

Learning about the Learning Curve

Among the pleasures of the “researcher lifestyle” is the joy of always learning something new. If the shock of the new is the jolt that gets you going, then you probably have the mindset of a researcher, or maybe an artist. Art and science share some aspects as lifestyles, but there are essential differences, too. Paul Gaugin commented that “art is either plagiarism or revolution,” but science certainly does not have to be plagiarism if it is not revolutionary: in fact, it had better *not*. Researchers almost always work in the context of what has been discovered before.

Starting out on a new direction of research, we learn first from prior work, and then we start to learn from nature itself as our own theories and experiments take shape. We all talk glibly about the “learning curve” and some managers even use it as a decision-making tool in identifying fruitful areas to pursue. Since the readers of *MRS Bulletin* live on the learning curve, I thought it would be a good idea to learn some more about its background.

For most of us, the term “Learning Curve” conjures up a shape something like the Johnson-Mehl-Avrami-Kolmogorov curve that you might have learned in a class in phase transformations. In case you missed that class, it is illustrated in Figure 1. If this is a learning curve, the horizontal axis represents the amount of time or effort expended in learning some particular subject matter, and the vertical axis represents the amount that is known. Of course, there is no way to measure the expended effort or the extent of mastery of any new topic so this is not really a quantitative matter, but we all have a pretty good grasp of the issues involved on some intuitive level. The curve can apply to the knowledge of an individual, or the collective knowledge of a group, or all of humankind.

This kind of learning curve leads to a particular strategy for prioritizing research projects: If you think that the world is on one of the two plateau regions of the curve, then you have little to gain by expending any more effort than your competitors, and no risk from expending less. On the other hand, if you think that knowledge is accumulating quickly, on the rising part of the curve, then it pays to make special efforts, because a little extra work should result in a significant gain of mastery, and you have the opportunity to get ahead of the competition—or fall behind. The hard part is knowing which part of the curve you are on at any particular moment. If you are doing research, you do not know the answers: If you did,

then whatever you were doing would not be research. You cannot determine where you are on the learning curve, except from the vantage point of the upper plateau. The people who do well at this game are those who can figure out where they are on the learning curve before they get to the top.

Most of us think we have some idea of where we are, but it can be distressingly unreliable. Figure 2 is a schematic learning

curve where actual mastery is replaced with *self-perceived* mastery on the vertical axis: Most people starting out in a new subject area think they have some understanding, so the curve does not start at zero. If, as you learn more, you come to recognize the true depth of your ignorance then you may be ready to master the subject after all. You may have some experience of dealing with individuals who have particularly high starting plateaus; if not, then you have not had teenage children, an idiot boss, or paid attention to much of the media, and you should consider yourself lucky. The higher the starting plateau, the less likely it is that the curve will follow the simple shape shown here. Things can get pretty weird for some of those cases.

There is one more kind of learning curve, sometimes distinguished from those described so far, by calling it an “experience curve.” This one is quantitative and often very reliably predictive, though it applies more to manufacturing than to research. It was first noted by researchers at Wright Patterson Air Force Base in 1936, where it was applied to aircraft manufacture, but it also applies to ships, razor blades, and solar cells, and probably all sorts of other stuff too. Moore’s Law is a variant of it, and materials research nearly always contributes to it. The basic observation is that the cost of making any item declines as you make more and more of them, and the decline is described by a simple power law. This is called Henderson’s Law and it is best illustrated on a log-log plot of cost versus accumulated experience, as shown in Figure 3. The steeper the negative slope, the faster the cost declines, and if the cost is reduced by 10% for every doubling of production, the product is said to have a “90% experience curve.” For any particular product, the experience curve is remarkably constant over long periods of time, unless truly disruptive technologies emerge. Solar cell production costs have been riding an 80% experience curve since about the mid-1970s.

OK, it is time for a Pop Quiz:

1. Where are you on the learning curve about learning curves?
2. Where were you when you started reading this article?
3. Is your mastery perceived or actual?
4. Where do you think the world is on the learning curve for climate change?

ALEX KING




Figure 1.




Figure 2.




Figure 3.