TIME & COST ESTIMATION WITH LEARNING CURVES: NEW SOFTWARE FOR SMALL AND SERVICE BUSINESSES

by

Charles D. Bailey, CMA, Ph.D.
Professor of Accounting
School of Accounting
University of Central Florida
Orlando, FL 32816-1400

Cbarles.Bailey@bus.ucf.edu
http://www.bus.ucf.edu/acc/cbailey/

September 7, 2000

Funded by a grant from the University of Central Florida College of Business.

Table of Contents

TIME & COST ESTIMATION WITH LEARNING CURVES: NEW SOFTWARE FOR SMALL AND SERVICE BUSINESSES

What Is a "Learning Curve"?

It is no surprise that, as a person or group repeats an activity, they become faster at it. Learning occurs through practice. Thus an experienced worker or group should produce more output per hour, the cost of the tenth unit should be less than the cost of the first, and so on.

The *surprising* fact is that the relationship between practice and improvement is rather predictable. By knowing the time taken to produce each of the first few units, one can make useful predictions about the cost of future units. Recognition of this phenomenon is attributed to T. P. Wright (around 1935), and learning curves were important to the Allies' military planning during WWII.

Subsequently, aircraft and other defense contractors—as well as governmental agencies—have benefitted heavily from learning curve techniques. Unfortunately, the techniques have been less accessible to many users who could benefit.

Who Can Benefit from this Package?

Organizations of all sizes and descriptions can benefit from using learning curves. These include service businesses, small businesses, local governments, and light manufacturers. But one obstacle to wide usage has been that the available texts and resources are tedious, requiring the manipulation of logarithms or arcane tables. The purpose of this package is to present a clear explanation of the technique and its uses, while providing software 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 in this manual therefore emphasize smaller businesses and nontraditional settings. Here are a few of the types of questions that you can expect to answer:

Service Businesses

A new clerical department of an insurance agency has been processing policy applications for one month. Will they be able to keep up with anticipated growth?

A beginning fast-food sandwich maker requires one hour to prepare her first 20 sandwiches, 0.75 hr. for the second twenty, and 0.60 hr. for the third twenty. What should her production rate be after 24 hours of experience?

A new check-encoding clerk in a bank required one hour 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 1000 per hour?

Production Business

An electrical contractor has wired two identical homes. How long should the same team require to wire the third identical home? The tenth? (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, etc.?

When is the Learning Curve Useful?

For a typical production activity, learning occurs fastest during the early stages, then levels off. Although improvement theoretically continues forever, it eventually becomes almost imperceptible — a phenomenon called "plateauing." For this reason, learning curves are most useful during early phases of production.

Another point worth noting is that the curve can apply to an individual worker or a group of workers. Furthermore, it may apply even to the entire pool of costs associated with an activity. This is because learning takes place not only in labor skills but also in operating systems, paperwork systems, and so on.

A Basic Description of Learning Curves

Actually, 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 through a mathematical equation. (You need not become involved in the math, but an appendix covers it if you are interested.)

This package provides four alternative forms of learning curves, all of which have exactly the same purpose: to forecast future times (or dollars) for each unit of production. Two of the curves use the traditional, well known logarithmic form. The other two are of a form introduced by Bailey and McIntyre[1] and found to fit and predict well for the mechanical assembly tasks used in their experiments. **One motive for providing this software free of charge is that I am very**

---

[1] Bailey, C. D., and McIntyre, E. V., "Some Evidence on the Nature of Relearning Curves," *The Accounting Review,* **67** (1992), 368-378.

interested in feedback from you about how well the curves fit and predict *your* data. Please see "Sending Data to the Author for Research Purposes," on p. 14.

Let's take a look at one learning-curve relationship, the traditional log-linear marginal (unit) curve. It assumes that, every time the total (cumulative) number of units produced *doubles, the marginal cost decreases by a constant proportion* (called the learning curve percentage).

Here is a hypothetical example of an 80% **marginal (unit)** curve:

Cumulative Units Made	**Marginal Cost (000's)**	**Calculation**
1	$30.00	[Initial unit cost.]
2	$24.00	0.80 X 30 = 24
4	$19.20	0.80 x 24 = 19.20
8	$15.36	0.80 x 19.20 = 15.36
16	$12.29	0.80 x 15.36 =12.29
32	$9.83	0.80 x 12.19 = 9.83
64	$7.86	0.80 x 9.83 = 7.86

Graphed, this relationship appears as follows:

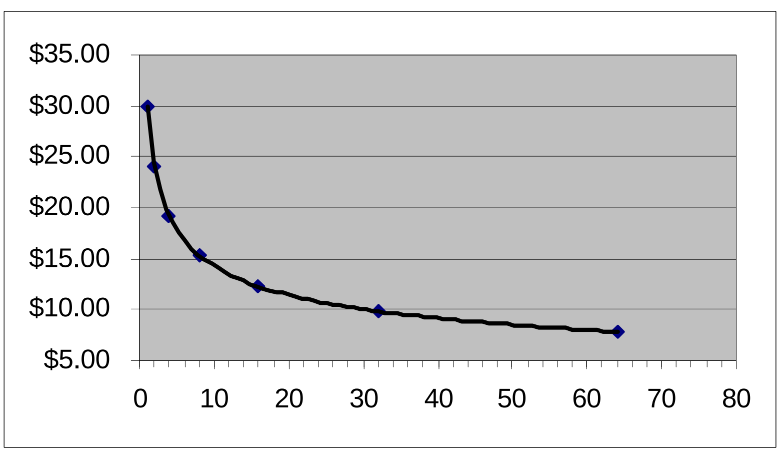

The tabular and graphical presentation demonstrated the nature of the relationship, but equations are necessary to generate forecasts about units other than the "doubling points" above.

From this basic relationship, we can derive other values of interest, such as the total cost, the marginal cost of any specific unit, the cost of a group of units, and so forth.

Reliability of Results

How much confidence should you place in the forecast? As implied in the introduction 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 the following:

< 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
< 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 r^2. called the “coefficient of determination.” It is a number between 0.00 and 1.00, indicating 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 70 percent of the fluctuation in production time is explained by the experience level (i.e., by the number of units produced so far). The remaining 30 percent of the overall fluctuation (variance) in production time is <u>un</u>explained by the prediction formula; it results from the errors in measurement, etc., listed above. The value of r^2 is based upon the sample data that you enter. The larger your sample (**N**), the more confidence you can place in the resulting estimates. The following is a rule of thumb (Draper and Smith, 1981, p.93):

N	Minimum R-square
2	Impossible to evaluate
3	0.99
4	0.94
5	0.88
10	0.63
20	0.40
30	0.29
40	0.23

Limitations and Caveats

This package is intended for users without formal training in statistics and therefore contains explanations that statisticians could criticize.

Statistical models such as this one are easily misused. The software uses “ordinary least-squares” regression to estimate model parameters, and the underlying assumption of this technique must be met. A discussion of these assumptions is beyond the scope of this manual, but there is good reason to believe that the model generally is appropriate for use in this setting.

It is important to remember that the software produces estimates based upon past sample data, and not to place too much certainty upon the numbers. Perhaps the biggest potential cause of bad prediction is a violation of the assumption that the environment is stable. To the extent that the task itself or the persons performing it have changed (other than through learning), the model's estimates based upon past performance will be inappropriate.

Software License

The **foresee.xla** program remains the property of the copyright holder, Dr. Charles Bailey. It is offered as “freeware” as a service to the business and academic community. You may use it and distribute it freely, as long as it is unmodified and the source is acknowledged. You may not sell it or incorporate any part of it into another product.

To help with further research, I ask that you share with me your learning-related data sets. Your confidentiality would, of course, be a high priority. If you have concerns about confidentiality, I would like to discuss them with you. There are many options, such as simply reporting goodness-of-fit information to.

Using the Software

This program functions as an add-in to Excel®, and was written using Visual Basic for Applications (VBA). To use it, open Excel and open the file called **foresee.xla**. The add-in module will be available until you close Excel.

When you open foresee.xla, you will see the dialogue box at the right:

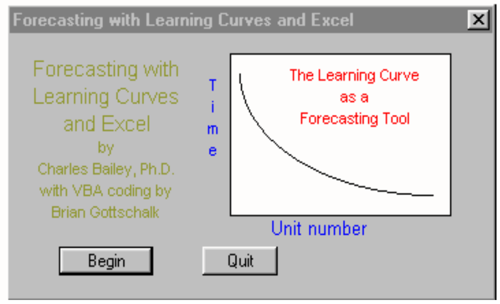

Clicking Begin will take you to a worksheet like the following:

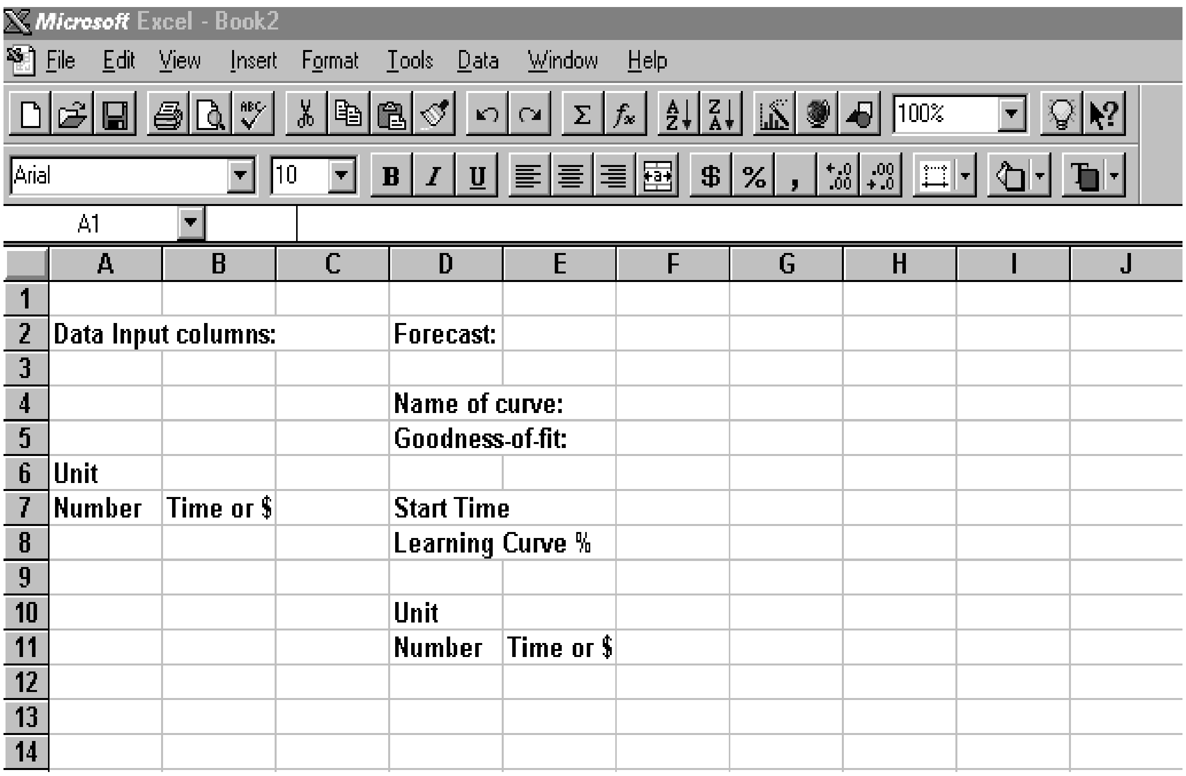

Columns A and B are for data entry. Forecasts will appear in columns D and E. The remaining columns are available for your use in supplemental calculations, using the regular capabilities of Excel.

Formulating Questions

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

To get you started in formulating and answering typical questions, here are several examples. Note that the program can base its answers upon two alternative sources of information: (1) data reflecting past performance, or (2) the initial time and learning rate (percentage) that you can enter directly. Option (2) is appropriate for users who wish to assume a given learning rate based on past data, typical figures for their industry, or the requirements of a government contract.

Examples

Here are several examples of problems and solutions. Question #1 uses actual data from a laboratory experiment that I conducted using paid subjects. The others use hypothetical data to address a variety of typical problems. (I look forward to you, as a user, sharing other real examples with me.)

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

After performing the steps to be described below, the output appears as shown on the following page.

	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	4	13.4		**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	

Solution to Question 1

The steps to achieve the above output were as follows:

1. Enter the past performance data in columns A & B.

2. Select "Tools; Learning Curve App; Fit Learning Curve," as shown on the menu below. (Your menu's exact appearance will depend on the version of Excel and its customization.)

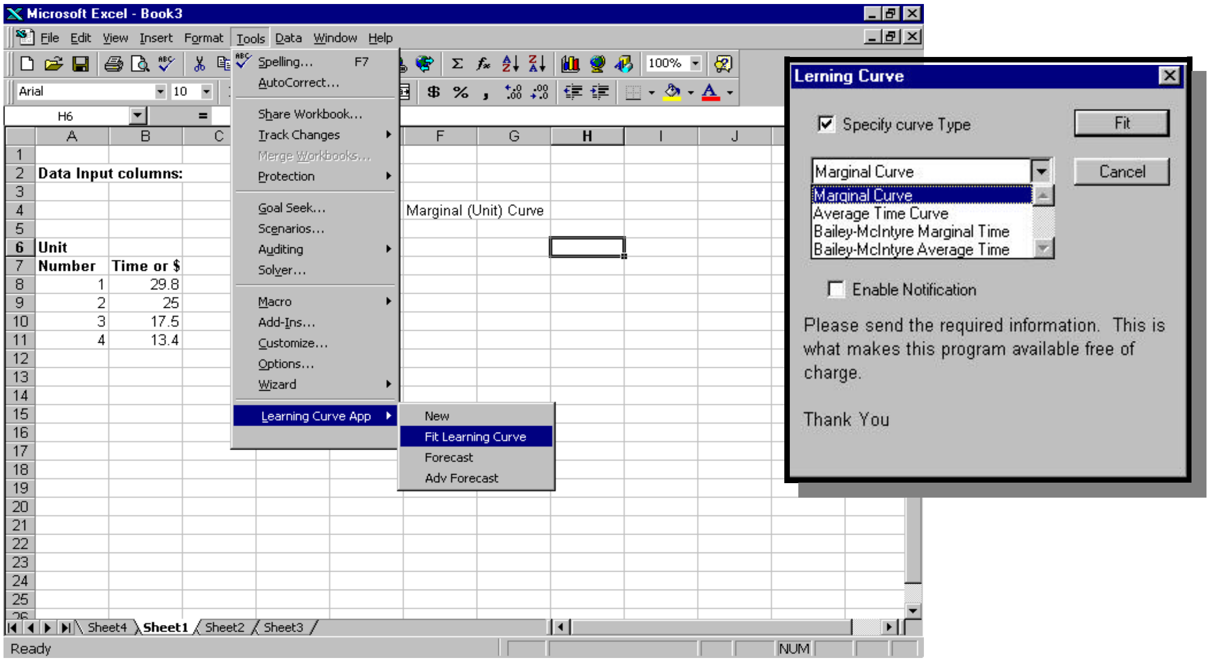

The Fit Learning Curve option fits a curve to the data on the current worksheet, ready to use for forecasting.[2] Checking the "Specify Curve Type" box allows you to select the curve form. Otherwise, the program chooses the curve that best fits the data.[3]

3. Enter the numbers 5 through 20 in the forecast column.
4. Select Tools; Learning Curve App; Forecast
5. Sum the *actual* times 1— 4 and the *forecasted* times 5 — 20. (We could have produced forecasts for times 1 — 4, but that would make little sense when we have the actual times.)

Question #2: How fast will John be after he has completed 100 units (that is, what is the projected time for unit #101)?

This is a simple forecast for unit 101, as shown here. Note that John's skill level is his projected *next* time.

Data Input columns:			**Forecast from best-fitting curve:**			
			Name of curve:		Marginal Unit Curve	
			Goodness-of-fit:		88.85%	
Unit						
Number	**Time**		**Start Time**		32.13	
1	29.8		**Learnig Curve %**		67.40%	
2	25					
3	17.5		**Unit**			
4	13.4		**Number**	**Time**		
			101	2.32		

[2] The New option opens a new worksheet. Adv[anced] Forecast allows forecasting with a known learning rate. 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.

[3] The curve that best fits these few data points may not be the one that an advanced user feels best describes the underlying typical improvement rate for the particular task. This is the purpose of allowing a forced choice.

Question #3: A beginning fast-food sandwich maker, required one hour to produce her first 20 sandwiches, .75 hr. for the second 20, and .60 hr. 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. (This is an occasion to point out a requirement of the program: it cannot accept data with values less than 1.0, and it will warn you if you enter such values. Converting the times from hours to minutes overcomes this limitation.) The result is as follows:

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**	Minutes	Hours
			4	31.99	172.99	2.88
			5	28.88	201.87	3.36
			6	26.55	228.42	3.81
			...	...	...	...
			...	...	...	...
			123	6.63	1431.94	23.87
			124	6.60	1438.54	23.98
			125	6.58	1445.12	24.09
			126	6.55	1451.67	24.19
			127	6.53	1458.20	24.30

Note deleted lines. The answer was reached after trial-and-error to find where the total reached 24 ^{2} hours.

I have added the two rightmost columns, labeled Minutes and Hours, to show cumulative total 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.*

These data are hypothetical, as I have no actual data from a fast-food setting. A physical constraint, such as cooking time or equipment limitations, might prevent the achievement of a 9.1 batches/hr. (6.58 min./batch) speed.

<u>Question #4A</u>: A new check-encoding clerk in a bank required 1.10 hours to encode his first 500 checks, 0.99 hour for the second 500, and 0.86 hour for the third 500. After how much time will this clerk reach the "standard" rate of 3 batches/hr. (0.333 hour/batch)?

As before, I have converted hours to minutes to avoid values less than 1.0. Here is the solution:

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		

[2] Some lines from trial-and-error deleted.

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 becomes available.

__Question 4B__: **Assume that, as in 4-A, checks are processed in batches of 500. We know the total number of checks that 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 4A, but batch 5 turns out to be 43.5, not the amount forecasted above. The "gap" in the data does not matter, and the solution approach is like 4A.**

Data Input columns:			Forecast:			
			Name of curve:		Marginal (Unit) Curve	
			Goodness-of-fit:		95.33%	
Unit						
Number	Time		Start Time		67.87	
1	66		Learning Curve %		83.32%	
2	59.4					
5	43.5		Unit			
			Number	Time		
			6	42.35		
			7	40.66		
			8	39.26		
			9	38.06		
		…	…			
		…	…			
			101	20.14		
			102	20.09		
			103	20.03		
			104	19.98		
			105	19.93		

This data leads to a forecast of 104 batches to achieve standard, versus the 264 batches forecasted in 4-A. This difference emphasizes the wide margin of error when relying on only 3 data points (see discussion under Goodness of Fit, above).

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

Data Input columns:			**Forecast:**			
			Name of curve:		Marginal (Unit) Curve	
			Goodness-of-fit:		100.00%	
Unit						
Number	**Time or $**		**Start Time**		8450.00	
1	8450		**Learning Curve %**		79.01%	
2	6676					
			Unit			
			Number	**Time**		
			3	5816.36		
			4	5274.44		
			5	4889.11		

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 2 data points, the program defaults to the traditional marginal curve unless you specify the curve form.

Even with only a few observations, however, the tool can be useful. Don't hesitate to apply some judgment, though. For example, if the team showed no improvement on the second house, it is unlikely that they learned nothing from the first one. More likely, something abnormal happened to slow the work. Thus you might discount the model's prediction of no improvement on the third house.

<u>**Question #6A**</u>: **When starting production of a new circuit board, the Galvanic Board Company encounters relatively many rejected boards at first. The number declines as the operators gain experience and equipment problems are ironed out. This process follows a learning curve, and they expect the number of rejects to decline according to a learning curve.**

Of the first 100 boards,	**7 were bad.**
Of the second 100	**6 were bad.**
Of the third 100	**5 were bad.**
Of the fourth 100	**5 were bad.**
Of the fifth 100	**4 were bad.**
Of the sixth 100	**3 were bad.**

What rejects should they expect in the future at this rate?

Data Input columns:			Forecast:			
			Name of curve:		Marginal (Unit) Curve	
			Goodness-of-fit:		86.01%	
Unit						
Number	Time or $		Start Time		7.61	
1	7		Learning Curve %		74.98%	
2	6					
3	5		Batch			
4	5		Number	Rejects	² Changed titles!	
5	4		7	3.39		
6	3		8	3.21		
			20	2.19		
			50	1.50		
			100	1.12		

Solution to Question 6A

<u>Question #6B</u>: 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. What is the benchmark for measuring improvement?

This calls for the Advanced Forecasting option. But what should we use as a definition of the first batch? 100 boards with 7 bad, or 200 with 13 bad? 600 with 30 bad? Any of the definitions would be acceptable. Let's see what an 80% cumulative average curve looks like assuming 7 were bad in batch 1.

Data Input columns:			Forecast:			
			Name of curve:		Average Time Curve	
			Goodness-of-fit:			
Unit						
Number	Time or $		Start Time		7.00	
			Learning Curve %		80.00%	
			Unit			
			Number	Time or $		
			7	2.60		
			8	2.48		
			20	1.82		
			50	1.35		
			100	1.08		

This output is the result of the following steps:
1) Entering the unit numbers 7, 8, etc. to be forecasted.
2) Using the <u>Tools</u> ... <u>Adv Forecasting</u> option to enter Average Curve, Start time 7, and 80%.

The results indicate that our forecasted rejects are slightly higher than the industry benchmark.

Solution to 6B

<u>Question 7</u>: The Hard Rock Insurance Company has opened an office in Florida. Based on sales forecasts and past claims experience, they expect the following numbers of claims to arrive for processing:

Quarter	Year	Demand
3	1997	2000
4		2200
1	1998	2420
2		2662
3		2928
4		3221
1	1999	3543
2		3897
3		4287
4		4716

The claim-processing department was established early in 1997, and has required the following amounts of time to process each batch of 50 claims received so far:

Batch	Min.
1	60
2	50
3	45
4	42
5	39
6	35
7	32
8	33

Ultimately, this team can devote up to 800 hrs./quarter to processing claims. Currently, until demand increases, they also work on other projects. *Will the claims processing team be able to keep up with demand?*

The solution appears as follows:

Data Input columns:			**Forecast:**							
			Name of curve:		Marginal (Unit) Curve					
			Goodness-of-fit:		98.04%					
Unit										
Number	**Time or $**		**Start Time**		61.58					
1	60		**Learning Curve %**		80.96%					
2	50									
3	45		**Unit**							
4	42		**Number**	**Time or $**	Hours to produce demand		Quarter	Year	Demand	Cum. Prod'n
5	39		48	18.93	757.2798		3	1997	2000	2400
6	35		92	15.67	689.6832		4		2200	4600
7	32		140.4	13.92	673.8242		1	1998	2420	7020
8	33		193.64	12.76	679.2029		2		2662	9682
			252.204	11.90	696.8699		3		2928	12610
			316.6244	11.23	723.2255		4		3221	15831
			387.4868	10.68	756.6063		1	1999	3543	19374
			465.4355	10.21	796.2172		2		3897	23272
			551.1791	9.82	841.7138		3		4287	27559
			645.497	9.47	893.0161		4		4716	32275

Sending Data to the Author for Research Purposes

1. **Check the "Enable Notification" block.**

2. **Fill out your identifying information and description of where the data came from.**

3. **Give the file a name and save it.**

4. **Attach it to an e-mail message and send it to Charles.Bailey@bus.ucf.edu or mail a diskette to me.**

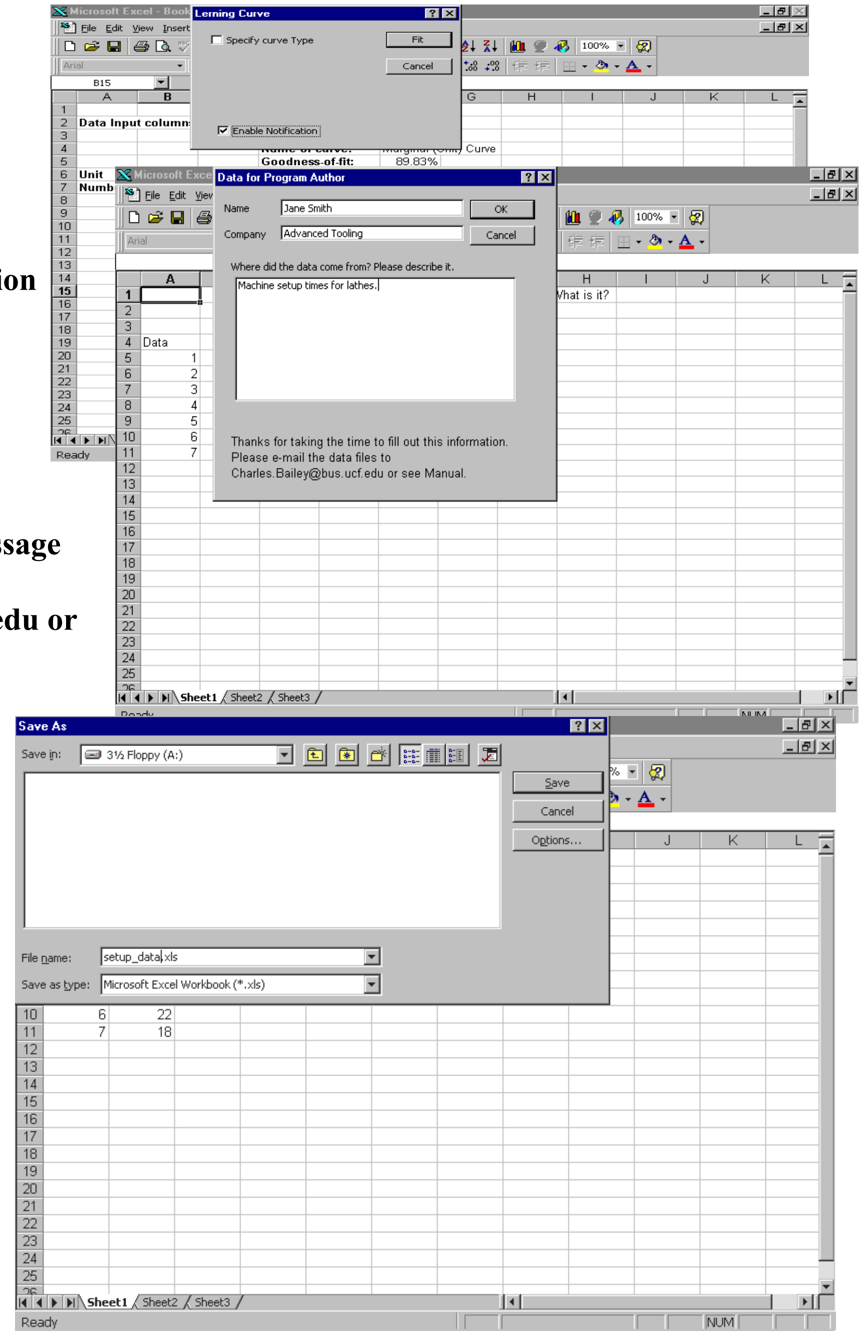

Please feel free to contact me at any time. I will certainly keep any data you send me confidential. ***If you have any concerns about the confidentiality of data, we can work out a way to disguise it.***

Thanks for your help and cooperation in return for this software.

Technical Assistance

I will be pleased to help you if this manual is unclear, if you find anything mysterious about the software, or if you need help formulating and solving a problem. Please also give me the benefit of any editorial comments, suggestions, or “bugs” you identify. Feedback that the software has been useful to you would be especially heartening. You may reach me at

Charles.Bailey@bus.ucf.edu (preferred mode of communication, checked regularly)
(407) 823-0151 (with voice mail)
(407) 823-3112 (FAX)

Best Wishes!

Appendix

The learning curve relationship can be modeled by an exponential formula. We may choose to fit either marginal or cumulative average data to this equation.

Average Model

If we let y be cumulative average time (or cost), then

$$y=ax^b$$

where a = the estimated time for the first unit (or batch)
x = cumulative units (batches) made
b = the slope parameter— i.e., log(LC rate)/log(2).
For an 80% curve, log(.8)/log(2) = !0.3219

Cumulative total time under this model after x units is $xy = ax^{b+1}$

Unit (Marginal) Model

If, on the other hand, we define y as the marginal time, we can use the same exponential equation,

$$y=ax^b$$

to predict y. However, the total time formula mentioned above no longer would apply. Now, total time would be the summation of all marginal times (although there are formulas to approximate this number more easily).

Distinction between Unit and Average Models

The distinction between two forms of learning curves—cumulative average and unit (incremental, marginal)—is in the choice of "what to model." We could fit a similar curve *form* (such as the popular $\mathbf{y=ax^b}$) to either average times (or dollars) or to unit times. From the model's basic estimate, one can estimate the other objects of interest, as summarized in the following table:

Object of Estimation	Average Form: Begins by Modeling Average Time	Unit Form: Begins by Model-ing Unit Time
Average Time [or $]	$\mathbf{y=ax^b}$	TT/x
Marginal (unit) Time	$TT_x - TT_{x-1}$	$\mathbf{y=ax^b}$
Total Time (TT)	$x(ax^b)=ax^{b+1}$	3y

To reiterate, we can define y as either cumulative average time or as marginal time. This choice of data will determine the a and b parameter estimates. This is why the same formula can apply to either the cumulative average or the marginal curves: the parameters will differ.

Using either curve form, one can fit the curve either to *average* data or to *marginal (unit)* data. By forecasting either the average or the unit data, one can easily convert to unit times, totals, etc. As explained below, 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, it *is* hard to grasp. I have yet to find a textbook that distinguishes clearly between the curve forms, and academic articles have addressed the topic.

Bailey-McIntyre Model

In addition to the log-linear models described above, this software offers the Bailey-McIntyre Model,

$$\log M = a[\log(x+1)]^b .$$

Bailey and McIntyre (1992, 1997) found that this model fit their assembly-task data better than the traditional model, and that it predicted as well or better.

The only difference from the discussion above is that the Bailey-McIntyre form of curve is used to predict either the average or the unit time, instead of the log-linear curve.

Measure of R^2

All curves were made linear by a logarithmic transformation, and ordinary least-squares regression was used to estimate parameter values. Because the dependent variables are different for each of the three types of curves, the R^2 values from the regressions are not comparable. To make the R^2 values comparable we used each equation to estimate marginal times for each subject and used these estimates and actual marginal times to compute R^2 values, as recommended by Kvålseth (1985).

References

Bailey, C. D., and Cocco, A., "The Effects of Compensation Method on the Parameters of Industrial Learning Curves." Working paper, University of Central Florida and University of Nevada, Las Vegas, 1997.

Bailey, C. D., and Sanjay Gupta, "Judgment in Learning-Curve Forecasting: A Laboratory Study." Working paper, University of Central Florida, 1997.

Bailey, C. D., and McIntyre, E. V., "Some Evidence on the Nature of Relearning Curves ," *The Accounting Review,* **67** (1992), 368-378.

Bailey, C. D., and McIntyre, E. V., "The Relation Between Fit and Prediction for Alternative Forms of Learning Curves and Relearning Curves ," *IIE Transactions* (forthcoming 1997).

Draper, N. R., and H. Smith, *Applied Regresssion Analysis.* New York: John Wiley & Sons, 1981.

Kvålseth, T.O. (1985) Cautionary note about r^2. *The American Statistician*, 39(4), 279-285.

Smith, Jason, *Learning Curve for Cost Control*. Norcross, Georgia: Industrial Engineering and Management Press, 1989.

Teplitz, C. J., *The Learning Curve Deskbook*. New York: Quorum Books, 1991.

Yelle, L. E. 1979. "The Learning Curve: Historical Review and Comprehensive Survey." *Decision Sciences* 10: 302-328.