

When the Gallup Organization applied Six Sigma principles to sales and service groups at several companies, it learned how much performance variation exists between seemingly similar work groups. Managing that variability can raise overall performance by orders of magnitude and can create organic growth.

MANAGE YOUR HUMAN

SIGMA

by John H. Fleming, Curt Coffman, and James K. Harter

“Quality” is easy to measure and manage in some contexts, and extremely difficult in others. Businesspeople have a pretty good idea how to judge the manufacturing process that yields a snazzy new handheld device, for example. But what about the retail employee’s attempts to sell the gadget? Or the call center employee’s efforts to help the customer navigate its eccentricities? Businesses aren’t especially good at measuring and managing the quality of those processes—or indeed of most work done by non-manufacturing businesses and units.

Yet it's essential that organizations learn to measure and manage quality in *all* kinds of business settings. In manufacturing, value is created on the factory floor. In sales and service organizations, and in many professional service firms, value is created when an employee interacts with a customer. Indeed, the employee-customer encounter *is* the factory floor of sales and services. If these organizations are going to achieve meaningful operational and financial improvements, the employee-customer encounter must be managed with great care.

Quality improvement methodologies such as Six Sigma are extremely useful in manufacturing contexts, where ingredients with predictable properties are repeatedly combined in the same ways, but they're less useful when it comes to the employee-customer encounter, with its volatile human dimensions. To address this problem of fit, we've developed a quality improvement approach that we call Human Sigma. Like Six Sigma, Human Sigma focuses on reducing variability and improving performance. But while Six Sigma applies to processes, systems, and

• To improve the quality of the employee-customer interaction, organizations must conduct both short-term, transactional interventions (such as coaching) and long-term, transformational ones (such as changing the processes for hiring and promotion). In addition, the company's organizational structure often must be adjusted so that the employee-customer encounter can be managed holistically.

Human Sigma grew out of a multiyear, research-based initiative designed to map the terrain of the employee-customer encounter. We identified ways to measure the effectiveness of the encounter, explored how those metrics could best be used, and assessed the benefits that could result from their application. This work was based on direct experience with hundreds of companies and millions of customers and employees. We then tested and cross-validated our findings in 1,979 business units—involved in financial services, professional services, retail, and sales—within ten companies. The results thus far have been extraordinary. The ten companies, all of which have applied

IT'S POSSIBLE to arrive at a single measure of effectiveness for the employee-customer encounter; this measure has a high correlation with financial performance.

output quality, our approach looks at the quality of the employee-customer encounter, weaving together a consistent method for assessing it and a disciplined process for managing and improving it.

As we developed our thinking about Human Sigma, we arrived at several core principles for measuring and managing interactions between customers and employees:

• It's important not to think like an economist or an engineer when you're assessing the employee-customer interaction. Emotions, it turns out, inform both sides' judgments and behavior even more powerfully than rationality does.

• The employee-customer encounter must be measured and managed locally, because there are enormous variations in quality at the work-group and individual levels.

• It's possible to arrive at a single measure of effectiveness for the employee-customer encounter; this measure has a high correlation with financial performance.

John H. Fleming (fleming_hs@gallup.com) is the chief scientist for customer engagement, Curt Coffman (coffman_hs@gallup.com) is a global practice leader, and James K. Harter (harter_hs@gallup.com) is the chief scientist for employee engagement at the Gallup Organization. They are based in Princeton, New Jersey; Denver; and Omaha, Nebraska, respectively.

the best-practice principles for managing the employee-customer encounter, together outperformed their five largest peers during 2003 by 26% in gross margins and by 85% in sales growth. We can't guarantee readers comparable results, but we believe that closely monitoring the health of a firm's employee-customer relationships will result in dramatic performance improvements.

Emotions Frame the Encounter

Six Sigma processes are data driven, rational, and analytic. They focus on conformance to requirements, which are generally specified in functional terms. Does the product have any defects? Are its parameters within specified manufacturing tolerances? Is it delivered on time? Widespread use of Six Sigma and TQM methodologies has resulted in vastly improved product quality over the past two decades.

Inspired by these improvements, businesses have tried to apply Six Sigma principles in sales and service settings. In early attempts, researchers and managers alike assumed that the customers in those settings would be as focused on conformance to requirements as the engineers on the factory floor were. Had this been the case—had customers been rational creatures who judged their interac-

tions with company representatives using rigorous, analytical standards – then simple flawlessness on the company's part would have resulted in satisfied, profitable, lifelong customers.

But nothing human is ever that simple. People may think that their behavior is purely rational, but it rarely is. Twenty years of research in two very different fields – neuroscience and behavioral economics – has established quite clearly that people base their decisions on a complicated mixture of emotion and reason. Indeed, recent work suggests that emotions may play a larger role than analysis.

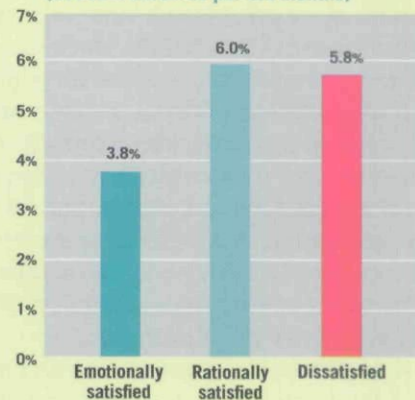
Customer Engagement. That work in neuroscience and behavioral economics is borne out by research into customer satisfaction and engagement. Results from a large and growing number of case studies suggest that “extremely satisfied” customers (people who provide the highest rating of overall satisfaction with a company's products and services) fall into two distinct groups: those who have a strong emotional connection to the company and those who do not. When we examine indicators of customer behavior (such as attrition, frequency of use, total revenue, and total spending), a clear and striking pattern emerges. Emotionally satisfied customers contribute far more to the bottom line than rationally satisfied customers do, even though they are equally “satisfied.” In fact, the behavior of rationally satisfied customers looks no different from that of *dissatisfied* customers. The pattern shown in the exhibit “Emotional Satisfaction Matters Most” has emerged in every study we have examined.

Imagine that you could peek inside the heads of your customers as they thought about your company. Would people with a strong emotional connection to the firm show different brain activity than other customers? As it turns out, the answer is yes. We studied three groups of customers of a luxury retailer in Japan. One group was strongly attached emotionally (according to our measure of emotional attachment), one was moderately attached, and the third had little or no attachment. While inside a functional magnetic resonance imaging (fMRI) machine, the customers responded to simple agree-or-disagree statements about the retailer, about their bank, and about various aspects of daily life. The brains of customers who had the strongest levels of emotional attachment to the retailer were significantly more active while the subjects were thinking about the company. The increased activity was concentrated in parts of the brain related to emotion, emotional-cognitive processing, and memory. Moreover, the enhanced brain activity was company specific; customers who were passionate about the retailer but not the bank did not show the same enhanced levels of neural activity when thinking about the bank. (The attitude survey that had been used to separate the subjects into three groups proved to be a good proxy for the fMRI study, in that it reliably predicted which individuals would

Emotional Satisfaction Matters Most

At a large U.S. retail bank, the attrition rate of dissatisfied customers was scarcely different from that of “rationally satisfied” customers, those who described themselves as extremely satisfied but scored low on an emotional-attachment metric that measures four dimensions – confidence, integrity, pride, and passion. By contrast, the attrition rate of people who were “emotionally satisfied” by the bank was, on average, 37% lower. Similarly, dissatisfied customers of an international credit card provider were virtually indistinguishable from rationally satisfied cardholders in their purchase behavior, while customers who were emotionally satisfied by factors such as service, features, and brand image spent more, on average, than people in the other groups. (The emotionally satisfied group also increased its spending by 67% over 12 months, compared with 8% for the rationally satisfied group; there was a small decrease within the dissatisfied group.)

ATTRITION RATES OF BANK CUSTOMERS
(account closures per six months)



AVERAGE MONTHLY SPENDING BY CREDIT CARD CUSTOMERS



>> THE HIGH-PERFORMANCE ORGANIZATION

show the enhanced activity levels). Even more striking was the relationship between emotional attachment and self-reported share of spending, which were strongly correlated at 0.6 on a scale of -1 to +1. This suggests to us that there is an underlying neurological mechanism that links emotional attachment to subsequent behavior.

Clearly, a Six Sigma approach to measuring and managing the quality of the employee-customer interaction needs to take customers' emotions into account. Building on the work of psychologist Ben Schneider and management professor David Bowen, we have developed just such a measure of customer engagement. It combines traditional metrics of customer loyalty (overall satisfaction, likelihood to repurchase, and likelihood to recommend) with a short battery of items that assesses the emotional nature of customers' commitment. The first dimension it looks at is *confidence*. Does this company always deliver on its promises? Are its people competent? The second is *integrity*. Does this company treat me the way I deserve to be treated? If something goes awry, can I count on the company to fix it fast? The next element is *pride*, a sense of positive identification with the company. The fourth dimension is *passion*. Is the company irreplaceable in my life and a perfect fit for me? Truly passionate customers, by the way, are relatively rare. They are customers for life, and they are worth their weight in gold.

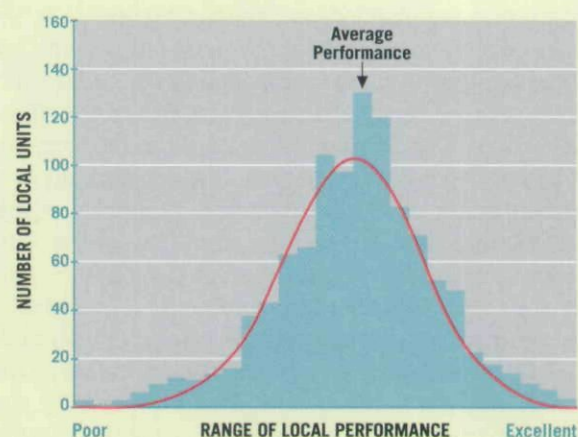
Our research suggests that for all kinds of companies, fully engaged customers—those who score in roughly the upper 15% to 20% on our measure—deliver a 23% premium over the average customer in terms of share of wallet, profitability, revenue, and relationship growth. Actively disengaged customers—those who score in the bottom 20% to 30%—represent a 13% discount on the same measures. And within a given company, business units whose levels of customer engagement are in the top 25% tend to outperform all other units on measures of profit contribution, sales, and growth by a factor of 2:1.

Employee Engagement. Every interaction an employee has with a customer represents an opportunity to build that customer's emotional connection—or to diminish it. Obviously, these interactions are not the only way to the customer's heart, but they are a large and largely untapped resource. In the United States, just 29% of employees are energized and committed at work, according to Gallup Poll data. Perhaps more distressing is that 54% are effectively neutral—they show up and do what is expected, but little more. The remaining employees, almost two out of ten, are disengaged.

Work groups whose members are positively engaged have higher levels of productivity and profitability, better safety and attendance records, and higher levels of retention. Not surprisingly, they're also more effective at engaging the customers they serve. Disengaged employees have a profound impact, too. We estimate that they cost companies \$300 billion per year in lost productivity in the

It All Depends on Which Store You're In

Levels of customer engagement vary widely across the 1,100 stores of a retail chain we studied. Each bar represents the number of stores that fall into one of 28 customer-engagement performance bands, with poorly performing stores on the left and exceptional performers on the right. The top stores' performance is 3.5 times as strong as the poorest stores'. The shape of the curve (a normal distribution) suggests that the variability is unmanaged.



United States alone. They also destroy customer relationships with remarkable facility, day in and day out.

Performance metrics that acknowledge the importance of emotional engagement—on the part of both customers and employees—provide much stronger links to desired financial and operational outcomes. But deciding which metrics to use is just the first step toward effective management of the employee-customer encounter. Deciding *how* to deploy them is equally important. Unfortunately, in many companies, metrics designed with the right intentions are often deployed in the wrong ways.

The Encounter Must Be Measured Locally

We have all seen the claims: A major airline touts itself as an industry leader in on-time performance and has the flight departure and arrival data to prove it. A cellular provider claims to be a leader in customer satisfaction, citing an independent study of customers. A retailer announces that it has won an award for being one of the country's best places to work for the fifth year in a row. Each of these summary claims—based on the results of

surveys—may be legitimate, but quick reviews of the on-time performance of specific flights, or candid conversations with cellular customers, or visits to several stores in the retail chain, inevitably reveal a considerable range of performance hidden behind the averages. Some flights are never on time; some always are. Some customers experience nothing but problems; others are routinely delighted. And some stores are exceptional places to work, while others are awful. High-level averages of company performance may provide good marketing copy, and they may make executives feel better about their position in the marketplace. But because they obscure the considerable variation from location to location within a company, they don't give managers and executives the information they need to improve performance.

Local variability shows up on virtually every performance metric we have examined. And it tends to be vast. In fact, the variations within a company easily dwarf the differences between competitors. Also, performance roughly follows a normal distribution, suggesting that local variability is largely unmanaged. (See the exhibit

lot about organizational performance. Let's say you manage one of several customer service call centers operated by a large telecommunications provider that we'll call Telecom A. Like its sibling centers, yours is a state-of-the-art facility, with an integrated CRM system that allows your CSRs to access each customer's relationship with the company—including account activity, revenue, and profitability—in real time. Calls are routed automatically to make the most efficient use of capacity. Every CSR is comprehensively trained, monitored, and coached, and there's little variation in the reps' pay from center to center.

To assess how well it is meeting its customers' requirements, Telecom A measures satisfaction at the company level by regularly surveying, and providing feedback from, a random sample of people who have recently called. Telecom A also conducts an annual employee survey. When you receive your copy of the quarterly customer satisfaction scorecard, you find that 88% of callers were satisfied with the service they received. The employee survey, meanwhile, reveals that just 40% of workers companywide feel they are adequately compensated.

FULLY ENGAGED CUSTOMERS DELIVER a 23% premium over the average customer in terms of share of wallet, profitability, revenue, and relationship growth.

“It All Depends on Which Store You're In.”) For sales and service organizations, unmanaged variability in the quality of the customer experience represents a significant threat to the enterprise's sustainability, because customers experience variation, not averages. Exactly the same pattern of performance variability emerges on employee measures, as well, with similar implications.

The only way to improve local performance is to provide feedback at the level where the variability originates. Suppose that instead of assessing your own heart rate, your physician based treatment on a measurement of the average heart rate for your entire town. It sounds absurd, but in many companies, something akin to this happens every day. The employee-customer encounter is assessed at the wrong level of specificity for the measurement to be useful. What does a cellular provider's description of itself as “an industry leader in customer satisfaction” mean to a customer who is routinely confronted with sub-par service at a local level? And what does a company's label as “one of the country's best places to work” mean to an employee whose local workplace is miserable and depressing?

When the employee-customer encounter is assessed at the level of the local work group, executives can learn a

What exactly does this information tell you? Not very much. To truly understand the totality of the employee-customer encounter, you need metrics that go deeper into the organization. Fortunately, Telecom A has deployed just such metrics, and they have produced some startling insights.

One insight—and this is borne out by one of the largest CSR-level studies ever conducted (including some 5,000 reps)—is that the customer's experience still depends almost entirely on the particular rep who takes the call. The best 10% of CSRs produce six positive interactions for every negative one, based on postcontact interviews with customers. The worst 10% yield only three positive for every four negative encounters. Critical information of this type was hidden behind the overall summary score of 88% customer satisfaction. Without the deeper metrics, you as a call center manager would have been unable to identify or manage the sources of both poor and exceptional performance.

Or consider Bank B. Some time ago, its top executives recognized that employees affect profitability through two separate paths. The first might be described as *direct cost efficiencies*. Committed employees generate greater output at a higher quality level than uncommitted workers.

They also stay longer with the firm, reducing training and replacement expenses. These efficiencies translate directly into enhanced profitability. The second path could be called *indirect customer outcomes*. Productive and committed employees generate stronger customer connections, which lead to higher levels of customer retention, profitability, and growth.

Early in their efforts to understand how to boost employee productivity and commitment, Bank B executives routinely assessed workers' opinions by surveying a random sample. They hoped to identify a key set of issues that, if improved, would make employees happier and more productive. The results were disappointing. It was not until they assessed worker attitudes at the branch level that they started to make progress. At the branches, employee attitudes ran the gamut from delight to disgust. Because Bank B measured at the correct level of specificity, it discovered that some local work groups epitomized the highest standards of excellence, while others were totally demoralized.

Local performance variation is the scourge of organizations that aspire to high performance. While it is in the nature of performance distributions to show variation (after all, "average" is simply a summary that represents almost no one's actual experience), the magnitude of the variability is a critical measure of organizational health. More than two decades ago, W. Edwards Deming and Joseph Juran noted that variability on critical performance metrics is a threat to the vitality of an enterprise because it is evidence that the business is not being managed effectively. And intuitively, it makes sense that the greater the range of performance on critical performance measures, the more costly the business is to operate.

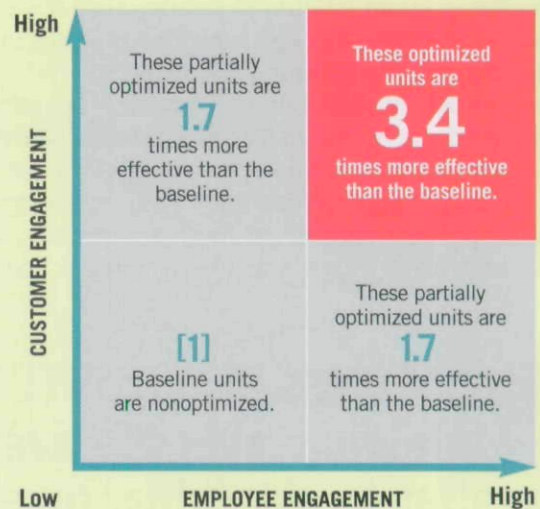
Unfortunately, in most organizations, variability in the effectiveness of the employee-customer encounter goes largely undiagnosed. As a result, revenues and profits are bled off, and growth is anemic. The extensive range of local performance variation that exists in every company we've studied means that there is really no such thing as a single corporate culture or unified brand. There are as many cultures and "brands" as there are local work groups and customer touch points.

Local managers sometimes blame variability from location to location on factors such as store size, age, or locale that are beyond their control. Our research doesn't back them up. For example, within a chain of retail stores, controlling for those and other "immutable" variables – including local demographics and the presence or absence of competitors – eliminates only a portion of the performance variation among stores.

What explains this local variability? We've controlled for the factors that can't be changed. And the factors that are common across the enterprise – product, price, processes, policies, and so forth – can't, by definition, explain local variability (they often play a critical role in driving

The Interaction of Employee and Customer Engagement

Local business units with even moderately high levels of both worker and customer engagement are, on average, more effective financially than units with very high levels of only one form of engagement.



customer engagement, of course). If these factors don't differ from place to place, the only remaining culprit is the way those processes and policies are implemented locally. But that brings us to a consideration of exactly who is doing the implementing and how the implementation is being managed. To reduce variability in the customer experience, businesses must focus on reducing variability in local "people" processes (the "who" and "how" of implementation). The power of a local focus on reducing variability lies in its simplicity and flexibility. Each unit can identify and correct its own problems.

The Link to Financial Vitality

Conventional analyses of employee attitudes, customer requirements, and financial performance have emphasized the linearity of the relationships among them: Employee attitudes affect customer attitudes, and customer attitudes affect financial performance. We believe that the three factors also interact in complicated ways. Our Human Sigma metric combines employee and customer engagement into a single measurement that, we believe, provides a more comprehensive way of capturing and understanding this dynamic system.

The Human Sigma model grew out of a partially failed experiment. Several years ago, we were working with a large, multisite retailer on two separate initiatives to measure and improve its relationships with its employees and its customers. By surveying all workers as well as a sample of customers at each store, we were able to provide metrics for both relationships at the local level. We also found, not too surprisingly, that scores on both measures were strongly linked to the stores' financial performance.

As the project evolved, we wanted to understand what the top performers on each measure did differently from their less-stellar counterparts. We first identified the ten highest-performing stores on the basis of employee engagement, then did the same for customer engagement. Our working assumption was that at least a few of the top employee-engagement stores would also be top customer-engagement stores. We were wrong. Just one store appeared on both lists.

As we thought about that finding, we returned to the data and noticed two things: As we expected, stores that performed well (defined as simply being in the top half, rather than the top ten) on both employee and customer

in earnings per square foot of retail space than the remaining stores – a difference that translated into more than \$32 million in additional annual profits for the entire chain. The exhibit “The Interaction of Employee and Customer Engagement” shows how the average net gain per business unit is associated with low and high engagement of workers and customers.

As we have refined the Human Sigma concept, we have developed a method for combining employee and customer engagement scores at the local unit level to yield a single score that is reliably related to the unit's overall financial vitality. (See the sidebar “The Math Behind the Human Sigma Score.”) This score allows us to classify units into six broad performance levels. Units in the lower two levels are in dire need of improvement: Those that engage employees without engaging customers have become too inwardly focused and have lost direction. Those that engage customers without engaging employees are living on borrowed time; over the long term, customer engagement will tend to erode. We consider units in the top three levels to be optimized. Obviously, we believe that sales and services companies should strive to

HIGH-LEVEL AVERAGES of corporate performance may provide good marketing copy, but they obscure the considerable variation within a company.

engagement produced considerably better financial results than those that did poorly on both measures. But stores that performed well on both metrics also outperformed stores that scored high on one but not the other. This observation suggested that customer and employee engagement interact to promote financial performance.

Our subsequent research has confirmed that customer and employee engagement augment each other at the local level, creating an opportunity for accelerated improvement and growth of overall financial performance. Our meta-analysis of the financial performance of the 1,979 business units in the ten companies in our present study reveals that local business units that score above our database median on both employee and customer engagement metrics are, on average, 3.4 times more effective financially (in terms of total sales and revenue, performance to target, and year-over-year gain in sales and revenue) than units that rank in the bottom half on both measures. The doubly stellar units are also roughly twice as effective financially as units that are high performers on one – but not both – of these critical vital signs. In one luxury retail chain, for example, the stores that scored high on both measures generated an average of \$21 more

move all of their local units into the top performance level. This means that, over time, local performance variability must be reduced and overall performance increased. While difficult, such improvement is indeed possible. And the movement of units into successively higher Human Sigma levels brings with it enhanced financial performance.

How to Get There

A detailed look at how to manage and reduce variability at the local level would turn into a lengthy discussion, so we will make just three quick points.

Responsibility for Human Sigma must be centralized. Since employee and customer engagement are intimately connected – and since, taken together, they have an outside effect on financial performance – they need to be managed holistically (at the same time that they're managed locally, which we'll get to in the next paragraph). That's easier said than done. In most companies, data about customers stay inside the marketing or quality department. Data about employee well-being reside, for the most part, in the HR department. And financial data,

of course, live in finance. But only when these data are brought together on a single platform can a true picture of the health of the employee-customer encounter be drawn. It is simply not sufficient to provide managers with

the employee-customer relationship must reside within a single organizational structure, with an executive champion who has the authority to initiate and manage change.

LOCAL PERFORMANCE variation is the scourge of organizations that aspire to high performance.

a “dashboard” of seemingly unrelated gauges and dials drawn from various and dispersed quarters of the organization. What this means in practice is that the responsibility for measuring and monitoring the health of

The local manager is nonetheless the single most important factor in local group performance. Local-level managers have a huge role to play, for better or worse, in local performance. Earlier Gallup research suggested that employees join great companies but leave poor managers. That is, employees join a company for a variety of both high-minded and practical reasons. But, invariably, their working lives revolve around local environments that can either nourish them and foster their learning or starve them, causing them ultimately to leave the company—or to hang around, unproductively waiting for retirement. Local managers whose work groups show suboptimal performance should be encouraged to use the familiar tool kit of interventions: targeted training, performance reviews, action learning, and individual coaching. And managers themselves should be supported in similar ways. If none of these interventions leads to better performance, the local manager should be replaced.

Some companies will need to overhaul their HR practices. A set of longer-term, transformational interventions may be necessary in some instances. Executives or outside consultants may need to reexamine how local leaders do their jobs, how these managers are being developed, and how decisions are made and executed at the local level. If the Human Sigma numbers throughout the organization are lower than expected, or if parts of the organization sustain low numbers over time, then a broader intervention may be needed. The company may need to look at how it selects employees, promotes people into management, does performance appraisals, approaches succession planning, and recognizes performance.

...

Ask any chief executive to list his or her most pressing business challenges, and you will no doubt hear concerns about customer and employee retention, authentic and sustainable growth, eroding margins, and cost efficiencies. Clearly, there is no single solution to those challenges. But we are confident that measuring and managing two simple factors—employee and customer engagement—can lead to breakthrough improvements in all aspects of your business. ▢

Reprint R0507J; HBR OnPoint 1533
To order, see page 195.

The Math Behind the Human Sigma Score

A business unit's Human Sigma score is computed by first converting its mean scores on employee and customer engagement into percentile equivalents (based on the observed distribution of scores for each metric). If a unit's converted scores on *both* metrics are above the median value for the distribution, the Human Sigma score is the square root of the product of the two percentile values, corrected for certain boundary conditions. (This correction value is equivalent to the ratio of the two percentiles—highest over lowest—raised to the 0.125 power.) If a unit's converted score on *either* metric is below the median value for the distribution, the Human Sigma score is the square root of the product of the two percentile values divided by 2. This produces a single bimodally distributed score that is then used to establish threshold values that define each of six Human Sigma levels, HS1 through HS6. The HS4 threshold is defined at 50. The HS3 threshold is defined as one standard deviation (SD) below that (using the standard deviation of the Human Sigma score distribution). The HS5 threshold is defined as one SD above the HS4 threshold. Successive thresholds are one SD away from the adjacent level. In algebraic terms: If employee engagement percentile and customer engagement percentile are both above 50, then:

$$HS = \sqrt{(EE\text{percentile} \times CE\text{percentile}) \times \left(\frac{\text{percentile Max}}{\text{percentile Min}}\right)^{0.125}}$$

If either employee engagement percentile or customer engagement percentile is less than or equal to 50, then:

$$HS = \sqrt{\frac{(EE\text{percentile} \times CE\text{percentile})}{2}}$$

Harvard Business Review and Harvard Business School Publishing content on EBSCOhost is licensed for the individual use of authorized EBSCOhost patrons at this institution and is not intended for use as assigned course material. Harvard Business School Publishing is pleased to grant permission to make this work available through "electronic reserves" or other means of digital access or transmission to students enrolled in a course. For rates and authorization regarding such course usage, contact permissions@hbsp.harvard.edu