5 on “goodness of fit”).

| Data Input columns: | | Forecast: | |
|---------------------|------|--------------------------------------|--------|
| | | Name of curve: Marginal (Unit) Curve | |
| | | Goodness-of-fit: 94.80% | |
| Unit | | | |
| Number | Time | Start Time | 66.80 |
| 1 | 66 | Learning Curve % | 86.08% |
| 2 | 59.4 | | |
| 3 | 51.6 | Unit | |
| | | Number | Time |
| | | 4 | 49.49 |
| | | 5 | 47.16 |
| | | ... | ... |
| | | ... | ... |
| | | 262 | 20.03 |
| | | 263 | 20.02 |
| | | 264 | 20.00 |
| | | 265 | 19.98 |

Figure 5: Solution to Question #4a

Question #5: An electrical contractor has wired two identical homes using the same approach and the same team of electricians. The total labor costs were \$8,450 for the first job and \$6,676 for the second job. For bidding purposes, what should the contractor expect the third, fourth, and fifth units to cost?

Two observations are not much data, and the “goodness of fit” is not meaningful because any curve form can fit two data points *perfectly*. With two data points, the program defaults to the traditional marginal curve unless you specify the curve form. The following data are inserted in cells A8 to B9:

- 1 8450
 - 2 6676
- The marginal curve produces forecasts as follows:
- 3 5816.36
 - 4 5274.44
 - 5 4889.11

Even with so few observations, the tool should be useful, although the user needs to exercise judgment.⁶

Question #6A: When starting production of a new circuit board, the Galvanic Board Company at first encounters many rejected boards. The number declines following a learning curve as the operators gain experience and equipment problems are ironed out. What should the company expect given the following data (which are hypothetical, but learning curves are used to track quality improvement)?

- Of the first 100 boards, 10 were bad.
- Of the second 100, 7 were bad.
- Of the third 100, 6 were bad.
- Of the fourth 100, 6 were bad.
- Of the fifth 100, 5 were bad.

The solution appears in Figure 6, which presents straightforward forecasts after selected numbers of units. The failure rate should be cut in half after about 100 units of production and to less than a third by 1,000 units.

Question #6B: Modify the question in #6A to assume that, because of industry benchmarking information, Galvanic Board Company expects the number of rejects to decline according to an 80% cumulative average curve. Assume also that they have made only 200 boards and thus have only the data on the next page.

| Data input columns: | | Forecast: | |
|---------------------|------------|------------------|---------------|
| | | Name of curve: | Average Curve |
| | | Goodness-of-fit: | 98.33% |
| Unit | | Start Time | 10.00 |
| Number | Time or \$ | Learning Curve % | 84.81% |
| 1 | 10 | | |
| 2 | 7 | | |
| 3 | 6 | Unit | |
| 4 | 6 | Number | Time or \$ |
| 5 | 5 | 6 | 5.08 |
| | | 7 | 4.69 |
| | | 8 | 4.73 |
| | | 9 | 4.58 |
| | | 10 | 4.47 |
| | | 20 | 3.77 |
| | | 40 | 3.18 |
| | | 70 | 2.78 |
| | | 100 | 2.56 |
| | | 200 | 2.17 |
| | | 300 | 1.97 |
| | | 1000 | 1.48 |

Figure 6: Solution to Question #6a

| Data input columns: | | Forecast: | |
|---------------------|------------|------------------|--------------------|
| | | Name of curve: | Average Time Curve |
| | | Goodness-of-fit: | |
| Unit | | Start Time | 10.00 |
| Number | Time or \$ | Learning Curve % | 80.00% |
| | | Actual | Unit |
| | | Units | Number |
| | | 100 | 1 |
| | | | Time or \$ |
| | | 200 | 2 |
| | | | 6.00 |
| | | 300 | 3 |
| | | | 5.06 |
| | | 400 | 4 |
| | | | 4.54 |
| | | 500 | 5 |
| | | | 4.38 |
| | | 600 | 6 |
| | | | 3.92 |
| | | 700 | 7 |
| | | | 3.71 |
| | | 800 | 8 |
| | | | 3.55 |
| | | 900 | 9 |
| | | | 3.41 |
| | | 1000 | 10 |
| | | | 3.29 |

| Data input columns: | | Forecast: | |
|---------------------|------------|------------------|--------------------|
| | | Name of curve: | Average Time Curve |
| | | Goodness-of-fit: | |
| Unit | | Start Time | 17.00 |
| Number | Time or \$ | Learning Curve % | 80.00% |
| | | Actual | Unit |
| | | Units | Number |
| | | 200 | 1 |
| | | | Time or \$ |
| | | 400 | 2 |
| | | | 17.00 |
| | | 600 | 3 |
| | | | 8.61 |
| | | 800 | 4 |
| | | | 7.71 |
| | | 1000 | 5 |
| | | | 7.11 |
| | | 1200 | 6 |
| | | | 6.66 |
| | | 1400 | 7 |
| | | | 6.31 |
| | | 1600 | 8 |
| | | | 6.03 |
| | | 1800 | 9 |
| | | | 5.79 |
| | | 2000 | 10 |
| | | | 5.59 |

Figure 7: Alternate Solutions to Question #6b

Of the first 100 boards, 10 were bad.
Of the second 100, 7 were bad.

What is the benchmark for measuring future improvement?

First, what should we use as a definition of the first batch—100 boards with 10 bad or 200 with 17 bad? Either definition is acceptable. The two results appear in Figure 7, using the “Advanced” option to specify the learning rate. This option requires an initial bit of data, in this case the error rate for an initial batch, the form of curve, and the curve percentage (80%). Note that the spreadsheet contains no data in columns A and B, as no data are used to fit a curve in this example. Only this example follows the approach typically taught in cost accounting textbooks, that is, estimation using a known percentage rate and starting point.

HOW MUCH CONFIDENCE SHOULD YOU PLACE IN THE SOLUTION?

As implied above, learning-curve estimates generally carry a high level of reliability or predictability. This level will vary, however, from application to application. Some factors that affect reliability include:

- ◆ Errors in measurement or collection of data;
- ◆ Random factors, unrelated to learning, that affect performance for a particular repetition;
- ◆ Subtle changes in the production activity itself, such as design changes; and
- ◆ The number of data points used to obtain the estimate.

A useful indication of the reliability of the estimate is the “goodness of fit” of the formula used in the estimate. In statistical terminology, this is called r^2 , the “coefficient of determination.” Stated simply, it is a number between zero and one that indicates the proportion of variation in production time (or cost) that is explained by knowing the number of units produced. For example if $r^2 = .70$, then the experience level, or the number of units produced so far, explains 70% of the fluctuation in production time. The remaining 30% of the overall fluctuation (variance) results from the errors in measurement and so on listed above. For small samples, the curve must fit very well (explain a lot of the variance) to be reliable. While a larger r^2 always is better, the following is a rule of thumb for the absolute minimum to reflect reliability: for two data points it is impossible to evaluate reliability; for three, the minimum acceptable r^2 is 0.98; for four, 0.81; for five, 0.65; for

10, 0.30; and for 40 or more, 0.10.

THE POWER OF LEARNING CURVES

The potential applications of learning curves far outstrip their current usage. Businesses from fast-food restaurants to banks can determine how fast their employees will obtain on-the-job skills. Local governments can improve budgeting processes. Manufacturers can estimate production rates. With software to aid in learning curve use, all the benefits of the technique are within reach. Organizations cannot afford to ignore such a powerful tool. ■

Charles Bailey, Ph.D., is a professor in the School of Accounting, University of Central Florida, Orlando, Fla. You can download his software and the related manual for free from his website, www.bus.ucf.edu/bailey. The Mid-Florida Chapter of the IMA has provided funding for a revision of the software, and that new version also is available on the website. He would appreciate feedback, which may be sent to Charles.Bailey@bus.ucf.edu.

- 1 Articles in *Management Accounting* that have addressed learning curve use include Eugene A. Imhoff, Jr., “The Learning Curve and Its Applications,” February 1978, pp. 44-46; and Jackson F. Gillespie, “An Application of Learning Curves to Standard Costing,” September 1981, pp. 63-65. A more recent article that illustrates the broad applicability of the learning phenomenon by applying it to the learning of new computer systems is K.M. Boze, “Measuring Learning Costs,” August 1994, pp. 48-52.
- 2 C.D. Bailey and E.V. McIntyre, “Some Evidence on the Nature of Relearning Curves,” *The Accounting Review*, April 1992, pp. 368-378.
- 3 The “advanced” title is something of a misnomer. The program to fit a curve to new data really is a more “advanced” feature. The name implies a need for caution, though; the user needs advanced knowledge of a learning rate before using this option.
- 4 The option of selecting the curve form lets advanced users override the program’s choice of a best-fit curve. Sometimes a few data points may not reveal the underlying improvement rate.
- 5 An average curve is not feasible because it requires a cumulative average that would include the missing fourth observation.
- 6 For example, if the workers showed little or no improvement, the analyst should not blindly extrapolate this trend because temporary setbacks may be followed by large improvements. Workers may have learned without any improvement showing in the most recent performance. One reason may be because of errors committed. See C. D. Bailey and S. Gupta, “Judgement in Learning-Curve Forecasting: A Laboratory Study,” *Journal of Forecasting*, January 1999, pp. 39-57.