dbkay
+associates

David Kay

DB Kay & Associates

Eleven Metrics Myths
(and What to Do About Them)

Metrics Myths Lead to Metrics Mistakes

Service and support organizations love their metrics. We install monitors throughout the office to display them to workers and managers. We discuss them in quarterly operations reviews and daily stand-ups. We hold people accountable to them. We send out weekly reports with them—in one organization's case, with a six megabyte Excel file attached! Metrics are a really big deal.

It's good to measure our business. But it's essential that we measure the right things, and use metrics in the right way. And unfortunately, over the last two decades, we've observed some organizations use metrics to drive counterproductive behavior, to reward activity over quality and value, and to demotivate workers.

The astonishing thing about the dysfunctional use of metrics is that it's almost never an accident. Leaders believe they're creating value with metrics when they're not. Most often, they're managing their teams the way they were managed as they came up through the ranks—no matter that the business of support and the nature of the workplace have changed so profoundly in the last twenty five years.

Because metrics mistakes are often based in deeply-held beliefs, it's important to examine these beliefs. In many cases, they're myths—beliefs that have no basis in reality.

In this e-book, we lay out eleven of the most pernicious myths we see in the industry, and suggest how to combat them. I hope you don't recognize your organization in these pages, but if you do, we hope we can help you identify and bust your organization's myths.

Metric Myth One: Goals Should be SMART

That is, individual performance goals should be Specific, Measurable, Attainable, Relevant, and Time-Bound…SMART. Doesn't that sound SMART? Don't I have to be a little DUMB to argue against SMART goals?

Well, not so fast. We're not on the assembly line any more.

F.W. Taylor, stopwatch and notebook in hand, invented the practice of scientific management. He did painstaking research on the best way for men to shovel material (really!) and extrapolated general principles for making workers more efficient from these time and motion studies. He advocated paying workers by the piece, rather than a fixed salary. He believed that there is One Best Way of doing any task, and that management's job is to make sure that One Best Way is exactly how workers do it.

F.W. Taylor would have absolutely loved SMART goals.

The nature of work has changed in the last hundred years since Taylor did his groundbreaking (given his fascination with shovels, *literally* groundbreaking) analyses. There's now far less correlation between the units of work accomplished and the value that work creates. Would you measure a programmer by how many lines of code she creates? Or an advertising copywriter by the word? I expect you wouldn't like what you got.

KCS practitioners know that, when it comes to activity, more is not necessarily better. *As the KCS Practices Guide says,*

> *During the KCS adoption process, we have seen organizations put goals on KCS article creation (everyone should create five KCS articles a week) or KCS article reuse (analysts will be measured on how often they reuse KCS articles). The goals for these leading indicators may have been met, but the quality of the knowledge base has been seriously compromised. Invalid and duplicate KCS articles are created, because the focus is on the activity, not the outcome. Worse, emphasis can shift to gaming the system rather than generating real value. Inevitably, quality and morale suffer, management looks less competent, and the value of the knowledge is diminished.*

Unfortunately, it's these activity measures that are just the kind of thing that are Specific and Measurable. We believe there is no single golden metric for the value created in the knowledgebase. Collaboration, which is notoriously hard to measure, is an important piece of employees' contribution. We believe managers must do the hard work of providing feedback and assessing value by considering activities, outcomes, and quality measures in the context of the job and the environment—a judgment that's driven by data, but is ultimately a *human* judgment, made in an environment of uncertainty. Anyone who tells you there's a simpler way is selling you something.

SMART goals are comfortable. They're reassuring. You make them or you don't, and it's all up to you.

If you want comfortable and reassuring work, it's probably time to pick up a shovel. Knowledge work has many advantages, but being easy to measure isn't one. For support delivery staff developing knowledge, it's time to bury SMART goals.

Metrics Myth Two: Support Leaders Should Have Direct Control Over their Metrics and their Goals

"Men at some time are masters of their fates." — Cassius, to Brutus, Julius Caesar, Act I, Scene 2.

It's almost axiomatic that we should sign up for goals only if we have control over achieving them. In *Julius Caesar,* Cassius berates Brutus for allowing himself to be treated as an underling; so too, conventional wisdom says that executives shouldn't allow themselves be held responsible for the performance of other organizations.

It's a nice theory. Unfortunately, it doesn't work in the real world.

First, like it or not, Support can't be measured separately from the rest of the company. Perceptions of Support are deeply enmeshed with perceptions about the product, its quality, and its supportability; by company policy; and by margin requirements levied by the CFO and Wall Street. Studies show that the single most significant factor determining Support customer satisfaction is product quality. No organization is an island, and least of all Support.

Secondly, what's the point in optimizing the performance of Support—or any other organizational silo—if has the side effect of sub-optimizing the performance of the business as a whole? For example, Support can deliver great First Contact Resolution (FCR)…if they're solving the same, well-known problems over and over again. This is great for Support's numbers, but terrible for customers and the company.

Patrick Lencioni is a careful student of corporate culture. In his book Silos, Politics, and Turf Wars, he dramatizes a number of environments in which various organizations were at each others' throats. While industries and details differed, the common theme was that each organization was working to meet its own objectives independent of those of other organizations. This put the various orgs at cross-purposes. Predictable conflict ensued.

This is old news to those of us who have been working in high tech for any length of time. Sales struggles to make their numbers, then blames product delays, feature deficiencies, and Marketing ineptitude. Product implementations go off the rails; Implementation Services blames Development, while Development wonders who hired these wet-behind-the-ears implementation teams.The Product Management team wonders why no one ever builds the impossible products it dreams up, and every-one knows what happens to Support. It's a mess.

In this environment, does it really make sense for us to isolate ourselves from each other's metrics? Because clearly, our successes can't be independent from that of other organizations. As a CEO once told me early in my career, "David, there's no way your department can succeed if the company fails."

Lencioni proposes a two-part solution: a rallying cry (or overarching vision) for the whole organization, and cross-functional metrics to ensure that we're all working together and supporting each other in executing on that vision.

This is a reason I love the Net Promoter Score (NPS), at least when it's used right. (It isn't always used right; we have a Metrics Myth about that below.) No one organization can own NPS, and everyone needs to work together to improve it. Unlike our siloed metrics of FCR, service level adherence, or customer satisfaction, something like Net Promoter Score gives Support a seat at the table. It gives a reason for Development to listen to us when we talk about customer experience improvements in the product, and it gives Marketing a reason to pay attention when we make the case for exposing known issues.

As with SMART metrics, keeping our heads down and focusing on Support metrics is more comfortable than living in a world where we're dependent on others. But, since we can't be successful without the rest of the enterprise, let's pick cross-functional metrics that line us up in the same direction instead.

Metrics Myth Three: Managers Should Set Ambitious and Inspiring Goals

Yes, but. This is one of these partial truths where getting the details right really matters.

Followers of popular business books will recall the BHAG—the Big Hairy Audacious Goal—proposed by James Collins and Jerry Porras in Built to Last. These goals are long-term, viewed as a long shot by outsiders, but considered possible by insiders. They act as a unifying vision for the company, much as Lencioni's "Rallying Cry" did in the last myth.

This kind of overarching, companywide, long-term goal is great, although, given how "goal" is commonly used at work, I'd prefer to call it a "measurable vision." Great executives set these, communicate them, and stick with them.

But most of the goals we have in companies aren't like that. They tend to be more tactical, more operational, and shorter term. You can't have a BHAG on service level compliance—BHAGs are a brass ring for the company overall, which service levels most decidedly are not. So executives should stay away from these operational goals.

Why? Because they're a means to an end—a way of achieving the BHAG. And it's the people actually doing the work who are in the best position to know how to get there, and how to hold themselves accountable to the work.

Those of you who have been to a KCS Foundations or Design Workshop may remember the Process and Change exercise we do with a ball. While I'm not going to spoil it for future participants, I'll say that the team estimates its own ability to improve its processes, and, especially after a few iterations, it's remarkably good at setting its own goals. The team is better than I as the facilitator would be, and I've led the exercise tens of times. But they have a secret ingredient that I don't have: they know each other, and they have hands-on experience with the process they're measuring.

Forget the director in the corner office. It's way better for the people doing the work to set their own goals.

Now a cynic might ask, won't they sandbag? Won't they set easy goals for themselves? My experience is, unless the organization has completely stamped initiative and self-respect out of a team (in which case you need a different e-book, or maybe just a different job), teams will challenge themselves and each other. Almost all the people we work with want to do well, they want to feel a sense of accomplishment, and they want to be part of a high-functioning team.

Not only will teams do a better job of setting goals than their managers will, they'll have that special something that comes from owning a goal. If you tell me I have to do something, I might just argue with you. But if I tell you I'm going to do something, just try and stop me.

So, executives should set an ambitious, exciting, and plausible vision. When it comes to setting the goals along the way, let the people doing the work take the lead. Support them. And get out of their way.

Metrics Myth Four: If We're Doing a Great Job of Delivering Customer Satisfaction and First Contact Resolution at High Efficiency, We're Doing a Great Job of Support

Because, you know, that's what support does. Close cases. Quickly. In a way that satisfies customers. Mission accomplished; let's grab a beer!

Not so fast.

This myth begs the question of why the Support organization exists. Because most of the activity, headcount, and budget are tied up in closing cases, it's easy to assume that the business of Support is closing cases, full stop. But that's far too limited a view of our business.

Support is in the customer success business—as a colleague likes to say, helping customers *receive* and *perceive* value from the product. Closing cases can be part of that—a means to an end—but it's not the most important part, and it's certainly not the most valuable.

First, let me back up my statement about why Support exists. When I ask a group of Support executives why they do what they do, they often start out with (true!) statements like, "the product breaks." Using a gentle version of the five whys, if I keep asking why that's important, we always get to some version of, "if the customer doesn't get value from the product, they won't buy more, and they won't tell their friends to buy, either." Bingo. There's no intrinsic value to the business in closing cases; it's all about customer retention, loyalty, lifetime value, and referenceability. Support feeds the company's engines of growth and sustainability. The Technology Services Industry Association has made this topic a focus of their research; their findings are summarized in Consumption Economics.

But even if we believe that Support's mission is to make customers successful, it's still easy to fall back on our old operational measures: cases closed, service levels, customer satisfaction (CSAT), and the like. What's wrong with that?

The problem is that it dramatically overstates the significance of the assisted case-closing channel as a way of delivering customer value. Cases are both more rare, and deliver less value, than most support organizations think.

If customers have a problem with our products, or are struggling to use them, opening a case is surprisingly low on their list of next steps. If you consider your own experience with technology, other approaches tend to happen first. Think about it a minute…what did you do last time you had a problem? If you're like most people, you Googled it. Or you asked a friend or colleague. (I should know. I'm often that friend or colleague. And then I generally Google it!)

Sometimes I go to the vendor's site and seek out an answer there. Sometimes I'll start a new thread in a community, although I'll do so only after looking for a while, because I don't want to be That Guy who re-posts an FAQ. I do open support cases, but not usually.

Now you might say, as a consumer and small business owner, I'm not entitled to great assisted support, so I tough it out, making do with free resources. But you know what? I'd really rather figure it out on my own—even if the vendor is there to help me, I'd rather not jump through their hoops and explain it all to someone else if I can be self-sufficient. Many people feel the same way.

At least, that's what the numbers show. As we look at our clients' data, enterprise customers open a case only a small single-digit percentage of the time they want help. (The number is much, much smaller for B2C support.) The rest of the time, they're using Google, the self-service site, in-product help, and social networks of various kinds, including of the sneakernet variety. Sure, they tend to open cases for higher severity issues, but when it comes to customer success, every moment of need is important.

Do a perfect job on 100% of the cases you receive, and you're only making a difference 5% of the time. Or less.

So, when **do** customers open a case? Generally when things have hit the proverbial fan. When there has already been significant value erosion. The system is down; something needs to be replaced; it's throwing an error and not completing a task; it isn't doing something it is supposed to be able to do; or, it's doing something it shouldn't be doing. By the time the customer picks up the phone or logs in to our support portal, we have a problem, Houston, and we're in recovery mode.

Closing a case quickly and well simply makes the problem go away. It gets customers back to where they expected to be. It doesn't make things better than expected. Certainly, herculean recovery efforts can engender a sense of personal gratitude, which is nice, but it's doing a job well that never should have had to have been done in the first place.

Of course, once a case has been opened, closing it quickly and well is the best outcome. But that doesn't always happen. Then things get really ugly, as you know from having been on both sides of this experience.

So, closing a case can be important, but only in so far as it gets us back to where we started. What else can support organizations do?

- Provide feedback to Development to help them improve the product and the customer experience, avoiding problems in the first place
- Expand into a Customer Success Management role that proactively advises customers to help them get the most value from their products
- Build and implement configuration health checks and other utilities that help customers avoid problems in the first place
- Create knowledgebase articles, video tutorials, and automated fixes to empower customers with the expertise they need to be successful themselves, and to get themselves back to value more quickly when needed

All of these tasks help customers receive and perceive value more than simply closing a case. What else do they have in common? Most support organizations are too busy to do them. Why? **They're closing too many cases.** Hmm.

Metrics Myth Five: There are Good Industry-Standard Measures. Use Them.

It's a cliché that you get what you measure. So before we start picking from the standard list of reports provided by your KM, CRM, or BI system, let me ask you: *what do you want to get?*

If what you want is the same thing as everyone else, then the standard measures should do just fine. But if you have ambitions to stand out…if you have a new insight about how to do support…if you're seeking competitive advantage, I'd say, not so fast.

When Southwest Airlines emerged into prominence, there were many standard airline metrics to choose from: occupancy, revenue per available seat mile, and the like. But Southwest's then-CEO, Herb Kelleher, recognized that their chief limit to growth was the number of extremely expensive airplanes they could afford to own. Or, to be more precise, their revenue is bounded by the number of airplanes in the air. This led to a simple insight: an airplane on the ground isn't making money, so let's minimize the time spent on the ground.

To put this insight into action, Southwest started measuring and improving their *turnaround time*, or the length of time their planes spend idle at the gate. Much of what's unique about Southwest, from flight attendants who clean the planes as passengers are leaving, to boarding by groups and unassigned seating, are the result of Southwest's obsession with turnaround time. The result is sustained profitability in a historically unprofitable industry, and the equivalent passenger capacity of 35 additional multi-million dollar 737s. "Strategic" trumped "standard" for Southwest.

Other industry leaders have taken the same path. Dell focused on its Cash Conversion Cycle (CCC), the length of time between paying vendors for components and receiving money from customers. Dell's remarkable build-on-demand infrastructure, leverage with suppliers, and streamlined operations resulted in a negative CCC—that is, Dell gets paid by customers several weeks before it has to pay its suppliers. This cash float is a real engine for growth, and it's there because Dell focused on its own, slightly non-standard, metric.

We've seen this in the support industry, too. When innovative support and customer experience leader Brad Smith was at the helm at Openwave Systems, he had his team focus on margin dollars per case. He reasoned that he wanted more margin (through efficiency and customer retention), and fewer cases (through product improvements, self-service, and proactive support). The more margin, and the fewer cases, the better their margin per case. I don't know if any other organization has measured margin per case—and it certainly isn't in any standard Support dashboard. But Brad and his team were completely focused on higher renewals, more efficient case handling, and helping customers avoid cases…all great things for Openwave and its customers.

"Transformational measures can help organizations focus on what is most important today and for the future," says Dr. Dean Spitzer in his outstanding book Transforming Performance Measurement. "When we change our ways of measurement, the fundamental 'lens' used to view things changes. Organizational transformation is what happens when people begin to see their organization through the new lens."

If you're not sure what your strategy should be, or how you're going to transform your organization for the better, then by all means, use the same measures as everyone else. But if you have a vision for where you want to go—and I hope that you do—then pick what Spitzer calls transformational measures to track and motivate your progress on that journey.

Metrics Myth Six: Do an ROI Analysis Before Investing in a Major Purchase or Initiative

There's nothing wrong with doing an ROI analysis—as a matter of fact, I think it's a great idea. It's just that the way that we, as an industry, typically use them is all backwards. They're shouldn't be a check-point; they're a *process*.

Let's back up and define our terms. A Return on Investment (ROI) analysis is a financial model that calculates what net present value (NPV) implementing something, generally a technology, will deliver relative to its cost. This can be expressed as a percentage ("returns will be 183% of investment within 24 months") or a time ("We'll break even after 14 months, and after that everything is upside.") Benefits need to be quantified in hard dollars: "better customer experience" doesn't count as a return in an ROI model. A thoughtful ROI captures all sorts of investments—not just software licenses or subscriptions, but also hardware, implementation costs, support contracts, ongoing maintenance, training, and project management.

Right away, it's clear that ROIs are tricky. Technology investments, by themselves, almost never generate benefits. Benefits only accrue from technology when it supports process improvements, and when people follow the new processes. Predicting the costs and benefits of an entire process change initiative is harder, and more speculative, than just adding up the costs of technology.

Still, as a sanity check, it's a reasonable plan to come up with a rough estimate of ROI. In particular, you might find out that the implementation is so expensive that you won't make up the investment in any reasonable time period. In this case, it's good to stop and reconsider your initiative.

So, why do I list ROI as a "Metrics Myth?"

1. The complexity of program-wide ROIs means that non-technology costs tend to get swept under the rug, and we pretend that the new knowledgebase, search engine, or CRM tool will magically shorten handle times or deflect contacts by itself. Needless to say, this kind of magic thinking isn't conducive to good decision-making.

2. Also, as an industry, we tend to know what good ROIs look like. A 10,000% or even 1,000% near-term return on investment isn't seen as credible (although knowledge programs can, in fact, do this!) A near break-even ROI isn't compelling. Spreadsheets make it easy to change assumptions to see what happens. So, of course, program teams all turn in ROIs that are high enough to be interesting, but not so high as to look bogus. I'm not accusing anyone of fudging the numbers, and I'm sure we all have the best of intentions, but it's remarkable how…consistent…the ROIs I see are. This is probably why KM and HR thought leader Hubert St. Onge describes ROI analyses with a single, earthy, Anglo-Saxon word.

3. At any given time, there are probably 50 initiatives that have positive near-term ROI. But the organization only has the ability to absorb so much change in the near term. So, which of these 50 initiatives will it implement? I guarantee you that the VP of Support and CFO will **not** rank order the candidates on the basis of calculated ROI. They'll find the financially attractive opportunities that best line up with the organization's strategic goals and initiatives. In executive decision-making, estimated hard dollars are less important than strategic alignment, perhaps because no one really trusts the hard dollar estimates, and perhaps because the strategic objectives are seen as the big opportunity for making money

The bottom line is, do an ROI to sanity check your initiative, but spend more time on the strategic framework—how you explain how the initiative, process and technology together, will move the needle for the organization's strategic goals.

Perhaps the saddest fact about ROI analyses is that, once the program is approved, they're generally filed away and forgotten about. This is a missed opportunity! The ROI analysis contains specific

assumptions and models about how the business will improve. Once the program goes live, you have a wonderful opportunity to test those assumptions in the real world. Are things working as you expected? Great! Now figure out what to do to make it even better. Are things falling short of expectations? Great! Your ROI gives you a model that will help you troubleshoot and make incremental improvements, or perhaps invalidate certain assumptions and help you pivot to a better approach. Either way, facts are friends: by overlaying reality on your ROI, you'll learn lots about how your business really works, and how to make things even better.

Metrics Myth Seven: Find Industry Peers and Benchmark Against Them

As a consultant, I find it rare to spend fifteen minutes with a capable executive before she asks me to assess how they're doing compared with other companies. It's a fair question: how can you know if you're best in class unless you know how well other organizations are doing? And if you're not yet best in class, how can you know how to get there without data?

Benchmarking works best on truly objective outcome goals: revenue per headcount or per square foot of retail space. Gross margin. Hotel occupancy. Share of wallet.

Unfortunately, support organizations own few of these. Arguably, for some organizations, maintenance renewal rates are a clear outcome, although even then, this depends on many cross-functional groups.

No, support leaders tend to want to benchmark other, more operational metrics: first contact resolution (FCR), customer satisfaction (CSAT), website "did this help" ratings on knowledge base articles, contact avoidance, and the like. And here's where cross-company comparisons can confuse as much as enlighten.

First, on a very practical level, everyone measures these things differently. You might think that there would be an industry-standard way of measuring FCR…but you'd be wrong. (I was brutally disabused of this notion during a—I am not making this up—two hour conversation among ten service desk managers comparing and contrasting the details of their FCR measurement techniques. These were among the longest, yet most enlightening, two hours of my life.)

CSAT, in particular, sounds like it should be straightforward to compare between companies, and yet it isn't. Survey results are notoriously sensitive to all kinds of details: how the question is worded and what scale is used for the answer can both substantially change the results. More subtle forms of bias can creep in, too. In our business, the most common is non-response bias—in other words, the small percentage of people who answer our surveys aren't necessarily representative of all the people we want to reach. Also, some companies exclude certain cases or customers from surveys, while others mix up transactional (post-case) surveys with relationship (periodic, overall experience) surveys. Comparing your "4.3 CSAT" against my "4.1 CSAT" really doesn't necessarily teach me anything unless I really, really understand the details of how you measure CSAT. If I'm simply using data from an industry benchmarking database, I have absolutely no idea how that particular brand of sausage was made.

If you have a deep and trusted relationship with a peer company, and you really understand how their measurement techniques match to yours, you might learn something from benchmarking. Unfortunately, ethical access to the fine details of peer company operations is rare.

More strategically, measures reflect your philosophy of business. Zappos is well known for not measuring handle time, because they believe personal relationships in customer service fuel their corporate engine of growth. So, what point is there in measuring your handle time against Zappos's?

That's an extreme example, but different business decisions drive different metrics, and obscure insights that might be gained by comparison. I know some support leaders are fixated on time to first technical contact, others on contact rate, others on service levels, and still others on NPS. None of these is wrong, but they will affect the measures.

So, benchmarking against other companies is fraught with difficulties. Still, we need aspirational goals. What to do?

First, sometimes benchmarks can be good to cause a deep rethinking of process. For example, I worked once with a company that was very proud that they had engaged additional outsourced staff in order to move their time to publication for knowledgebase articles from 23 to 14 days. When they heard that some companies do it in minutes, that data point spurred useful conversation. (Clearly, this requires more than hiring some new reviewers.)

Secondly, the best benchmarks are self-benchmarks. When I get on the scale, I don't compare myself with the average male my height; I'm interested in where I am relative to where I was the week before. Similarly, if I've improved my ability to capture, reuse, and improve knowledge, I'd like to see if my contact rate is going down, or my ability to close cases faster with reused knowledge is increasing. Any metric I report on is a candidate for self-benchmarking, especially if I've been taking actions to change something. Not every change is meaningful, but if you look at your own trends over time, in measures that are strategic to your business, you'll have the best possible, least-gameable, most informative metrics of all.

Metrics Myth Eight: Hold a Weekly Ops Review Where Changes are Scrutinized and Punishments are Meted Out

Operations Reviews, or ops reviews, are common management tools in Support organizations. If you're going to measure something, goes the logic, it's best to act on those measurements. Making course corrections sooner means that things get fixed sooner, with less customer impact. Bad habits don't have time to take root. Quick, decisive action at a weekly ops review can keep the ship on an even keel.

As with most of these myths, there's a kernel of truth. It's important to stay ahead of problems. The problem is a matter of execution. And, in fact, human nature makes it very, very hard for weekly ops reviews to be conducted productively.

The discipline of behavioral economics has taught us a great deal about the biases that humans exhibit. Unlike the rationally maximizing human beings who play the starring role in most economics textbooks, real humans make decisions using many mental shortcuts—shortcuts that provided evolutionarily advantage. If your ancestors had performed a careful, logical analysis when saber-tooth tigers prepared to pounce on them, you wouldn't be here. But because we unwittingly fall back on these mental shortcuts, we're not always as rational as we might be. For a comprehensive and wonderful treatment of this, read Thinking, Fast and Slow by Daniel Kahneman, cofounder of Behavioral Economics; for a faster, thoroughly enjoyable version, try Predictably Irrational by Dan Ariely.)

What do saber-tooth tigers and behavioral economics have to do with weekly ops reviews? Well, when humans are confronted with data, they can't avoid finding patterns and explanations—even if variations in the data are 100% consistent with chance. (One last book recommendation: The Drunkard's Walk by Leonard Mlodinow lays out many examples where even mathematically sophisticated people can't help but believe in explanations that don't exist.)

This bias towards pattern recognition turns ops reviews into the Theater of the Absurd. We have all this data, so we need to perform analysis and take actions! Page views increase, and some minor change in the website UI is given credit. CSAT goes down, and managers are called on the carpet and demanded to explain themselves. Culprits are identified and chastised. Yet, rarely is there someone in the room with the statistical acumen, the fortitude, and the mandate to say, "hey, but maybe that's just a normal variation."

It's not just any patterns we find, either. "Confirmation bias" means we find the patterns we're looking for. (If you need proof of this, just listen to two people at opposite ends of the political spectrum explain how a particular news report supports their respective positions.) If I change the website's color, and knowledgebase page views go up the following week, is there any doubt in my mind where the credit lies? This doesn't make me a bad human; it just makes me human.

Also, a statistical curiosity called regression to the mean is responsible for reinforcing mistaken notions of cause and effect in ops reviews. If my metrics are atypically bad, by chance, they're likely to be more typical—in other words, better—next week. This is true whether or not I get yelled at…but it certainly makes yelling at me look like a highly effective strategy!

I know of some businesses where the weekly ops review is anticipated much as a visit to the dentist: you can be pretty sure there will be pain somewhere in the office, and you mostly hope it's inflicted on someone else.

So, what's the alternative?

- In a normal environment, there's no reason to have a full-cast review of operational metrics more than once a month. This at least reduces the amount of metrics "jitter" and makes the signal stronger relative to the noise. Two caveats: Note that individual operations analysts and managers may scan the metrics more frequently to look for large unexpected movements that may require interventions. And during periods of rapid change, such as a short intense busy season or while rolling out new initiatives, more frequent focused meetings are appropriate.

- Be aware of the natural tendency to assign cause where a cause may not exist. Be especially aware of conclusions that confirm your own instincts, or that explain a return to normalcy. Confirmation bias and regression to the mean may be in play. It's not a bad idea to find someone who can calculate confidence intervals on your data: you may be surprised how little of the variation that your business experiences is statistically significant.

- Finally, remember that the purpose of metrics is to inform the conversation and to aid learning… not to reward and punish. As Dean Spitzer points out, the effectiveness of metrics as learning tools is directly related to the health of the culture of measurement. And nothing damages the culture of measurement more quickly than beating people up based on their numbers.

"The beatings will continue until morale improves" is a funny t-shirt, but it's no fun in real life. Keep your ops reviews grounded in reality, positive, and informative.

Metrics Myth Nine: NPS is Whatever I Say It Is

I'll just say it right now: I love the Net Promoter Score. It may not be perfect, but in my experience, it's the most powerful tool that Service and Support executives have to demonstrate their value to the entire enterprise, become more strategic in how they think about their organization's mission, and focus the entire company on the customer and their experience. The NPS is an All Access Pass to the boardroom.

Here's a brief definition for those unfamiliar with Net Promoter. Industry research led by Fred Reichheld and his team at Bain & Company determined that customers fall into one of three buckets: they're promoters who advocate for your brand, detractors who will badmouth you in the marketplace, or neutrals—people who may be satisfied, but who don't rise to the level of advocate. If we think of the companies we're customers of, we can probably easily slot ourselves into one of those three buckets relative to them.

After much quantitative analysis, Reichheld and team determined that the most accurate way of predicting whether a customer is a promoter, detractor, or neutral is to ask them how likely they would be to recommend the company, on a zero to ten scale. Nines and tens are promoters; sixes and below get classified as detractors; and sevens and eights are neutrals. If you survey your customers, take the percentage of promoters and subtract the percentage of detractors, you get your Net Promoter Score—that is, the number of promoters net of the number of detractors. Neutrals don't count for or against you. A company that scores all 9s and 10s gets a 100%; a company with all low scores is -100%. Bain's research suggests that most companies are in the low positive territory, while leaders can be plus 50% or higher. (If you're in negative territory…sell your stock.)

Unfortunately, because of its power, NPS has also gotten trendy. And trendy measures tend to be misused, or to put it more nicely, "reinterpreted." It's these creative interpretations that are the NPS myth. Here are some examples I've seen:

Redefining the Scale
It's a small tactical issue, but Net Promoter is based on the zero to ten scale with break points between six and seven, and eight and nine. I see companies talk about the "NPS," but pick entirely different scales or endpoints. There can be good reasons for picking a different scale, but you can't call what you calculate from it Net Promoter Score. It's something different, and the research on NPS may not apply directly, and comparisons with actual NPS scores aren't valid.

Measuring NPS for Support
The Net Promoter Score is designed to "measure how well an organization treats the people whose lives it affects—how well it generates relationships worthy of loyalty." What we're measuring is the customer's impression. And customers don't care about your internal divisions. They don't have one impression of support, another of product development, another of your CFO, and yet another of HR. By and large, they think about the company overall, and they'll certainly recommend the company overall (or not.) They can't say, "Well, buy your accounting software from Company X, but use Company Y's support organization."

So why would we measure the Net Promoter Score for Support? "How likely are you to recommend Company X's Support Organization to a Professional Colleague?" is an unhelpful question. It doesn't work that way. You either recommend Company X, or you don't.

As we mentioned earlier, it's understandable to want to try to measure things that are only under your control, but that's not how customers experience things, so it's not a very productive exercise.

Measuring NPS Transactionally
Even more extreme than measuring Net Promoter for Support overall, is measuring Net Promoter for a single support transaction. It's a great thing to send (very brief, recurrence-limited) surveys to people who have engaged with your support staff or online resources—but let's not pretend that their

willingness to promote your brand is primarily based on their last support interaction. Transactional and relationship surveys are different: in transactional surveys, you ask about the transaction, and only in the relationship surveys should you ask about overall impressions.

In transactional surveys, stick to asking about the case, or the web session, and leave it at that.

Compounding this error, I'm sorry to say that I've seen companies that selectively measure transactions, exempting entire classes of cases from surveys and giving individuals the option to strike particular troublesome interactions from surveys as well. Needless to say, there's enough sample bias and nonresponse bias in most survey programs that adding more with malice aforethought is really a terrible idea, and ultimately self-defeating. This almost always is a result of focusing too much on the numbers.

Focusing on the Numbers

The root cause for many of these reinterpretations is a mistaken belief that the NPS numbers you report, by themselves, mean anything.

NPS, as you'll recall, was developed as a way of measuring loyalty. That is, higher NPS scores mapped into higher customer retention, lifetime value, and referenceability. Since those are future behaviors, they can't be measured directly, so NPS at its best is a predictor of good things happening in the future.

There is no NPS Bowl where companies or executives get rewarded because their numbers are high. Whatever credit an executive hopes to get when he presents his "82% NPS" falls apart as soon as a skeptical audience member smells a rat and starts asking questions—at which point, credibility is irretrievably lost. And, more importantly, although independent research has been mixed in supporting the predictive power of NPS, at least there is some. Gamed numbers (based on modified scales, limited to support transactions, selectively sampled) have no research supporting their predictive power.

Fred Reichheld himself has said that the important thing about NPS isn't what the number is; it's what you do about it. In the case of Support leadership, it can help line up the entire company on Support's mission of the customer experience and customer success. This is a Good Thing. If we instead create unreliable, Support-only numbers to try to win a nonexistent NPS Bowl, we've done our customers, our companies, and ourselves a disservice.

Metrics Myth Ten: Contact Deflection = Clickthroughs Times the X-Factor

Industry surveys ask. Colleagues ask. IT asks. Executives ask, incessantly.

"What's your contact deflection rate?"

Let's leave aside for the moment that contact deflection isn't the primary reason to expose your knowledgebase to your customers. Let's also leave aside the fact that it's both practically and philosophically difficult to count things that **didn't** happen. (Do I get credit for all the times I wasn't snippy to my family and colleagues? No? But I counted 39 instances of snip-avoidance last week! I mean, if I hadn't been so nice, I would have been a real jerk.)

So, like it or not, we have to come up with an answer for how many contacts we've avoided. As an aside, I don't recommend that you characterize this as a rate—that is, a proportion of something else. People ask for this, but in my experience most don't really know how to think about it that way, and there also isn't universal agreement about what to divide it **by.** Saying "We had 10,000 self-service sessions this week that deflected approximately 500 cases" is clear enough; saying "we had 68% contact deflection" seems less obvious.

I've written elsewhere about techniques for estimating contact deflection; if you're interested, please look at this, and I won't wave mathematical formulas at everyone else reading this e-book.

What I do want to say is, it's not straightforward. And I cringe when I hear that there's some magical "X Factor" that you can multiply by page views or web sessions to calculate deflection.

For one thing, contact deflection is lower than many people expect, because the majority of customers using self-service, even entitled enterprise customers, wouldn't have opened a case. As charming as we all are, customers would usually rather not talk with Support if they can avoid it.

But whatever your customer's propensity to open a case after an unsuccessful self-service interaction (or page view, or Google search, or in-product link, or…what were we multiplying again?), you have to acknowledge that it's going to vary from customer segment to customer segment, and product to product, and perhaps by geography, even within your own company. Certainly, your customer behavior around opening cases after self-service will be different from that of other companies, even in your industry.

So the idea that there's a single multiplicand that can give you self-service contact deflection is, unfortunately, just another myth. If you must have an answer, calculate escalation rates and success rates as best you can, and let everyone know you're estimating. Please don't say that you've multiplied your web traffic by an "industry best practice" constant.

Metrics Myth Eleven: Any Initiative for Which You Can't Define Clear Outcomes Isn't Worth Doing

This is another inference from "you get what you measure." If you don't know how to measure an initiative, you don't know what you're trying to achieve. And if you don't know why you're undertaking an initiative, you shouldn't do it. Fair enough.

But isn't it sometimes the case that we have an intuition that something will be valuable, even though we're not quite sure how to measure it?

Collaboration springs to mind. I've heard many support executives say things like, "if we could just get the right person with the right knowledge on the right case the first time, we'd be in much better shape." And I believe that, too. But trying to construct a measurement framework to capture the health and value of collaboration is really hard. ("Ask me how I know.")

In general, if what we're trying to do is very innovative, we probably won't be sure how to measure it, at least at first. And that's probably OK.

But does that mean that people running innovative initiatives have no accountability? Absolutely not. And to discover how, we'll look at another place where disruptive innovation is the norm: the startup.

The Lean Startup by Eric Ries is the most influential book of the decade in Silicon Valley. It's hard to imagine a book since Crossing the Chasm that's had more of an effect on how people think and talk about startup strategy. (Have you noticed people saying "pivot" all the time recently? Blame Ries.) Measuring innovative businesses in the midst of uncertainty is central to The Lean Startup's message.

Ries argues that, in an innovative or uncertain environment, the goal should be validated learning—learning how to create sustainable value. He contrasts the purposeful acquisition of validated learning through a build/measure/learn loop in opposition to the kind of "learning" (in air quotes) that is the consolation prize from a failed initiative. "Sure, it didn't work like we expected, but we've learned some lessons for next time" isn't validated learning. Rather, validated learning uses the scientific method and a rapid cycle of experiments to propose, validate, or discard hypotheses, then move on to the next experiment.

Validated learning happens inside a framework of what Ries calls "innovation accounting" that holds innovators accountable for learning and improving. And, if sufficient improvement isn't forthcoming, to rethink their assumptions and pivot to another innovation, until the value they're expecting is forthcoming.

So, for really new initiatives, hold yourself accountable for validated learning. For example, in collaboration, enumerate your key success factors: people need to use it frequently, and the number of users should grow over time, for starters. Release a minimum viable collaboration program and see if people use it, and see if demand grows. If it does, improve it. If it doesn't, figure out what to change and test that. Eventually, you'll have a program that's viable, or you'll know you were on the wrong track and need to pivot to a new one.

By the time you've acquired sufficient validated learning, and you have a functioning program, then you're in a position to figure out how to measure the value. It's easy to measure the value of a program that isn't working—it's zero. But measuring the value of an effective program…now that's interesting, and very worthwhile.

What Do We Do About All Of These Myths?

Goodness me, listen to me go on about these metric myths. If you've made it this far, you can be forgiven for asking, "That's all fine, but what do **I** do about it?" Metrics exist in a context, and that context isn't going to change itself just because some consultant writes a few blog posts. Or even a lot of blog posts.

Here's some guidance for taking on your culture of measurement.

Address Technical Mistakes

Because measurement and leadership are so closely entangled, it's hard to critique measures without also appearing to critique their owner. But, a relatively safe critique is one that is purely technical. For example, if people are miscalculating self-service metrics or leaving people and process costs out of an ROI, it's pretty nonthreatening to offer to tweak the measure a bit…especially if it's new and hasn't been publicized widely in the past.

Similarly, sometimes just renaming measures can help. For example, leaders might be reluctant to stop reporting a "Net Promoter Score" that's support only, transactional, and calculated with bogus cutoffs but perhaps you can have it renamed to a "Support Recommender Index" or some such. This gives people the metric they want without the possibility of misleading others and losing credibility.

Ask Good Questions

It's always easier to change attitudes with questions rather than opinions. (Even **I** won't walk into your offices and start telling people they're doing metrics wrong. Probably.)

Most people at work are smart and well intentioned, and as we've seen, there's at least a grain of truth in each of these myths. With careful and appreciative inquiry, we can perhaps get people to move away from the unhelpful parts of their measurement strategies while holding on to the good core.

For example, if I see an ops meeting that attendees despise, I might ask,

- What are the most useful insights you've gotten from these reviews this year?
- What were some of the most effective corrective actions you've taken, and how long do those actions take to implement?
- How do you assess what changes in the data are meaningful, and which are just usual variations?
- What do you think is the normal variation in the data? Are there patterns around seasons or new releases?

If I see only activity-based measures (SMART goals) on review forms, I might ask,

- Who are your superstars here? Why? How does that show up in these goals?
- Have you ever seen people attempt to game these metrics?
- Have you had employees score well on these measures, despite impressions or feedback that they're not the best performers?
    - What guidance has HR given you on the nature of your goals?
    - How do your knowledge-working counterparts in development and marketing measure their staff?

I hope answering these questions might cause some introspection and different ways of thinking about a desirable measurement strategy. At the very least, having heard the answers, I have more ways of framing what I'm going to recommend in a way that's responsive to my colleague's challenges, values and interests.

Implement Better Alternatives

Living well is the best revenge…and it's also a great way to get people to follow you! Generating and analyzing metrics may not be in your formal job description, but everyone needs to measure their work at some level, so consider yourself officially empowered to be master of your own metrics.

Of course, if you manage a team or program, you have an even clearer mandate, and more opportunity to make a bigger difference.

For example, if you're part of a team taking on a new initiative, why not experiment with some Innovation Accounting, perhaps tracking progress against your original ROI model? Or, you could implement a self-benchmark? However you choose to measure yourself, make sure that people will know what you're learning and what you're doing as a result of your measures. And see if the things that you're doing don't start to spread across the organization.

One additional note on metrics innovation: the people you report into may not be ready to let go of their old measures, at least until they have significant experience and comfort with your new ones. So it's useful to "keep two sets of books," as we've discussed before—the traditional numbers to report up, and your new metrics to learn from and to start to expose alongside the traditional numbers. It may seem like extra work, but it's hard for people to lose metrics they're used to, so you'll probably have to operate in parallel for a while.

Use a Strategic Framework

Measures are in service of a bigger goal or measurable vision.The document that connects the operational metrics to the bigger goal is a strategic framework.

The strategic framework is a great way to communicate the "why" of your measures. Because each measure supports an action that supports a strategic objective, there's a clear reason for why we're looking at what we're looking at.

The Strategic Framework is simple—so simple, it's tempting to skip. Please don't. Your metrics, your executives, and the rest of your colleagues will thank you.

Outlast 'em

Sometimes, it's just not possible to change leadership's measurement practices, especially if they've been doing things the same way for a very long time. In this case, just wait for them to move on. It'll happen eventually—sometimes, gratifyingly quickly. If you see dysfunction, you're likely not the only one.

It's a satisfying feeling, especially if you've been planting seeds with your own more effective measurement strategies. Keep the faith!

DB Kay provides:

Strategic Consulting in Metrics
Assessments
Training
Technology selection
Adoption consulting services
for market-leading service and support organizations

DB Kay's customers include IBM, Farmers Insurance,
Extreme Networks, NetApp, Veritas and Intuit.

David Kay, principal of DB Kay & Associates, has been a leader in applying technology to knowledge-intensive business processes like customer service since 1984. He is a frequent industry speaker, and has been recognized as a customer service Innovator by the Consortium for Service Innovation.

David Kay holds five patents covering the use of next-generation technology in customer service and support.
He is co-author of *Collective Wisdom: Transforming Support with Knowledge,*
available on Amazon.com