BY CHRISTINE M. ANDERSON-COOK AND LU LU

Divide and Conquer

Understanding trade-offs by studying separate relationships

IN HIS BEST-SELLING book, *David and Goliath: Underdogs, Misfits, and the Art of Battling Giants*,[1] Malcolm Gladwell presents examples in which benefits for many desirable attributes don't follow the assumed ever-increasing pattern, but rather can be described by what he calls an "inverted U-shaped curve."

He describes, for example, how many of us might assume that having more money would make parenting easier. We might think that having $200,000, instead of $50,000 or $100,000, would take some pressure off of child rearing. To dispel this idea, he describes how having too much money can lead to children lacking motivation to work for anything and feeling entitled.

Hence, a plot of income (the x-axis) versus ease of child rearing (the y-axis) has a bell (or inverted U) shape with a peak at about $75,000, according to sources quoted by Gladwell. Through examples—including how school class size relates to the effectiveness of teaching, and how alcohol consumption affects people's health—Gladwell describes a common principle that applies to many things in life: Having too much of a good thing can actually be less desirable.

Competitive situations

These examples speak to the primary message of Gladwell's book of how we tend to exaggerate advantages and disadvantages when looking at competitive situations.

The idea is an important one, and it can lead to improved thinking about daily choices we make. An important question, however, is how we can find the best solution to achieve the optimal gain—that is, the peak of the inverted U-shaped curve. This requires an accurate and precise quantification of the true relationship of the outcome and the input factor—that is, the inverted U-shaped curve.

There is, perhaps, one more step to take to help further unravel the complexity of finding the optimum. At the heart of all of Gladwell's examples, and many others that we can think of for ourselves, is the idea of competing relationships between the driving factor and the result for different objectives.

In the child-rearing example, more money always will allow parents to offer their children more opportunities. There is no natural maximum on the opportunity scale, but perhaps diminishing returns. What pulls the inverted U-shaped curve down after the peak, however, is a different relationship. We might call it the "entitlement scale," which suggests that children lose motivation to achieve things or strive for excellence on their own.

More relationships

With the example of classroom sizes, the two competing relationships might be labeled "attention from teacher" (in which smaller class size is more beneficial for attention) and "diversity of opinion and group learning" (in which opportunities to hear from differing views and to develop cooperation skills) are lost with smaller classroom size. When both aspects play important roles in influencing teaching,

Choosing the number of production lines / FIGURE 1

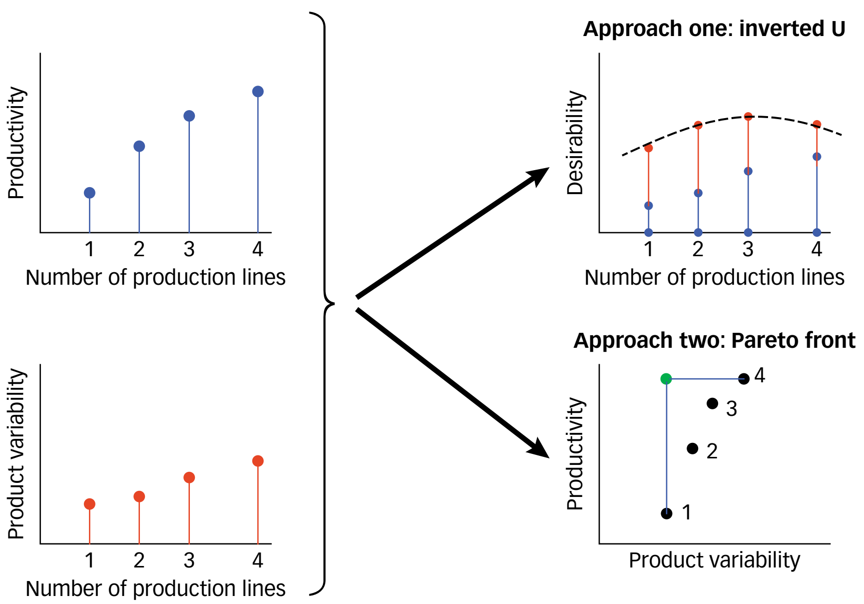

These are two approaches for choosing the best number of production lines for maximizing productivity and minimizing product variability. Approach one uses a desirability function to combine the two objectives with the integrated relationship matching Gladwell's inverted U-shaped curve to summarize. Approach two identifies a Pareto front consisting of all the superior solutions and examines their trade-offs and robustness to subjective choices.

reducing the class size improves teaching overall, before the class size gets too small. After reaching the optimum, reducing the class size further could lead to diminishing effect. So the essence of the inverted U-shaped curve is not that any individual relationship has a natural peak, but rather that there are competing relationships with optimal values in different regions.

With alcohol consumption, consider the positive effects of additional intake to be the established heart, cholesterol and longevity benefits, as well as potential creativity boosts. Competing against that are the detrimental effects on a person's liver, blood pressure and weight. In addition, there is potential impairment to good cognitive function and, potentially, addiction. Hence, some alcohol consumption is generally considered beneficial for many, but overconsumption is undesirable.

Consider a simple illustrative quality management scenario in which an organization has a small number of production lines. It has the opportunity to expand its production and add more capacity with more production lines. A simplistic version of the decision-making process would say that more production would always be better, but let's see how Gladwell's inverted-U shaped curve might manifest itself for this situation.

The managers consider two aspects when deciding whether and how much to expand—possibly from one line up to as many as four lines:

1. Productivity as measured by the amount of product that can be produced from having more production lines.
2. Consistency of product by producing items with smaller variability.

As shown in the left panel of Figure 1, increasing the number of production lines will increase the productivity (desirable), but also makes the products more variable (undesirable). Hence, the decision involves making a trade-off between the two competing relationships.

The top right plot in Figure 1 shows how to use a desirability function approach,[2] labeled as approach one: inverted U. First, because it is difficult to compare different objectives measured on different scales, the method converts each of the performance measures on individual objects and criteria into desirability values on a 0-1 scale. The desirability value 1 corresponds to best performance, and value 0 corresponds to worst performance.

For the production line example, the desirability on productivity increases as the number of production lines gets larger (corresponding to the increasing height of the blue lines), while the desirability on the consistency measure decreases with having more production lines (corresponding to the decreasing height of the red lines).

Note when the goal is to minimize an objective, smaller values on the original scale are mapped to larger values on the desirability scale. In the second step, the overall desirability is obtained by combining the two individual desirability values—stacking the blue and red lines. As a result, the overall benefit from different numbers of production lines can be assessed.

Diminishing returns

Notice that there are often diminishing returns as we get to the extremes of our choices—too many production lines (for productivity) or too few (for consistency). When we combine the measures into a single desirability number, moderate values from among our choices look quite inviting.

Hence, the resulting inverted U-shaped curve—if we look at the combined scores when combining the heights of the red and blue line segments for an overall desirability score—is obtained for capturing the effect of the number of production lines.

To use the desirability approach, we must make a few user-specific choices at the beginning of our decision-making process:

- Decide how to convert the performance measures on individual criteria into a desirability scale (what values get assigned 1 and 0) and map the values in between based on a linear transformation.
- Choose the metric to use for integrating multiple objectives into a single summary. Here, we have used an additive method to combine the contributions from individual measures to allow good performance on one criterion to overcome the disadvantage of having poor performance on another criterion.
- Select the appropriate relative importance of the objectives or the relative range of heights for the red and blue lines.

Among these three aspects, how we value the importance of the objectives is difficult to decide upon precisely, and different subject matter experts can have different preferences. These subjective choices can have different effects on the decisions and can vary on a case-by-case basis.

You can certainly imagine how making the range of lengths of the blue line segments larger or smaller (depending on how important it is to deliver a consistent product) would have an effect on where the peak of the inverted U-shaped curve is located.

Hence, it is important to recognize that the decision depends on the subjective choices made when we combine different characteristics on different scales into a single summary.

It is also important to understand the impacts of our choices. Consider again the child rearing and income example. How much we value opportunity versus the fear of entitlement can substantially alter where we perceive the optimal income level to be.

Pareto front approach

An alternative approach is to look at many promising solutions before narrowing them to a specific set of subjective choices to understand their potential impacts. The Pareto front approach[3,4] (labeled as Approach two: Pareto front in Figure 1) offers an objective way to find the subset of promising solutions. It considers the objectives directly without initially converting them to a common desirability scale.

The bottom right plot in Figure 1 with the four solutions (labeled for production lines one to four) shows possible solutions with productivity summary and consistency measure each displayed on their own axes.

For this example, the ideal solution would be to maximize production with minimal variability. This green circle (known as the Utopia point in the Pareto front literature) shows the best combinations of objectives, but it is rarely achievable in practice. It does provide some calibration for how to assess the possible solutions under consideration.

The Pareto front identifies superior solutions that are not outperformed by any other solutions based on only the observed values for all criteria under consideration.

For the production line example, suppose the four solutions are the only options we have. For each one, we cannot find another solution that is strictly better on both criteria. Hence, the Pareto front shown in the bottom right panel of Figure 1 consists of all four solutions, and each of them can be optimal for different subjective choices.

In many other applications, there are a lot more candidate solutions (some may not be possible to enumerate), and the Pareto front offers an efficient approach to reduce the set of promising solutions for further consideration.

After finding the Pareto front, the second stage uses a structured approach with a rich set of graphical tools[5-7] to evaluate different options and their trade-offs to make an informed decision. Specifically, all the solutions on the Pareto front can be ranked based on their proximity to the Utopia point for different weighting preferences between the two objectives.

Next, the best solutions for different weighting choices are identified, and their robustness and relative performance to optimum based on the user priorities are considered for making an informed decision.

Advantages of the Pareto front approach include:

1. It allows the decision-maker to eliminate noncontenders from further consideration. This point is not so important in with the production line example with limited options, but it's often a substantial simplification for cases with a large set of options or uncountable options.
2. It shows all the superior options before narrowing the consideration based on the subjective choices.
3. It offers a flexible and efficient structure for evaluating the potential impacts from the subjective choices.

Making better decisions

The two approaches covered here showcase different ways and philosophies to seek optimal solutions for multiple objectives. No matter which approach practitioners choose for their particular applications, it is important to start with untangling the overall pattern into the individual relationships that drive changes in the outcome of interest, quantifying these separate drivers and showing how they interact with one another.

It is certainly easier to understand child rearing, alcohol consumption, classroom size or production capability when the relationships associated with a positive result are separately articulated and quantified. Based on this understanding, a final decision can be made by strategically combining the competing relationships and optimizing the integrated relationship based on goals and priorities.

Gladwell's inverted U-shaped curves are broadly applicable for so many daily decisions. How to manage our own inverted U-shaped curves for making the best decisions requires us to separate and understand the relationships driving changes in the outcomes and appropriately balance the trade-offs between competing objectives.

There have been other examples in statistics literature in which the idea of studying objectives separately has led to dramatic improvements. Perhaps one of the best known is robust parameter design in which Genichi Taguchi's signal-to-noise ratio was disentangled to look at factor effects on the mean and variance individually.[8]

A clearer understanding of advantages and disadvantages can help us to more realistically quantify the balance of life and realize that more is not always better. The more you understand the underlying relationships, however, the better decisions you can make for achieving ideal results. QP

REFERENCES

1. Malcolm Gladwell, *David and Goliath: Underdogs, Misfits, and the Art of Battling Giants,* Little Publishing, 2013.
2. George Derringer and Ronald Suich, "Simultaneous Optimization of Several Response Variables," *Journal of Quality Technology,* Vol. 12, No. 4, October 1980, pp. 214-219.
3. Christine M. Anderson-Cook and Lu Lu, "Weighing Your Options," *Quality Progress*, October 2012, pp. 50-52.
4. Christine M. Anderson-Cook, "Let's Be Realistic," *Quality Progress,* March 2013, pp. 52-54.
5. Lu Lu, Christine M. Anderson-Cook and Timothy J. Robinson, "Optimization of Designed Experiments Based on Multiple Criteria Utilizing a Pareto Frontier," *Technometrics,* Vol. 53, No. 4, pp. 353-365.
6. Lu Lu and Christine M. Anderson-Cook, "Rethinking the Optimal Response Surface Design for a First-Order Model With Two-Factor Interactions, When Protecting Against Curvature," *Quality Engineering,* Vol. 24, No. 3, pp. 404-422.
7. Lu Lu, Jessica L. Chapman, Christine M. Anderson-Cook, "A Case Study on Selecting a Best Allocation of New Data for Improving Estimation Precision of System and Sub-System Reliability Using Pareto Fronts," *Technometrics,* Vol. 55, No. 4, pp. 473-487.
8. Vijayan N. Nair, ed., "Taguchi's Parameter Design: A Panel Discussion," *Technometrics*, Vol. 34, No. 2, pp. 127-161.

CHRISTINE M. ANDERSON-COOK is a research scientist in the statistical sciences group at Los Alamos National Laboratory in Los Alamos, NM. She earned a doctorate in statistics from the University of Waterloo in Ontario. Anderson-Cook is a fellow of both ASQ and the American Statistical Association.

LU LU is a visiting assistant professor in the department of mathematics and statistics at the University of South Florida. She was a postdoctoral research associate in the statistical sciences group at Los Alamos National Laboratory. She earned a doctorate in statistics from Iowa State University in Ames, IA.