What Shape Does Progress Take?

Don’t assume it’s a straight line.

By Lee Anne Fennell

I*s this worth doing?* The question arises in every domain of life, at every scale, from the smallest and most personal of decisions to the largest and most public. For assessing what—and how much—is worth doing, one useful conceptual tool is a *production function*. It maps the relationship between units of input (like money, time, or effort) and outputs (whatever you are trying to achieve, from social change to completing a research paper).

People often assume, without thinking about it much, that the relationship between inputs and outcomes will be linear, like figure 1(a), where the output rises by the same incremental amount for each unit of input. If this were true, it would provide clear guidance about what is worth doing. You could make a few inputs, study the results, and then extrapolate outward to predict the full pattern, as in figure 1(b). And if you were getting a flat line, as in figure 1(c), you could just call it a day and move on.

FIGURE 1, LINEAR PRODUCTION FUNCTIONS

(a) (b) (c)

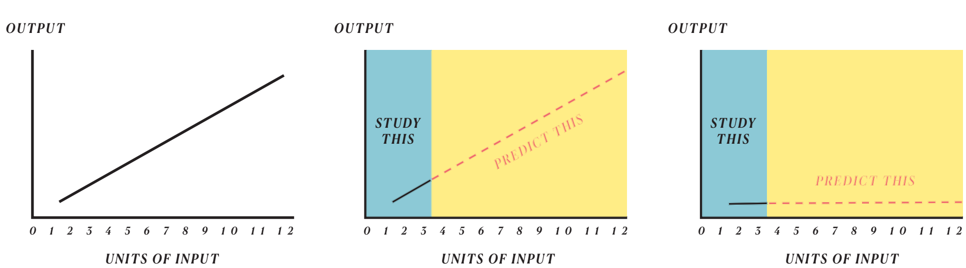

BS2-BrainMeetsWorld-Interior-toPress.indd 36 3/4/22 2:55 PM

However, production functions are commonly *nonlinear*.

Think of a bridge. If the inputs are bridge segments, we get no output at all—at least not in the domain of "bridge usefulness"—until we have put together enough segments to span the full chasm or river or alligator pit. And continuing to add extra bridge segments after the span is complete does no further good. The production function looks like figure 2(a), a step function.

The value of our bridge-in-progress remains flat as the first nine segments are added, and then jumps up all at once, when the tenth and final segment is put into place. If we were to assess the potential of our bridge-to-be based on the returns we get from the first few segments, we would get a misleading answer, as in figure 2(b).

FIGURE 2, STEP PRODUCTION FUNCTIONS

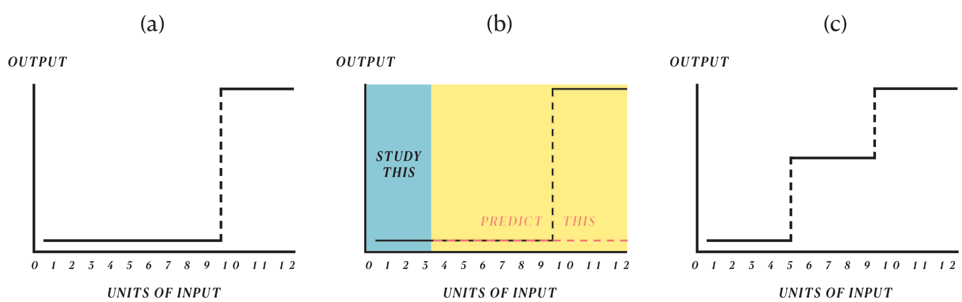

Of course, everyone knows how bridges work, so we would never make this mistake with an actual bridge. But how often do we make similar mistakes in other contexts—in work, in life, in public policy—by expecting linear results from what is really a nonlinear production process?

To be sure, a single-step production function is an extreme example of nonlinearity, one that captures things that operate in an all-or-nothing fashion: a machine that needs all its parts to function, a pass-fail test, or an election result. Sometimes there is more than one discrete step, as in figure 2(c), like making the cutoff for different grades or teams, or achieving milestones like job promotions. Many other production functions don't have sharply defined steps but do have areas of dramatically increasing or decreasing returns, like the examples in figure 3.

Figure 3(a) is an S-curve that might describe phenomena like social movements, learning curves, or the gains from urban clustering. Suppose you are organizing a rally, learning a new language, or trying to develop a downtown arts district in a city that lacks one. Things go slowly at first,

FIGURE 3, MORE PRODUCTION FUNCTIONS

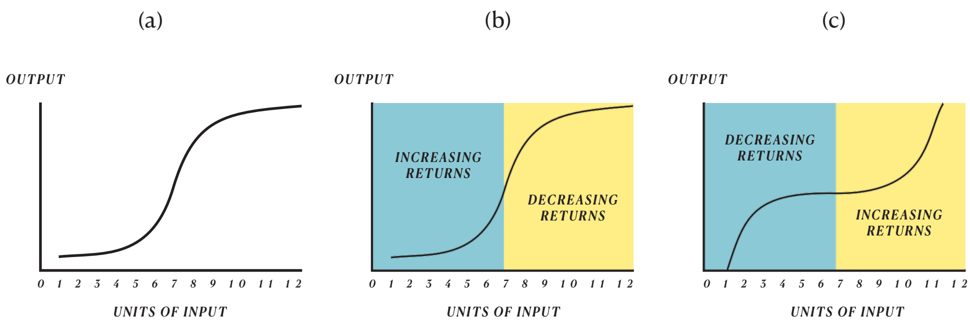

and the returns seem meager. The early ralliers, the simple phrases painfully strung together, the first small gallery with limited hours and few visitors, may all have a discouraging drop-in-the-bucket quality. But as you keep going, if you keep going, additional inputs start to bring increasing returns, and then the gains snowball for a time. Eventually, you may build a robust core of participants, a decent vocabulary, or a thriving arts district. At some later point, as success continues, things level off, and each input brings smaller and smaller marginal returns.

In fact, figure 3(a)'s S-curve combines two different curves, as shown in figure 3(b): an early range of increasing returns and a later range of decreasing returns. Some situations might involve *only* increasing returns or *only* decreasing returns, at least within the range that is relevant to our experience on the ground. Or the regions of increasing and decreasing returns might be inverted as in figure 3(c), where the first units of an input bring the largest gains, then level off, and then begin to gather steam again, taking things to a whole new level.

Many shapes are possible, all of which underscore the error of assuming linearity. Researchers have flagged the hazards of what Jordan Ellenberg calls "false linearity" in interpreting statistical trends and making predictions based upon them. Analogous cautions apply to our own projects, as we determine where to allocate effort and money, when to keep going, and when to give up. Different production functions call for different strategies.

But how can we tell whether we are dealing with an S-curve, a step, or just a true flat line? Sometimes we can't, though careful analysis and past experience can often shed light on what combinations of inputs are most likely to produce returns, and whether those results are likely to appear little by little or in large lumps of value. We can ask ourselves some questions: Are there fixed costs that need to be covered up front to get things going, or a certain critical mass of participants who have to be on board before things take off? Is the first stab at this type of problem usually the most significant, with later follow-ups helping significantly less? Are there good reasons to think that a policy might produce different results in the long run from those that are visible in the short run?

Consider the Moving to Opportunity studies, which investigated the effects of housing mobility and neighborhood environment on families who had been living in public housing. The experiment, which ran in five large U.S. cities in the 1990s, randomly assigned some families to receive experimental housing vouchers that, unlike ordinary Section 8 vouchers, could only be used in low-poverty neighborhoods. The observed effects of this intervention were initially underwhelming. Although the intervention *did* deliver some important benefits (like improved subjective well-being, family safety, and better mental and physical health), the hoped-for gains in earnings and employment rates didn't seem to materialize. But in 2016, Raj Chetty and his coauthors published a new study that traced the long-term effects on children who were younger than age 13 when their families received vouchers to move to low-poverty neighborhoods. These children, unlike their older counterparts, were more likely to attend college and earned more in adulthood than those whose families received regular Section 8 vouchers or no voucher. Longer exposures to low-poverty neighborhoods, starting earlier in childhood, produced results that exposures starting later in life did not. Extrapolating from the initial data would have been as misleading as judging a bridge's potential on the performance of the first few segments. Even when we remain uncertain about the precise shape that a production function takes, keeping the possibility of nonlinearity in mind can help keep us patient and humble.

Nonlinear production functions can be both exciting and daunting. The possibility of a bridge-completing breakthrough can be highly motivating, but the prospect of never being able to put together all of the necessary elements to produce one can keep some would-be bridge-builders from ever starting at all. Often, we must send our bridge segments out into the world without knowing yet how, or whether, they will fit together with other inputs—our own or those of others—to create something of value. At the same time, we must always be on the lookout for opportunities to connect our own contributions to the bridge segments built by others.

These lessons of nonlinearity can also inform our everyday lives and shed light on otherwise puzzling aspects of human behavior. Continuing to pour money or effort into an enterprise that is not generating any apparent returns might look like a sunk cost fallacy from the outside, but may feel like bridge building from the inside. Conversely, pursuing the predictably productive day by day might mean missing out on a larger gain that can only come from stepping away from the treadmill long enough to engineer a larger leap.

Before embarking on any undertaking—and especially before giving up on it—ask yourself what the production function looks like. And when assessing what is worth doing or supporting, or when gauging what counts as success, stay attuned to the prospects of *nonlinear* production in your own life and those of others. ■

Lee Anne Fennell is the Max Pam Professor of Law at the University of Chicago Law School. She is the author of *The Unbounded Home* and *Slices and Lumps: Division and Aggregation in Law and Life*.

For Further Reading

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz (2016). "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106: 855 - 902.

Ellenberg, Jordan (2014). *How Not To Be Wrong: The Power of Mathematical Thinking*. Penguin Press.

--

Fennell, Lee Anne (2019). *Slices and Lumps: Division and Aggregation in Law and Life*. University of Chicago Press.

--

Oliver, Pamela, Gerald Marwell, and Ruy Teixeira (1985). "A Theory of the Critical Mass. I. Interdependence, Group Heterogeneity, and the Production of Collective Action." *American Journal of Sociology* 91: 522 - 56.