

Linear Thinking in a Nonlinear World

By Bart de LangheStefano PuntoniRichard Larrick

FROM THE MAY–JUNE 2017 ISSUE

Test yourself with this word problem: Imagine you're responsible for your company's car fleet. You manage two models, an SUV that gets 10 miles to the gallon and a sedan that gets 20. The fleet has equal numbers of each, and all the cars travel 10,000 miles a year. You have enough capital to replace one model with more-fuel-efficient vehicles to lower operational costs and help meet sustainability goals.

Which upgrade is better?

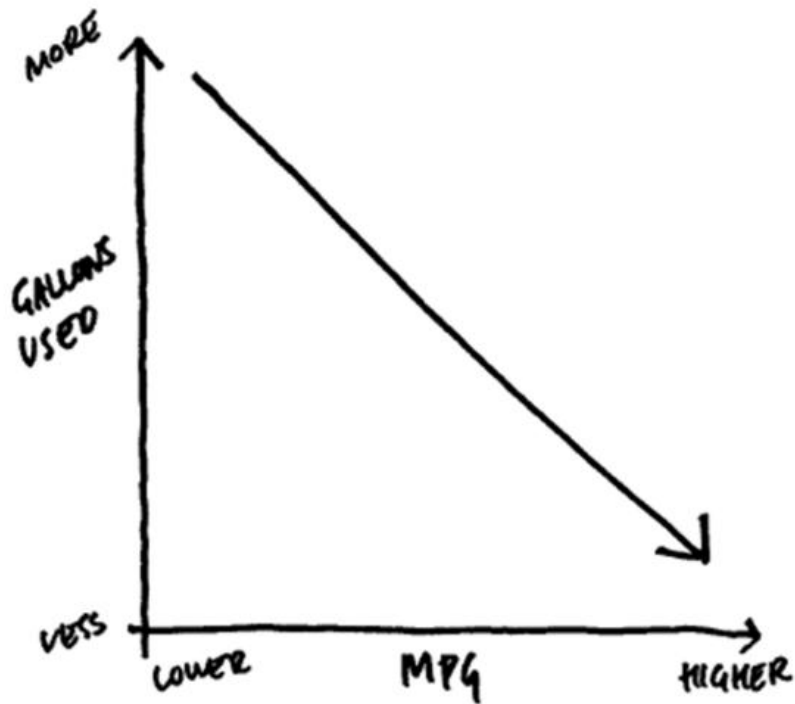
- A. Replacing the 10 MPG vehicles with 20 MPG vehicles
- B. Replacing the 20 MPG vehicles with 50 MPG vehicles

Intuitively, option B seems more impressive—an increase of 30 MPG is a lot larger than a 10 MPG one. And the percentage increase is greater, too. But B is not the better deal. In fact, it's not even close. Let's compare.

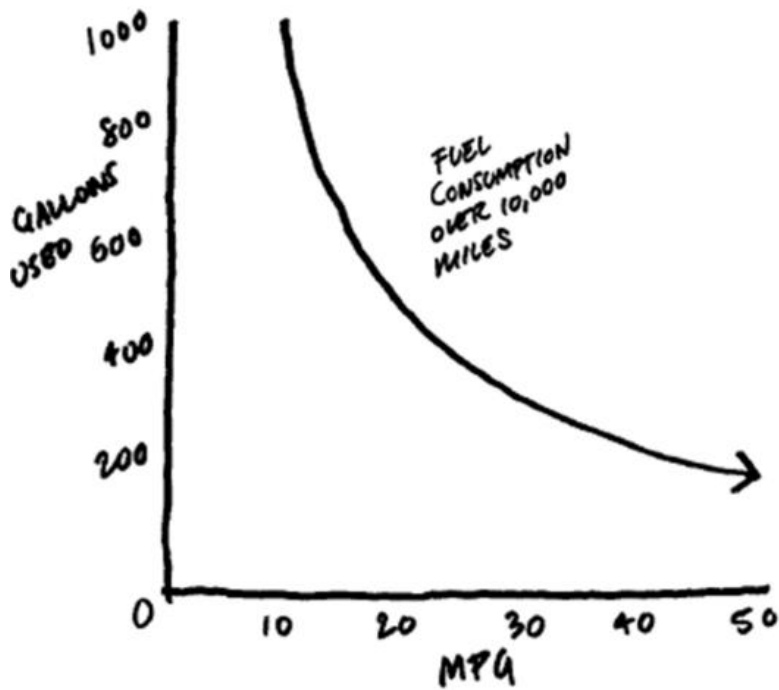
GALLONS USED PER 10,000 MILES

	CURRENT	AFTER UPGRADE	SAVINGS
A.	1,000 (@10 MPG)	500 (@20 MPG)	500
B.	500 (@20 MPG)	200 (@50 MPG)	300

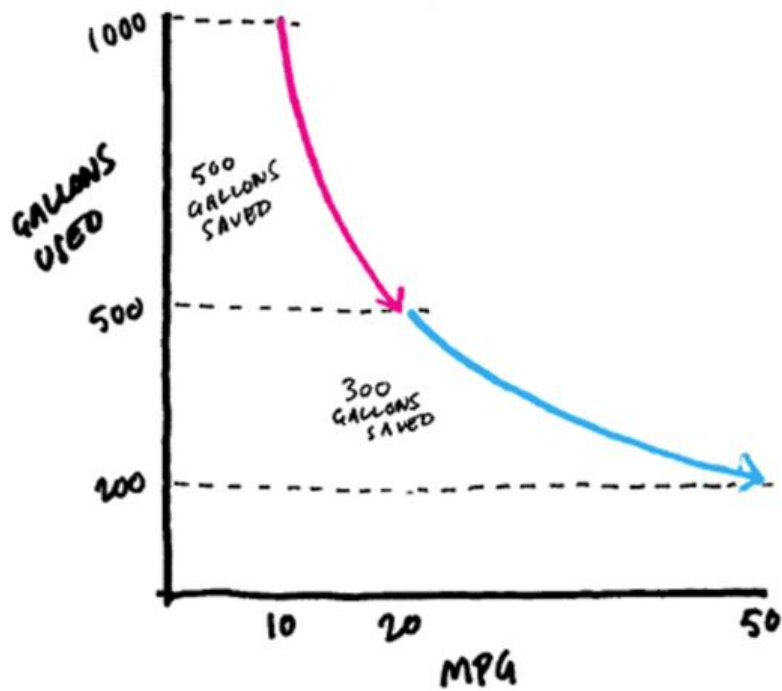
Is this surprising? For many of us, it is. That's because in our minds the relationship between MPG and fuel consumption is simpler than it really is. We tend to think it's linear and looks like this:



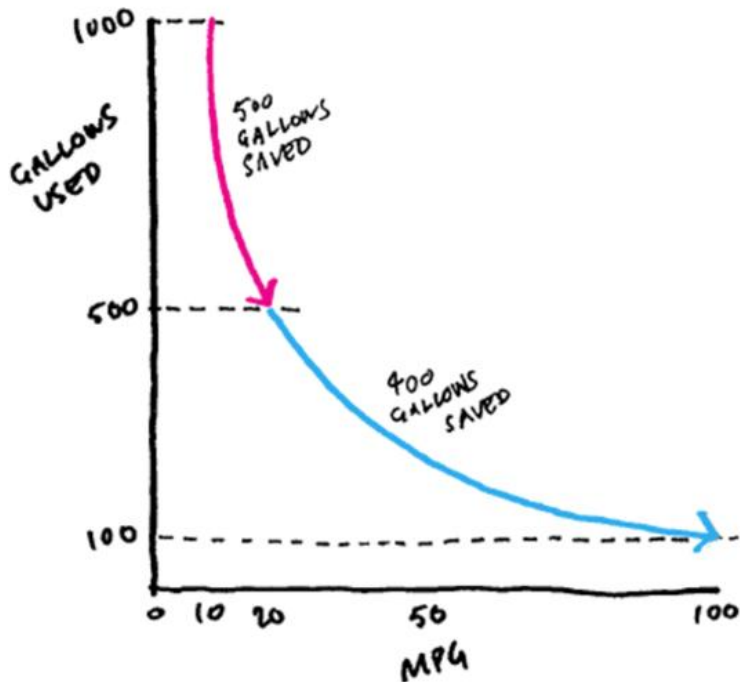
But that graph is incorrect. Gas consumption is not a linear function of MPG. When you do the math, the relationship actually looks like this:



And when you dissect the curve to show each upgrade scenario, it becomes clear how much more effective it is to replace the 10 MPG cars.



Shockingly, upgrading fuel efficiency from 20 to 100 MPG *still* wouldn't save as much gas as upgrading from 10 to 20 MPG.



But choosing the lower-mileage upgrade remains counterintuitive, even in the face of the visual evidence. It just doesn't feel right.

If you're still having trouble grasping this, it's not your fault. Decades of research in cognitive psychology show that the human mind struggles to understand nonlinear relationships. Our brain wants to make simple straight lines. In many situations, that kind of thinking serves us well: If you can store 50 books on a shelf, you can store 100 books if you add another shelf, and 150 books if you add yet another. Similarly, if the price of coffee is \$2, you can buy five coffees with \$10, 10 coffees with \$20, and 15 coffees with \$30.

But in business there are many highly nonlinear relationships, and we need to recognize when they're in play. This is true for generalists and specialists alike, because even experts who are aware of nonlinearity in their fields can fail to take it into account and default instead to relying on their gut. But when people do that, they often end up making poor decisions.

Linear Bias in Practice

We've seen consumers and companies fall victim to linear bias in numerous real-world scenarios. A common one concerns an important business objective: profits.

Three main factors affect profits: costs, volume, and price. A change in one often requires action on the others to maintain profits. For example, rising costs must be offset by an increase in either price or volume. And if you cut price, lower costs or higher volumes are needed to prevent profits from dipping.

Unfortunately, managers' intuitions about the relationships between these profit levers aren't always good. For years experts have advised companies that changes in price affect profits more than changes in volume or costs. Nevertheless, executives often focus too much on volume and costs instead of getting the price right.

Why? Because the large volume increases they see after reducing prices are very exciting. What people don't realize is just how large those increases need to be to maintain profits, especially when margins are low.

Imagine you manage a brand of paper towels. They sell for 50 cents a roll, and the marginal cost of producing a roll is 15 cents. You recently did two price promotions. Here's how they compare:


	NORMAL	PROMO A: 20% OFF	PROMO B: 40% OFF
PRICE/ROLL	50¢	40¢	30¢
SALES	1,000	1,200 (+20%)	1,800 (+80%)

Intuitively, B looks more impressive—an 80% increase in volume for a 40% decrease in price seems a lot more profitable than a 20% increase in volume for a 20% cut in price. But you may have guessed by now that B is not the most profitable strategy.

In fact, both promotions decrease profits, but B's negative impact is much bigger than A's. Here are the profits in each scenario:


	NORMAL	PROMO A: 20% OFF	PROMO B: 40% OFF
PRICE/ROLL	50¢	40¢	30¢
SALES	1,000	1,200 (+20%)	1,800 (+80%)
PROFIT/ROLL	35¢	25¢	15¢
PROFIT	\$350	\$300	\$270

Although promotion B nearly doubled sales, profits sank almost 25%. To maintain the usual \$350 profit during the 40%-off sale, you would have to sell more than 2,300 units, an increase of 133%. The curve looks like this:



The nonlinear phenomenon also extends to intangibles, like consumer attitudes. Take consumers and sustainability. We frequently hear executives complain that while people say they care about the environment, they are not willing to pay extra for ecofriendly products. Quantitative analyses bear this out. A survey by the National Geographic Society and GlobeScan finds that, across 18 countries, concerns about environmental problems have increased markedly over time, but consumer behavior has changed much more slowly. While nearly all consumers surveyed agree that food production and consumption should be more sustainable, few of them alter their habits to support that goal.

What's going on? It turns out that the relationship between what consumers say they care about and their actions is often highly nonlinear. But managers often believe that classic quantitative tools, like surveys using 1-to-5 scales of importance, will predict behavior in a linear fashion. In reality, research shows little or no behavioral difference between consumers who, on a five-point scale, give their environmental concern the lowest rating, 1, and consumers who rate it a 4. But the difference between 4s and 5s is huge. Behavior maps to attitudes on a curve, not a straight line.



Companies typically fail to account for this pattern—in part because they focus on averages. Averages mask nonlinearity and lead to prediction errors. For example, suppose a firm did a sustainability survey among two of its target segments. All consumers in one segment rate their concern about the environment a 4, while 50% of consumers in the other segment rate it a 3 and 50% rate it a 5. The average level of concern is the same for the two segments, but people in the second segment are overall much more likely to buy green products. That’s because a customer scoring 5 is much more likely to make environmental choices than a customer scoring 4, whereas a customer scoring 4 is not more likely to than a customer scoring 3.

The nonlinear relationship between attitudes and behavior shows up repeatedly in important domains, including consumers’ privacy concerns. [A large-scale survey](#) in the Netherlands, for example, revealed little difference in the number of loyalty-program cards carried by consumers who said they were quite concerned versus only weakly concerned about privacy. How is it possible that people said they were worried about privacy but then agreed to sign up for loyalty programs that require the disclosure of sensitive personal information? Again, because only people who say they are extremely concerned about privacy take significant steps to protect it, while most others, regardless of their concern rating, don’t adjust their behavior.

Awareness of nonlinear relationships is also important when choosing performance metrics. For instance, to assess the effectiveness of their inventory management, some firms track days of supply, or the number of days that products are held in inventory, while other firms track the number of times their

inventory turns over annually. Most managers don't know why their firm uses one metric and not the other. But the choice may have unintended consequences—for instance, on employee motivation. Assume a firm was able to reduce days of supply from 12 to six and that with additional research, it could further reduce days of supply to four. This is the same as saying that the inventory turn rate could increase from 30 times a year to 60 times a year and that it could be raised again to 90 times a year. But employees are much more motivated to achieve improvements if the firm tracks turnover instead of days of supply, [research by the University of Cologne's Tobias Stangl and Ulrich Thonemann](#) shows. That's because they appear to get decreasing returns on their efforts when they improve the days-of-supply metric—but constant returns when they improve the turnover metric.

Other areas where companies can choose different metrics include warehousing (picking time versus picking rate), production (production time versus production rate), and quality control (time between failures versus failure rate).

Nonlinearity is all around us. Let's now explore the forms it takes.

The Four Types of Nonlinear Relationships

The best way to understand nonlinear patterns is to see them. There are four types.

Increasing gradually, then rising more steeply.

Say a company has two customer segments that both have annual contribution margins of \$100. Segment A has a retention rate of 20% while segment B has one of 60%. Most managers believe that it makes little difference to the bottom line which segment's retention they increase. If anything, most people find doubling the weaker retention rate more appealing than increasing the stronger one by, say, a third.


Managers underestimate the benefits of small increases to high retention rates.

But customer lifetime value is a nonlinear function of retention rate, as you'll see when you apply the formula for calculating CLV:

$$\frac{\text{MARGIN} \times \text{RETENTION RATE}}{1 + \text{DISCOUNT RATE} - \text{RETENTION RATE}}$$

When the retention rate rises from 20% to 40%, CLV goes up about \$35 (assuming a discount rate of 10% to adjust future profits to their current worth),


but when retention rates rise from 60% to 80%, CLV goes up about \$147. As retention rates rise, customer lifetime value increases gradually at first and then suddenly shoots up.



Most companies focus on identifying customers who are most likely to defect and then target them with marketing programs. However, it's usually more profitable to focus on customers who are more likely to stay. Linear thinking leads managers to underestimate the benefits of small increases to high retention rates.

Decreasing gradually, then dropping quickly.

The classic example of this can be seen in mortgages. Property owners are often surprised by how slowly they chip away at their debt during the early years of their loan terms. But in a mortgage with a fixed interest rate and fixed term, less of each payment goes toward the principal at the beginning. The principal doesn't decrease linearly. On a 30-year \$165,000 loan at a 4.5% interest rate, the balance decreases by only about \$15,000 over the first five years. By year 25 the balance will have dropped below \$45,000. So the owner will pay off less than 10% of the principal in the first 16% of the loan's term but more than a quarter of it in the last 16%.



Because they're misled by their linear thinking in this context, mortgage payers are often surprised when they sell a property after a few years (and pay brokerage costs) and have only small net gains to show for it.

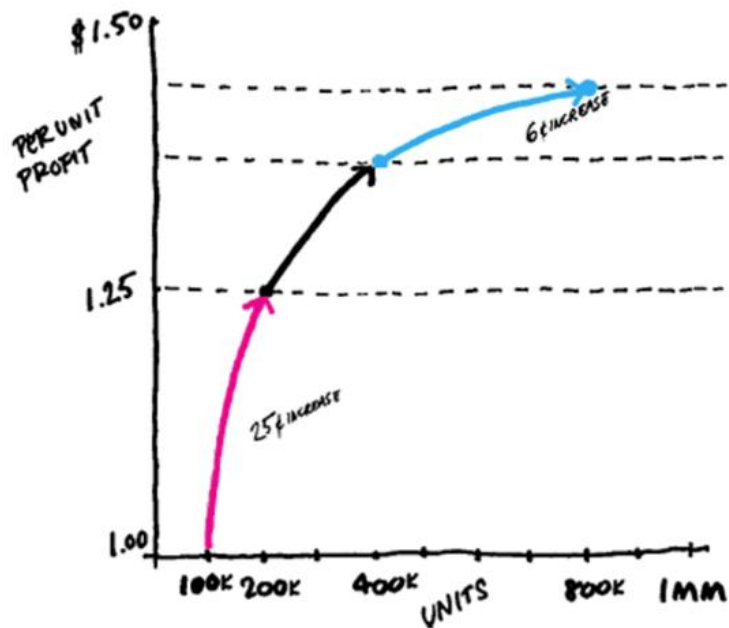
Climbing quickly, then tapering off.

Selling more of a product allows companies to achieve economies of scale and boost per unit profit, a metric often used to gauge a firm's efficiency. Executives use this formula to calculate per unit profit:

$$\frac{(\text{VOLUME} \cdot \text{UNIT PRICE}) - \text{FIXED COSTS} - (\text{VOLUME} \cdot \text{UNIT VARIABLE COSTS})}{\text{VOLUME}}$$

Say a firm sells 100,000 widgets each year at \$2 a widget, and producing those widgets costs \$100,000—\$50,000 in fixed costs and 50 cents in unit variable costs. The per unit profit is \$1. The firm can increase per unit profit by producing and selling more widgets, because it will spread fixed costs over more units. If it doubles the number of widgets sold to 200,000, profit per unit will rise to \$1.25 (assuming that unit variable costs remain the same). That attractive increase might tempt you into thinking per unit profit will skyrocket if you increase sales from 100,000 to 800,000 units. Not so.

If the firm doubles widget sales from 400,000 to 800,000 (which is much harder to do than going from 100,000 to 200,000), the per unit profit increases only by about 6 cents.



Managers focus a great deal on the benefits of economies of scale and growth. However, linear thinking may lead them to overestimate volume as a driver of profit and thus underestimate other more impactful drivers, like price.

Falling sharply, then gradually.

Firms often base evaluations of investments on the payback period, the amount of time required to recover the costs. Obviously, shorter paybacks are more favorable. Say you have two projects slated for funding. Project A has a payback period of two years, and project B has one of four years. Both teams believe they can cut their payback period in half. Many managers may find B more attractive because they'll save two years, double the time they'll save with A.

Company leadership, however, may ultimately care more about return on investment than time to breakeven. A one-year payback has an annual rate of return (ARR) of 100%. A two-year payback yields one of 50%—a 50-point difference. A four-year payback yields one of 25%—a 25-point difference. So as the payback period increases, ARR drops steeply at first and then more slowly. If your focus is achieving a higher ARR, halving the payback period of project A is a better choice.

Managers comparing portfolios of similar-sized projects may also be surprised to learn that the return on investment is higher on one containing a project with a one-year payback and another with a four-year payback than on a portfolio containing two projects expected to pay back in two years. They should be careful not to underestimate the effect that decreases in relatively short payback periods will have on ARR.

How to Limit the Pitfalls of Linear Bias

As long as people are employed as managers, biases that are hardwired into the human brain will affect the quality of business decisions. Nevertheless, it is possible to minimize the pitfalls of linear thinking.

Step 1: Increase awareness of linear bias.

MBA programs should explicitly warn future managers about this phenomenon and teach them ways to deal with it. Companies can also undertake initiatives to educate employees by, for instance, presenting them with puzzles that involve nonlinear relationships. In our experience, people find such exercises engaging and eye-opening.

Test Your Nonlinear Aptitude

READ MORE

Answers to Questions

READ MORE

Broader educational efforts are already under way in several fields. One is [Ocean Tipping Points](#), an initiative that aims to make people more sensitive to nonlinear relationships in marine ecosystems. Scientists and managers often assume that the relationship between a stressor (such as fishing) and an [ecological response \(a decline in fish population\)](#) is linear. However, a small change in a stressor sometimes does disproportionately large damage: A fish stock can collapse following a small increase in fishing. The project's goal is to identify relevant tipping points in ocean ecology to help improve the management of natural resources.

Step 2: Focus on outcomes, not indicators.

One of senior management's most important tasks is to set the organization's direction and incentives. But frequently, desired outcomes are far removed from everyday business decisions, so firms identify relevant intermediate metrics and create incentives to maximize them. To lift sales, for instance, many companies try to improve their websites' positioning in organic search results.

The problem is, these intermediate metrics can become the end rather than the means, a phenomenon academics call "[medium maximization](#)." That bodes trouble if a metric and the outcome don't have a linear relationship—as is the case with organic search position and sales. When a search rank drops, sales decrease quickly at first and then more gradually: The impact on sales is much greater when a site drops from the first to the second position in search results than when it drops from the 20th to the 25th position.

Other times, a single indicator can be used to predict multiple outcomes, and that may confuse people and lead them astray. Take annual rates of return, which a manager who wants to maximize the future value of an investment may consider. If you map the relationship between investment products' ARR and their total accumulated returns, you'll see that as ARR rises, total returns increase gradually and then suddenly shoot up.

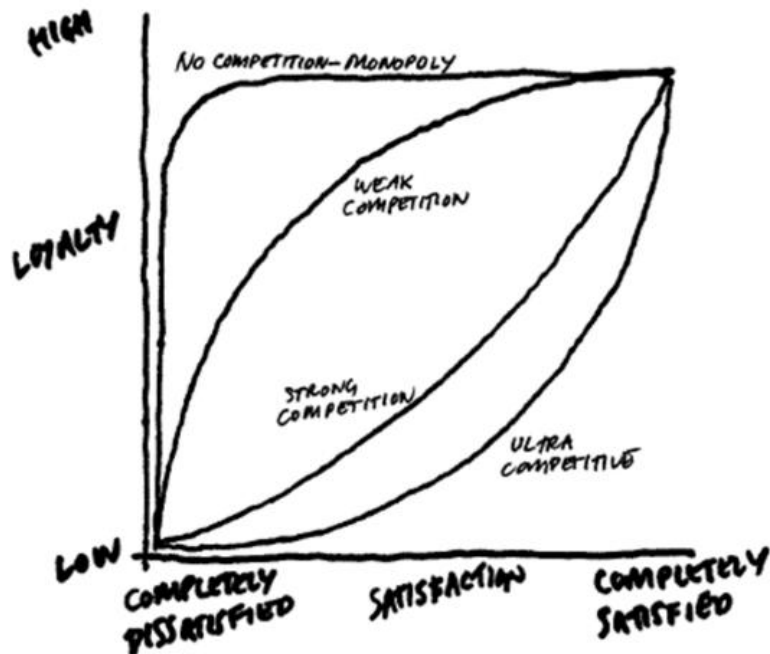
Another manager may wish to minimize the time it takes to achieve a particular investment goal. The relationship here is the reverse: As ARR rises, the time it takes to reach a goal drops steeply at first and then declines gradually.

Because ARR is related to multiple outcomes in different nonlinear ways, people often under- or overestimate its effect. A manager who wants to maximize overall returns may care a great deal about a change in the rate from 0.30% to 0.70% but be insensitive to a change from 6.4% to 6.6%. In fact, increasing a low return rate has a much smaller effect on accumulated future returns than increasing a high rate does. In contrast, a manager focused on minimizing the time it takes to reach an investment goal may decide to take on additional risk to increase returns from 6.3% to 6.7% but be insensitive to a change from 0.40% to 0.60%. In this case the effect of increasing a high interest rate on time to completing a savings goal is much smaller than the effect of increasing a low interest rate.

Step 3: Discover the type of nonlinearity you're dealing with.

As Thomas Jones and W. Earl Sasser Jr. pointed out in HBR back in 1995 (see "[Why Satisfied Customers Defect](#)"), the relationship between customer satisfaction ratings and customer retention is often nonlinear—but in ways that vary according to the industry. In highly competitive industries, such as automobiles, retention rises gradually and then climbs up steeply as satisfaction


ratings increase. For noncompetitive industries retention shoots up quickly and then levels off.



In both situations linear thinking will lead to errors. If the industry is competitive, managers will overestimate the benefit of increasing the satisfaction of completely dissatisfied customers. If the industry is not competitive, managers will overestimate the benefit of increasing the satisfaction of already satisfied customers.

The point is that managers should avoid making generalizations about nonlinear relationships across contexts and work to understand the cause and effect in their specific situation.

Field experiments are an increasingly popular way to do this. (See Eric T. Anderson and Duncan Simester's 2011 HBR article "[A Step-by-Step Guide to Smart Business Experiments.](#)") When designing them, managers should be sure to account for nonlinearity. For instance, many people try to measure the impact of price on sales by offering a product at a low price (condition A in the chart below) and a high price (condition B) and then measuring differences in sales. But testing two prices won't reveal nonlinear relationships. You need to use at least three price levels—low, medium (condition C), and high—to get a sense of them.



Step 4: Map nonlinearity whenever you can.

In addition to providing the right training, companies can build support systems that warn managers when they might be making bad decisions because of the inclination to think linearly.


Ideally, algorithms and artificial intelligence could identify situations in which that bias is likely to strike and then offer information to counteract it. Of course, while advances in AI make this possible in formal settings, it can't account for decisions that take place off-line and in conversations. And building such systems could eat up a lot of time and money.

A low-tech but highly effective technique for fighting linear bias is data visualization. As you've noticed in this article, whenever we wanted you to understand some linear bias, we showed you the nonlinear relationships. They're much easier to grasp when plotted out in a chart than when described in a list of statistics. A visual representation also helps you see threshold points where outcomes change dramatically and gives you a good sense of the degree of nonlinearity in play.

Putting charts of nonlinear relationships in dashboards and even mapping them out in "what if" scenarios will make managers more familiar with nonlinearity and thus more likely to check for it before making decisions.

Visualization is also a good tool for companies interested in helping customers make good decisions. For example, to make drivers aware of how little time they save by accelerating when they're already traveling at high speed, you could add a visual cue for time savings to car dashboards. One way to do this is

with what Eyal Pe'er and Eyal Gamliel call a “paceometer,” which shows how many minutes it takes to drive 10 miles. It will surprise most drivers that going from 40 to 65 will save you about six minutes per 10 miles, but going from 65 to 90 saves only about two and a half minutes—even though you’re increasing your speed 25 miles per hour in both instances.




The Implications for Marketers

A cornerstone of modern marketing is the idea that by focusing more on consumer benefits than on product attributes, you can sell more. Apple, for instance, realized that people would perceive an MP3 player that provided “1,000 songs in your pocket” to be more attractive than one with an “internal storage capacity of 5GB.”

Our framework, however, highlights the fact that in many situations companies actually profit from promoting attributes rather than benefits. They’re taking advantage of consumers’ tendency to assume that the relationship between attributes and benefits is linear. And that is not always the case.

We can list any number of instances where showing customers the actual benefits would reveal where they may be overspending and probably change their buying behavior: printer pages per minute, points in loyalty programs, and sun protection factor, to name just a few. Bandwidth upgrades are another good example. Our research shows that internet connections are priced linearly: Consumers pay the same for increases in speed from a low base and from a high base. But the relationship between download speed and download time is nonlinear. As download speed increases, download time drops rapidly at first and then gradually. Upgrading from five to 25 megabits per second will lead to time savings of 21 minutes per gigabyte, while the increase from 25 to 100

Mbps buys only four minutes. When consumers see the actual gains from raising their speed to 100 Mbps, they may prefer a cheaper, slower internet connection.



Of course, willfully exploiting consumers' flawed perceptions of attribute-benefit relationships is a questionable marketing strategy. It's widely regarded as unethical for companies to take advantage of customers' ignorance.

CONCLUSION

In recent years a number of professions, including ecologists, physiologists, and physicians, have begun to routinely factor nonlinear relationships into their decision making. But nonlinearity is just as prevalent in the business world as anywhere else. It's time that management professionals joined these other disciplines in developing greater awareness of the pitfalls of linear thinking in a nonlinear world. This will increase their ability to choose wisely—and to help the people around them make good decisions too.