

Introduction

THE END OF THE INFORMATION REVOLUTION

There was no fog of war over Basra on the seventh day of the invasion of Iraq. Instead of operating in an information vacuum like U.S. Civil War soldiers firing into a smoke-filled field, the American pilots of two A-10 Thunderbolts flying a mission to destroy Iraqi rocket launchers on March 28, 2003, were bathed in a sea of information as dense as the Mesopotamian sunlight. Supporting them were at least three Marine Corps forward air controllers attached to British units on the ground, an AWACS control ship called Sky Chief, a nearby British pilot, and a theater commander called Twin Act.¹

“Hey, I got a four ship,” called POPOV36, one of the A-10 pilots, spotting a suspicious convoy below. “Looks like we got orange panels on them though. Do we have any friendlies up in this area?” Coalition vehicles bore orange panels as markers.

A forward air controller called Manila Hotel checked that the pilot was eight hundred meters north of Basra. “That is an affirm,” he replied, affirming a probable target. “You are well clear of friendlies.”

“I see multiple riveted vehicles,” interjected POPOV35, the other A-10 pilot. “Some look like flatbed trucks and others are green vehicles. Can’t quite make out the type. Look like they may be ZIL157s.” POPOV36 repeated his sighting of a convoy of four vehicles, the “four ship,” for POPOV35’s benefit, adding that the vehicles were evenly spaced along a road. POPOV35 replied, “I don’t have your visual.”

POPOV36 provided more detail to help his wingman find the convoy: “Right underneath you. Right now, there’s a canal that

runs north/south. There's a small village, and there are vehicles that are spaced evenly there. They look like they have orange panels on though."

POPOV35 reminded him that Manila Hotel said there were no friendly vehicles in the area. "They've got something orange on top of them," pressed POPOV36. "Hey, tell me what type of rocket launchers you got up here?" he asked the air controller. "I think they're rocket launchers." And he added to POPOV35: "Roll up your right wing and look right underneath you."

POPOV35 replied, "I know what you're talking about."

"Okay, well they got orange rockets on them," said POPOV36.

"Orange rockets?" asked POPOV35.

"Yeah, I think so," answered POPOV36. Then he added, "I think killing these damn rocket launchers, it would be great." A minute later he asked POPOV35, "Okay, do you see the orange things on top of them?"

"I'm coming off west," POPOV35 replied. "You roll in. It looks like they are exactly what we're talking about."

"We got visual," said POPOV36. "Okay. I want to get that first one before he gets into town then."

"Get him—get him," urged POPOV35.

"All right, we got rocket launchers, it looks like," said POPOV36.

POPOV36 rolled his A-10 into a vertical dive to strafe what would prove to be a British column, hitting two Scimitar armored vehicles.

This book is about telling in an age of information overload what's relevant to what you're trying to do. The trouble with an overload of information isn't just that it's confusing. It's that the data have conflicting implications. There are no friendly vehicles in the area. But the vehicles below bear something like the coalition's orange marker panels. They could be rocket launchers. But orange is an odd color for a launcher. Our enormous data

resources are doing us no favors when they paralyze our decision making. This book helps find what information is most relevant despite the overload.

The transcript from POPOV36's cockpit tape goes on to etch in blood and tears just how futile a surfeit of information can be. Seven seconds after POPOV36 fired again, a second forward air controller, called Lightning34, warned, "Be advised that in the 3122 and 3222 group box you have friendly armor in the area. Yellow, small armored tanks. Just be advised."

A few moments later, a third forward air controller, called Manila 34, ordered: "Hey, POPOV34 [sic], abort your mission. You got a, looks like we might have a blue on blue situation." POPOV35 swore and asked him to confirm they were over friendly vehicles. "Okay, POPOV," called Manila 34. "Just west of the 3-4 easting. On the berm up there, the 3422 area is where we have our friendlies, over."

POPOV35 asked how the friendlies were. Manila 34 told them they were getting an initial brief that one was killed and one wounded. "I'm going to be sick," said POPOV35. A few moments later, he made the comment to his fellow pilot that lit up newspaper accounts of the incident: "We're in jail dude."

"Relaying for TWINACT," broadcast Sky Chief, "the A-10s are running against friendlies."

A British pilot repeated: "POPOV35, this is COSTA58. Relaying message for TWINACT. Abort. Abort."

Of course, the two pilots had already aborted their mission. Their curses ran from denial to anger to grief. "They did say there were no friendlies," moaned POPOV35.

And as Lance Corporal of Horse Matt Hull, a young Englishman from Windsor, Berkshire, just short of his twenty-sixth birthday, died below, POPOV36 wept.

Dimensions of an Information Crisis

The information revolution is drawing to a close as our expanding information resources bury decisions under ever higher piles of conflicting data. The next section blames this on our neglect of

relevance as raw data get cheaper to generate. The third section asks how we can learn from experience despite the glut of information and explains why I want to address this challenge by focusing specifically on performance experience—on the results of our projects, businesses, and ventures. The fourth section summarizes this book's answer to the challenge of learning from the reams of performance data our organizations collect. It calls for the development of performance strategies explicit enough to be testable, the derivation of performance metrics from the assumptions behind those strategies, and the use of performance results to reveal errors in our strategic goals and assumptions, as well as execution. This section outlines the dimensions of the problem.

The sources of evidence for an information crisis provide a decent list of people who should consider reading this book. One group it excludes would be all of those whose e-mail inboxes, performance scorecards, operating dashboards, and management reports are growing shorter. If your organization is asking you to cope with less information, these pages may not speak to you. Those feeling oppressed by operating data should read on.

Small business owners and general managers of the divisions of larger firms may suffer most. They often feel they have to monitor everything that could possibly go wrong in their organizations. If that's not bad enough, the general managers also have to watch the metrics their CEOs monitor—even if those metrics are irrelevant to their plans. This book proposes a three-part program to determine which metrics really matter.

It confronts the core challenge before finance professionals—including chief financial officers, controllers, divisional finance directors, audit directors, and financial analysts—of reporting financial information to improve performance without overwhelming their organization. The book's tools will also help human resource executives design performance goals to develop high-potential staff. And it addresses the increasingly intense performance assessment demands on information technology groups reflected in the proliferation of project management offices.

The book's answers to the learning challenge posed by exponential information growth are directly relevant to military organizations as they struggle to simplify the information cascading down on combat theater teams and rethink that information in terms of conflict between opposing networks—so-called net-centric warfare. Nondefense government officials are in equal need of ways to focus on what, of all the information at their disposal, really matters to policy plans and proposals.

Ironically it may be nongovernmental organizations that have undergone the biggest change in their use of performance data over the past five years. Impact assessment groups have stepped up efforts at the World Bank, the International Monetary Fund, and their sister institutions under the International Bank for Reconstruction and Development. Charities like the Gates Foundation, moreover, are applying fast-to-fail performance criteria to their initiatives to promote learning as well as husband resources. The growing focus on performance evaluation in a context of complex and often conflicting objectives is sharpening these organizations' appetite for a robust framework like this book's approach to strategic goals, assumptions, and success indicators.

The most dramatic development in performance assessment over the past decade, however, has been the effort to integrate functional measures for marketing, research and development, procurement, supply chain planning, process improvement, quality control, sales, and customer service into a coherent view of corporate effectiveness. Often taking advantage of enterprise resource planning systems, this effort has evolved in two directions. Numerically intensive continuous process improvement programs, usually found at larger manufacturers, tend to require heavy levels of up-front investment. Qualitative approaches like balanced scorecards are a little less expensive but often grow complex without measurable benefits. This book proposes an approach to defining and measuring performance that applies the discipline of continuous process improvement to strategy evaluation and development without disrupting existing systems or processes.

One could easily go beyond these examples of information overloads in organizations to list the kinds of financial, consumer, and legal information that bombard us in our personal lives and complicate our decisions about careers, spending, investment, households, and even networks of friends and relatives. The problem is familiar and pervasive. In fact, it's woven into the texture of contemporary life tightly enough to mark a period in the very information revolution that gave rise to it.

Relevance and the End of the Information Revolution

It's one thing to say our information resources are burying decisions under ever-higher piles of conflicting data, but it's quite another to declare an end to the information revolution. Even so, I think our neglect of relevance in information—encouraged by the availability of raw data that keep getting cheaper to generate—has drawn a curtain on it.

A major idea in this book is that the usefulness of any piece of information depends on two complementary values: its specificity—which you can think of as its content, or simply how surprising it is—and its relevance. Specificity gives you a sense of the power of a piece of information because you can draw more conclusions from precise outcomes than from vague ones. A result reflecting one out of ten possibilities, for example, tells you more than a result reflecting one out of just two.

Relevance complements specificity. While specificity is intrinsic to information—it's the same no matter what your goal is—relevance depends entirely on that goal, or at least on the assumptions behind it. As I use the term in this book, *relevance* has to do with how well a piece of information tests your expectations.

It's fairly clear that both specificity and relevance are important in evaluating information. The startling part of the idea is that they complement each other—they're all that matters. It's startling but plausible. To see how specificity and relevance

fit together, imagine you're throwing paintballs at a target. Specificity is like the size of the paintballs. Big ones resemble unspecific results; they don't gauge your accuracy. Relevance is more like their shape. Round ones that stay intact until they hit the target—like relevant information—really test your aim. Those that scatter or leave thin, oblong marks do not. To improve your aim, you need balls that leave precise round marks—information that's both specific and relevant.

The idea that specificity and relevance drive the value of information helps explain how the information revolution's economics may have undermined it. Technologies like microprocessors and optical fiber have basically cut the cost of specificity. Whether you run a division, compile financial reports, or manage service operations, you've been able to get more data faster. The relevance of those data has mattered less and less because increasingly you've been able to compensate with volume. And as we've grown used to less relevant data, I suspect that we've formed two bad habits.

The first bad habit is *requirements-based analysis*. Instead of devising a specific strategy to meet a financial goal like a sales or profit target, we're increasingly tempted to enumerate requirements for meeting the goal. Requirements might include minimum levels of staffing, funding, technology functionality, or process performance. Budget planning increasingly follows this course, working backward from profit targets to sales requirements and cost ceilings. And executives tend to use balanced scorecards, one of the most widespread management tools, as menus of common requirements.²

In this book, I liken these necessary conditions for meeting a goal to lists of ingredients. The comparison helps explain the appeal of requirements; it's as easy to find metrics for them as it is to determine whether you have all the required ingredients for the dish you're cooking for dinner. But just as you can know all the ingredients for a dish without having a recipe for it, you can meet every conceivable requirement for a goal without knowing how to achieve it. Simply put, requirements are not strategies. And unlike strategies, there's

little to learn when they fail to work. That may be another key to their appeal: you can be right about requirements even if a goal proves elusive. And it leads to the second bad habit: our growing reliance on *red herrings*.

Red herrings are results that appear to confirm your plans but in reality are merely consistent with them. You might think customers will prefer cell phones with Internet browsers, for example, and find they purchase more of a model that includes them. But if that model is also much easier to use than its predecessors, the increase in customer uptake may be a red herring as far as your assumption about browsers is concerned. The new model's results, in other words, are not necessarily relevant to expectations about cell phone browsers.

As our enterprise planning systems accumulate more and more data of limited relevance to our plans, red herrings are bound to proliferate. That's fine as long as you recognize them. The trouble is that you may think they confirm your strategy. Worse, you may be tempted to build a strategy largely out of requirements and rely on irrelevant metrics—on red herrings—to test its more explicit prescriptions. Experience won't teach you much about such a strategy: events are unlikely to prove the requirements wrong, and the red herrings are unlikely to challenge the explicit prescriptions.

Our declining ability to learn from experience as we accumulate conflicting data marks the end of the information revolution. Hastening its close is a characteristic style of performance analysis that I call the *analyze-execute system*. It has three tenets: you can reduce any account of performance to a set of basic facts, those facts imply a correct strategy, and so finding the right strategy reduces to collecting enough facts. These tenets contradict the way we formulate and test hypotheses in our conduct of science—scientists rarely just collect facts—and yet we embrace them in business and government.

Why do we do that? Perhaps it's not just that we want to be right—which is less likely if you lay out explicit strategies that events might prove wrong and track highly relevant metrics

that will tell you as soon as they do. Perhaps it's that we hope we've hit on an objective way of looking at our businesses and the world around us that captures the way they really are. We distrust the filter of our point of view.

But the question remains: Why do we want to be right in such a metaphysical way rather than just achieving a goal? The philosopher Richard Rorty offered the controversial explanation that we pursue a kind of objectivity beyond mere public evaluation because we fear we'll be forgotten. He wrote, "The picture of a common human nature oriented towards correspondence to reality as it is in itself comforts us with the thought that even if our civilization is destroyed, even if all memory of our political or intellectual or artistic community is erased, the race is fated to recapture the virtues and the insights and the achievements which were the glory of that community."³

Rorty may have put his finger on a profound reason that we're not content with making up business strategies and testing them for errors. Perhaps we really do crave insights and business strategies that not only work but are somehow true for all time and in all languages. There's a big practical advantage, however, to my proposal that cheap information has tempted us to neglect relevance and led us into some bad habits. It means you may not have to develop a view on the fate of the world to learn from experience despite the information overload our organizations produce. Paying close attention to relevance may be enough.

Why Focus on Performance Management?

If it's true that the information revolution has generated piles of data so complex we can hardly tell what matters, how are we to learn from experience going forward? Of course, broad questions like this invite answers so broad that they never get around to concrete suggestions. But a narrower question could fail to do justice to all the ways information overloads hamper decisions we need to make in everyday life. For this book, I thought the best bet

was to investigate a specific problem that raises the full spectrum of challenges posed by information gluts and try to provide an explicit solution.

It's probably no surprise that a business writer would identify a core problem in business learning—selecting the best performance indicators to accelerate it and shape our organizational strategies—as the most revealing way to address the phenomenon of information overload. I suspect many business readers will grant me the benefit of the doubt on this. Here I'll try to persuade the nonbusiness reader that the challenge of learning from an organization's performance experience uniquely illuminates the broader question of what really matters in the sea of information around us.

One reason is that organizations—and businesses in particular—are awash in a rich mix of information. Businesses were the earliest investors in information technology so they can gather a lot of data. They're under competitive pressure to figure out what works, so they collect information on factors like work effort and resources, as well as on results like profit. And they collect it frequently to speed up decisions. As a result, the performance reports businesses produce are especially rich in data with conflicting implications. That makes them a good model for the mixed signals that create so much confusion in our increasingly information-intensive daily lives.

The other reason is that while any solution to the information overload in business organizations may be more structured than we need for solving the kinds of personal problems that arise in career choice and managing busy households, that very structure should clarify principles for simplifying our lives. For example, the testable strategies, assumption-based indicators, and critical strategy reviews that I propose in this book also serve as models for precision in our personal expectations, milestones for testing those expectations, and realism about changing them in the light of experience. The principles for determining what matters in the performance experience of a business are clear and structured versions of those that can determine what matters in our own experience.

Before going on to a summary of the book's solution to the problem of organizational learning in the face of overwhelming and conflicting performance data, here are twelve questions that it should help you answer. The Conclusion takes a stab at summarizing the answers developed across its chapters:

1. How is organizational learning related to growth?
2. What's wrong with balanced scorecards, and what's the alternative?
3. How should we set performance goals?
4. How do testable strategies help us learn from experience?
5. Which metrics matter?
6. How can we measure relevance?
7. What kinds of acquisitions are relevant to our business?
8. Are forecasts necessary?
9. Why do traditional performance reviews destroy morale?
10. What's the relation between performance volatility and compensation?
11. What does relevance have to do with leadership?
12. What's a relevance revolution?

Relearning Learning from Experience

This book argues that learning from experience—when you're confronting too much of it—requires three things: the development of performance strategies explicit enough to be testable, the derivation of performance indicators from the assumptions behind those strategies, and the use of performance results to reveal errors in goals and assumptions as well as execution. Chapters One and Two deal with testable strategies, Chapters Three and Four deal with assumption-based metrics, and Chapters Five and Six deal with strategy reviews.

Chapter One uses a cautionary tale from BP and examples from Alcoa and GE to argue that learning from experience requires performance strategies explicit enough to be testable. A strategy is testable if it spells out goals and assumptions about how to achieve them that could conceivably prove wrong. Chapter Two proposes *eight-line strategies* for distilling the strategy relevant to a manager at any level of an organization to a short list of testable assumptions. These strategies are basically devices to help you identify the key unsettled assumptions or biggest bets you're making when you project a result. Laying the foundation for vastly simplified management reporting and performance reviews, these assumptions let you pick out the facts most relevant to your strategy from large amounts of conflicting performance data.

To learn from experience despite the noise of that conflicting data, Chapter Three derives performance indicators from key strategic assumptions—and not from balanced lists of output targets and input requirements. It draws examples from Media General, Saatchi & Saatchi, and Ingersoll-Rand—three companies that participated in Kaplan's and Norton's Balanced Scorecard Collaborative—to show that however important it may be to balance short- and long-term concerns, the pursuit of balance in what we measure reflects confusion about why we measure it.

Chapter Four helps you select a handful of critical performance indicators by defining their relevance with respect to your assumptions about what will achieve a strategy's goals. A performance indicator is relevant to a strategic assumption, according to the definition, if the assumption's truth or falsity greatly affects the results you expect. This definition marries aspects of Bayesian probability, which tracks how new evidence affects beliefs, and information theory, which basically measures what we don't know, to avoid false conclusions and mimics the skeptical use of performance data at Toyota and Capital One.

Requiring only a pen and piece of paper, a tool called the *metrics matrix* uses the definition to determine which indicators best test each of your key assumptions from an eight-line strategy.

The handful of resulting indicators are the only ones you need to track until your strategy changes. Balanced scorecards, in contrast, rarely test strategies. They have a natural tendency to test something quite different: requirements for success. The problem is that you can meet a list of requirements and still miss the goal. Since there is no end to the requirements you can identify, balanced scorecards tend to accumulate more and more metrics without ever defining a testable strategy.

To learn from increasingly complex experience, finally, we must use performance results to reveal errors in our goals and assumptions, as well as execution. Splitting the difference between actual and expected performance results into execution, uncontrollable, and strategy gaps, Chapter Five proposes a model *strategy review* that explicitly challenges strategic plans and assumptions in place of traditional reviews of performance results that focus only on execution. Strategy reviews look at every performance period as a controlled experiment. These experiments naturally test work effort and risk forecasts, but they also test the quality of the assumptions relating that effort and those forecasts to results. The biggest benefit of strategy reviews is motivational because organizations that don't look systematically for strategy errors in their performance results force their operating managers to protect themselves from being blamed for missing unrealistic goals.

Our organizations won't sustain growth, Chapter Six argues further, unless we use the pattern of gaps in our performance results to review strategy and revise tactics continuously. Nestlé provides a powerful example of a performance system that can telescope the review of that pattern into a single performance period. The pattern compresses your organization's performance experience into three numbers. By telling you whether your assumptions about execution, risk, or strategy are most in need of repair, the pattern of gaps lets you allocate scarce talent to your most persistent problems. The result is a fully engaged management team that evolves strategy continuously by testing and revising assumptions at every level of the organization.

Taken together, the testable strategies, assumption-based metrics, and strategy reviews proposed here turn traditional planning and performance systems inside out, basing them on your biggest bets about how to achieve a goal. These bets are more like recipes than lists of ingredients. They specify how to achieve an objective, may be wrong, and require testing and revision. But they focus you on the performance results from which you can learn the most—about your current strategy.

Call this alternative performance system the *guess-test system*—in contrast to analyze-execute systems that try to derive strategy from the raw data of a comprehensive scorecard and test only whether execution is faithful to it. While guesses in the form of assumptions may seem like a strange place to start learning from experience, the discipline of testing strategy rapidly hones them to the realities of your business.

The Larger Relevance of Relevance

The *Washington Post*'s Steven Pearlstein once complained to me that books offering business advice too often claim their frameworks solve every problem known to humans. This is the section where I run the risk of doing exactly that. But before I do, here are a few thoughts on what this book is not and what I think books like it should try to accomplish.

It's not a book of academic case studies, even though it provides plenty of business examples, opinionated accounts of developments at well-known companies, and nonbusiness examples as well. I think books laying out a potentially controversial argument like this one make a mistake if they try to prove their point with detailed case studies. The selection of details can easily make a case appear to support an argument that it actually undermines. The examples here are illustrations, not proofs.

Business books advancing new or neglected claims also need to offer advice that every reader could, in principle, follow. If that advice is to collect the most data, for example, those who follow it

will conflict with one another and may be better off ignoring it. That's not a limitation of the advice on these pages because it shows you how to find information relevant to your own assumptions. If your assumptions are creative, the book's implications for you will be unique. This is not a zero-sum book.

Nor will the advice in this book restore organizational morale, separate good mergers from bad ones, calculate optimal compensation levels, solve complex environmental threats, or identify great leaders. And yet its themes touch on all of these things. It's worth mentioning the connections as a reminder that the deepest problems in our businesses, organizations, and institutions echo larger problems in public life.

Relevance bears strongly on organizational morale. The reason is that performance indicators selected for their relevance to strategic assumptions will naturally evaluate the quality of plans as well as the quality of effort against them. And that's all you may need to rescue the morale of a division chafing under goals the staff feel they can't control—a clear understanding that strategies and goals are subject to criticism in the light of performance results as well as execution. Without such an understanding, the staff may feel they're being held responsible for someone else's planning errors.

Acquisition targets that are relevant to your business are good; others aren't. But there may be a deeper connection between relevance and acquisition planning. This book's definition of relevance provides a criterion for determining when the ongoing business results and experience of an acquisition candidate can help improve the strategy of your business. And that kind of cross-fertilization may be more important than cost synergies at a time when technology, outsourcing, and provider networks can help almost any scale of business reach world-class levels of operating efficiency.

Relevance and compensation are also related. Variability in the results of metrics relevant to core strategic assumptions reveals a lot about the volatility of your business environment. So it also reveals a lot about the difficulty of the business problems you must

solve. Under the right conditions, you can use that variability to make sure the difficulty of the jobs of similarly paid executives is comparable and compensation is fair.

Relevance even bears on complex problems like global warming and economic development. A primary reason for the complexity of these problems is the sheer difficulty of measuring the success of alternative environmental and development strategies. Both areas cry out for hard thought on how to find metrics that test what these strategies really assume. This book tries to supply some of that hard thought. But its larger public policy contribution may be to encourage a more experimental approach to problems of public choice. For example, the belief that our latest development theories were right and lacked only confirmation has prevented us from making headway on development problems for decades. It's time to try assuming that we're wrong.

Finally, relevance helps separate leadership traits such as resolve and flexibility from stubbornness and inconstancy. There are as many dramatic cases of CEOs who persevered and succeeded, for instance, as CEOs who sank with their ships. It may be easier to draw the line if we recognize that learning from experience requires leaders to be clear about their ideas precisely in order to subject them to the harshest relevant tests.

What these connections really suggest is that the testable strategies, relevant metrics, and reviews of results proposed in the following pages are more than a guide for coping with performance information overload. They embody an experimental approach to management and problem solving. It's an approach that reflects a thorough fallibilism about the strategies we construct to meet our goals and an optimism that we can always do better.

It may be too much to hope this approach spares the life of a future Matt Hull. But the A-10 cockpit transcript shows that even in the heat of battle, it's possible to ask what we're really assuming. And from there it's a short step to focusing on what could prove us wrong rather than just what suggests we may be right.