

The False Alignment Trap

Leaders often seem to agree when they really don't. That's when transformations fall apart. PAGE 46







Cartier

The AI era is a human one.

We believe data and AI should be used to improve lives, strengthen communities and help solve meaningful problems. As AI becomes more capable, the need for responsibility and deep industry understanding only grows.

Domain expertise matters more than ever in the age of generative and agentic AI, grounding innovation in real-world context and judgment. Our commitment is to keep people in charge of the outcomes, and to build technology that supports transparency and trust.

This promise reflects the kind of company we set out to be in 1976 and the company we choose to be 50 years later. As we look to the next 50 years, we remain committed to keeping people at the center, innovating responsibly and serving as a trusted partner – combining advanced AI with proven industry expertise to help power the world's decisions.



Jim Goodnight, CEO, SAS



The Power to Know®
for 50 Years





GAGGENAU

FOR THOSE WHO KNOW

Creators of luxury kitchen appliances.
Established 1683.

Experience it



The difference is Gaggenau



Turn AI hype into hype-worthy results.

Build an AI foundation that's progress over promises.

We can help you go beyond strategy to turn your AI investment into achievable results you can see, measure and scale.

Learn more at unisys.com/ai

Keep Breaking Through

 **unisys**

Contents

July–August 2026

45 Features

46 **CHANGE MANAGEMENT**

The False Alignment Trap

Leaders often behave as if they agree on why, what, and how to change—when they actually don't. That's why most transformation efforts fail.

Julia Dhar, Kristy R. Ellmer, and Philip Jameson

58 **ORGANIZATIONAL DECISION-MAKING**

What Companies Get Wrong About Decision Rights

Four common mistakes and how to avoid them.

Lindy Greer, Jennifer Jordan, and Maxim Sytch

68 **STRATEGY**

The Power of Strategic Centering

Choosing a clear organizing principle can help companies navigate uncertainty and act faster.

Rita McGrath

COVER ART

HBR staff using AI;
*illustrated facial
features by Ryan Snook*



46

“The biggest adjustment in PE is trusting that it really is all about the value-creation plan, not about appeasing stakeholders.”

“MAKING THE LEAP FROM CORPORATE LEADER TO PE-BACKED CEO,” PAGE 80



80 LEADERSHIP

Making the Leap from Corporate Leader to PE-Backed CEO

Not everybody can manage the transition. Here's what it really takes to succeed.

Samantha Hellauer et al.

88 CUSTOMER EXPERIENCE

Why the Best Immersive Experiences Succeed

They invite deep engagement by answering key questions people ask themselves.

Joseph C. Nunes and Wendy Heimann

100 LEADERSHIP

Six Ways Leaders Harness Stress

Executives perform best when they understand their default reactions—and broaden their range of responses.

Jon Miller and Drew Keller

112 ENTREPRENEURSHIP

How Agentic AI Supercharges Startups and Threatens Incumbents

A new operating model shrinks the time and capital needed to build companies, redefining how ventures launch and scale.

Vivian S. Lee, Linda Mantia, and Jon McNeill

122 BUSINESS & SOCIETY

Bring Back Managing for Value

As capital costs rise, companies need a disciplined way to allocate resources without abandoning commitments to employees, customers, and the environment. Here's how to find that balance.

Michael Mankins and Matthew Crupi

IF IT CAN'T MEET AI SCALE IT'S NOT REAL RESILIENCE.

It's time to get real about the impact of AI on your business—generating massive opportunity, yes. But equally out-sized complexity, threats, and scale.



A new reality that calls for a new approach to resilience. Unified to close gaps between security, cyber resilience, and recovery. Scalable to multiple clouds, and beyond. Able to recover cleanly, quickly, repeatedly. Shifting resilience from crisis backstop to a continuous operational process. The answer isn't more and more solutions. It's the one that does it all.



Our Commitment to Sustainability

We're proud that the paper we use in our print magazine is certified under the Sustainable Forestry Initiative® program, meaning that it comes from responsibly managed sources and is a renewable resource.



17 Idea Watch

New Research and
Emerging Insights

17 IN THEORY

How the Best Leaders Respond to Rule Breaking

They first seek to understand why it happened. **PLUS** Why you should take boring meetings, the cost of posting wide salary ranges, how effective managers encourage their employees to develop new skills, and more.

29 In Focus Understanding the Hidden Risks of AI Tools

30 AI Doesn't Reduce Work— It Intensifies It

*Aruna Ranganathan
and Xingqi Maggie Ye*

33

When Using AI Leads to “Brain Fry”

Julie Bedard et al.

37

LLMs Are Manipulating Users with Rhetorical Tricks

Thomas Stackpole

40

Research: Using AI Can Stifle Innovation, But It Doesn't Have To

Chengwei Liu et al.

131 Experience

Advice and
Inspiration

131 MANAGING YOURSELF

How Elite Sports Coaches Make High-Pressure Decisions

What they do before, during, and after critical choices.

Alan McCall et al.

136



136 HOW WE DID IT

The Former CEO of Procafecol on Building a Global Growth Strategy for Juan Valdez Coffee

By targeting a narrow set of international markets and seeking out values-aligned local partners, the company has put itself on a sustainable path to scale.
Camila Escobar

141 CASE STUDY

A Key Supplier Failed a Labor Audit. Now What?

A furnishings company must decide how to address violations at a partner's factory.
Sandeep Puri

148 LIFE'S WORK

José Andrés

*Camila Escobar, former
CEO of Procafecol*

Departments

12 FROM THE EDITOR

13 CONTRIBUTORS

146 EXECUTIVE
SUMMARIES

Lasting wealth is built through coordination.

We bring structure to fiduciary trust oversight, investment management, and multi-generational planning so every part of your wealth works together.



Investment Management & Consulting | Trust Services
Family Office | Philanthropy & Family Continuity | Real Estate

CONTACT

Tim McCarthy (626) 463-2545 whittiertrust.com

Whittier Trust generally requires a minimum of \$15 million in marketable securities for new client relationships. Investment and Wealth Management Services are provided by Whittier Trust Company and The Whittier Trust Company of Nevada, Inc. (referred to herein individually and collectively as "Whittier Trust"), state-chartered trust companies wholly owned by Whittier Holdings, Inc. ("WHI"), a closely held holding company. This document is provided for informational purposes only and is not intended, and should not be construed, as investment, tax or legal advice. Past performance is no guarantee of future results and no investment or financial planning strategy can guarantee profit or protection against losses. All names, characters, and incidents, except for certain incidental references, are fictitious. Any resemblance to real persons, living or dead, is entirely coincidental.

Are We All in Agreement?



Amy Bernstein

THAT MOST ORGANIZATIONAL change initiatives fail is no surprise to anyone who’s been paying attention. There are many reasons why this is true, touching on issues of culture, process, people management, and leadership. In this month’s cover story (page 46), Julia Dhar, Kristy R. Ellmer, and Philip Jameson—all consultants from BCG—explore one particular area of leadership dysfunction: false alignment.

To succeed, they write, every transformation needs clear answers to the following questions: Why are we changing our company? What are we changing about our company? (And what are we *not* changing?) And how will the changes occur?

Every member of the leadership team must truly agree on the answers. “When company leaders say, ‘We are aligned,’ what they usually mean is, ‘We are not in one another’s way,’” say the authors.

That’s not enough. “Leaders who settle for mere alignment typically find that it fails them in the end....By contrast, leaders who work hard to create detailed and explicit compacts—what we call true agreement—find that they can effectively make progress on shared priorities and hold one another to account.”

The intense cohesion and astringent honesty required to achieve true alignment may make some senior leaders uneasy; they’ll need to get past that. Change is never comfortable, but the discomfort is entirely worth it when it enables an organization to navigate the kinds of transformation so many of us are pursuing now.

Amy
Amy Bernstein
Editor in chief

Connect with HBR



JOIN US ONLINE
WWW.HBR.ORG
FACEBOOK @HBR
INSTAGRAM
@harvard_business_review

LINKEDIN
Harvard Business Review
TIKTOK
@harvardbusinessreview

WHATSAPP
Harvard Business Review
X @HarvardBiz
YOUTUBE
@harvardbusinessreview



CONTACT HBR
PHONE
800.988.0886

EMAIL
editors@hbr.org
customerservice@hbr.org

Elle Honnain

Contributors



In high school **Julia Dhar** represented Australia on the team that won the World Schools Debating Championships—an experience she says first revealed the transformative power of productive disagreement. Now a managing director and partner at Boston Consulting Group, Dhar has seen organizations succumb to the problem of “false alignment” as leaders hesitate to challenge one another’s thinking. In this issue she and her coauthors present a playbook, drawn from their book *How Change Really Works*, that will help leadership teams move from superficial consensus to true agreement.

46 The False Alignment Trap



After years studying authenticity in consumer behavior, USC Marshall School of Business professor **Joseph Nunes** began noticing a shift: Companies were no longer simply selling products. They were now staging “immersive experiences” like the walk-through destinations that Netflix built around *Stranger Things* and the wine tastings and tours of the Tuscan hills curated by Italian luxury retailers. In this issue, alongside Wendy Heimann, an expert in location-based entertainment, Nunes examines what truly makes such experiences compelling.

88 Why the Best Immersive Experiences Succeed



Growing up, **Lindy Greer** was fascinated by what made some teams—from soccer squads to equestrian duos—work while others floundered. Now a professor at the University of Michigan’s Ross School of Business, she researches how power can break teams when they view it as a fixed resource—or strengthen them when they see it as a flow of energy to be shaped. In her article in this issue, she and her coauthors explore how teams that are defining and assigning roles in decision-making can avoid power struggles and supercharge their performance.

58 What Companies Get Wrong About Decision Rights



Fresh out of a master’s degree program in the 1980s, **Rita McGrath** joined the City of New York’s Department of General Services just as it began digitizing processes long done with pen and paper. Watching a bureaucracy expand its capabilities in real time sparked what would become a central theme of her scholarship: how organizations develop new strengths amid disruption. In this issue the longtime Columbia Business School professor introduces *strategic centering*, a framework designed to help companies adapt faster as competition and change intensify.

68 The Power of Strategic Centering



Photographer **Cheuk Lun Alan Lo** uses his medium to explore the tensions between competing emotions, individual aims, and societal pressures. For his photographs in this issue, he collaborated with the floral artist Vanessa Lyu to dismantle fresh blossoms, stitch them back together with bright cotton thread, and photograph them against stark backdrops. The visibly mended blooms evoke both vulnerability and resilience, mirroring the work required to build the true agreement described in the article his images accompany.

46 The False Alignment Trap



SUMMER SPECIAL ISSUE

Break Down Silos

To innovate and compete, companies must enable collaboration across teams and organizations—but most struggle to make it work.

This special issue delivers HBR's most practical advice for leaders on building the relationships, systems, and cultures needed to solve today's complex problems.

Each quarterly *Harvard Business Review* special issue focuses on a single, timely theme and includes expert-authored articles from HBR's rich archives, along with concise article summaries.

Available on newsstands and hbr.org.



EDITOR IN CHIEF Amy Bernstein

EDITOR, HBR.ORG
Maureen Hoch

EDITOR, HBR
Daniel McGinn

EDITORIAL DIRECTOR, HBR PRESS
Dave Lievens

EXECUTIVE CREATIVE DIRECTOR
John Korpics

EXECUTIVE EDITORS
Alison Beard
Ania W. Masinter
Sarah Moughty

EDITOR AT LARGE
Adi Ignatius

EDITORIAL

MANAGING EDITOR, HBR.ORG Gretchen Gavett

SENIOR EDITORS Scott Berinato, Courtney Cashman, David Champion, Paige Cohen, Susan Donovan, Kevin Evers, Eben Harrell, Jeff Kehoe, Scott LaPierre, Toby Lester, Juan Martinez, Steven Prokesch, Thomas Stackpole, Erica Truxler

PRODUCTION DIRECTOR & MANAGING EDITOR, HBR PRESS Allison Peter

SENIOR ASSOCIATE EDITORS Stefanie Fernández, Alana Grambush, Kelsey Hansen, Ramsey Khabbaz

EDITORIAL COORDINATORS Rachel DuRose, Emily Lang

EXECUTIVE ASSISTANT Christine C. Jack

CONTRIBUTING EDITORS Eileen Ahuja, Karen Dillon, Sarabeth Fields, Amy Gallo, Kelly Messier, JM Olejarz, Anand P. Raman, Dana Rousmaniere

DESIGN

DESIGN DIRECTORS Stephani Finks *HBR Press*, Susannah Haesche *HBR*, Marta Kusztra *Product & UX*, Karen Player *Special Projects*

SENIOR GRAPHICS EDITOR Macaulay Campbell

SENIOR DESIGNERS Carey Bass *HBR*, Kim Lapan *UX*, So Young Park *Product*

SENIOR PHOTO EDITOR Sasha Patkin

DESIGNER Alex Belser *Multimedia*

PRODUCTION

MANAGER, EDITORIAL PRODUCTION Christine Wilder

SENIOR PRODUCTION EDITORS Kristin Murphy Romano, Jennifer Waring

SENIOR PRODUCTION SPECIALIST Robert Eckhardt

PRODUCTION DESIGNER Emily Peterson

PRODUCTION EDITORS Victoria Desmond, Anne Starr

PRODUCTION SPECIALIST Alexie Rodriguez

EDITORIAL OFFICES

60 Harvard Way, Boston, MA 02163
617-783-7410 | fax 617-783-7493
HBR.org

VOLUME 104, NUMBER 4
JULY-AUGUST 2026

Printed in the U.S.A.

Copyright 2026 Harvard Business School Publishing Corporation. All rights reserved.

COPRESIDENT, HARVARD BUSINESS PUBLISHING Sarah McConville

**VICE PRESIDENT,
DEPUTY PUBLISHER**
Erika Heilman

**VICE PRESIDENT, GLOBAL
ADVERTISING SALES**
Gail Day

**VICE PRESIDENT, NATIONAL
ADVERTISING & MARKETING**
Craig Catalano

**MANAGING DIRECTOR,
DIGITAL PRODUCT STRATEGY**
Dave Lefort

ADVERTISING

MANAGING DIRECTOR, ANALYTIC SERVICES Alex Clemente
DIRECTOR, U.S. SALES & PRODUCT DEVELOPMENT Daniel Cohen
DIRECTOR OF RESEARCH, ANALYTIC SERVICES Beth Tracton-Bishop
SENIOR EDITOR, RESEARCH & SPECIAL PROJECTS Todd Pruzan
MANAGING EDITOR, ANALYTIC SERVICES Anthony Baldo
ASSISTANT EDITOR, ANALYTIC SERVICES Amanda Yao
ASSOCIATE DIRECTOR, MARKETING Samantha Holloway
MANAGER, MARKETING Monica Ha
DIRECTOR OF PROJECT MANAGEMENT, ANALYTIC SERVICES Samantha Barry
SENIOR RESEARCHER, ANALYTIC SERVICES Mary Callaghan
SENIOR PROJECT MANAGER, ANALYTIC SERVICES Audrey White
COORDINATOR, MARKETING Caitlin Amorin
CONTRIBUTING CONSULTANT MaryAlice Holmes

MARKETING & COMMUNICATIONS

DIRECTOR, CONSUMER MARKETING Nini Diana
DIRECTOR, COMMUNICATIONS Amy Poflak
ASSOCIATE DIRECTOR, MARKETING FOR CUSTOMER JOURNEYS Caty Jackson
SENIOR MANAGER, LEAD GENERATION & CONVERSION Josh Getman
MANAGER, DIGITAL MARKETING Stephanie Kazemi

PRODUCT MANAGEMENT

DIRECTOR, PRODUCT MANAGEMENT Anne Bartholomew
SENIOR PRODUCT MANAGER Jon Goldberg
SENIOR PRODUCT MANAGER, MOBILE Gerrie Diaz
PRODUCT MANAGER, AUDIO Ian Fox
SENIOR TECHNICAL PRODUCT MANAGER Frederic Lalande
SENIOR PRODUCT MARKETING MANAGER, LEARNING Marqui Mapp-Taylor
AGILE LEADS Kara Berrini, James Monia

BUSINESS & CUSTOMER ANALYTICS

SENIOR DIRECTOR, BUSINESS ANALYTICS & INSIGHTS Jessica Avery
DIRECTOR, CUSTOMER & MARKETING ANALYTICS Emily Freeman Kroger
ASSOCIATE DIRECTOR, EMAIL AUTOMATION & OPERATIONS Jerry Knott
ASSOCIATE DIRECTORS, EXPERIMENTATION & ANALYTICS Chris Chaky, Abigail Dawson
SENIOR MANAGER, EMAIL PRODUCTION Frances Lee
MANAGERS, EMAIL MARKETING Milo Goodman, Cara Vaccaro
ANALYST, EXPERIMENTATION & ANALYTICS Tony Manev

FINANCE & OPERATIONS

SENIOR DIRECTOR, FINANCE & OPERATIONS Greg St. Pierre
ASSOCIATE DIRECTOR, OPERATIONS & CUSTOMER EXPERIENCE Greg Daly
SENIOR MANAGER, FINANCE Krishna Dudhwewala
MANAGER, CUSTOMER EXPERIENCE Danielle Weber
SPECIALIST, DATA INTEGRITY & OPERATIONS Edward D. Domina IV

HBR PRESS

COMMERCIAL DIRECTOR Jon Shipley
DIRECTOR, MARKETING & COMMUNICATIONS Julie Devoll
DIRECTOR, B2B & CUSTOM SALES Lindsey Dietrich
ASSOCIATE DIRECTOR, MEDIA & COMMUNICATIONS Sally Ashworth
SENIOR MANAGER, OPERATIONS Rick Emanuel
MANAGER, MARKETING & COMMUNICATIONS Alexandra Kephart
SENIOR ASSOCIATE, SALES Jordan Concannon

TECHNOLOGY

CONTENT SYSTEMS ARCHITECT Matt Wagner
SOFTWARE ARCHITECTS Bill Munyan, Kenneth Pereira
SOFTWARE ARCHITECT, MOBILE Ahmad Khan
QUALITY ASSURANCE ARCHITECT Matthew Han
FRONT-END ENGINEER Caitlin Smith

WORLDWIDE ADVERTISING OFFICES

NEW YORK
3 Columbus Circle, Suite 2210
New York, NY 10019
973-769-3631 | fax 212-956-0933
Maria A. Beacom *Sales Director,
Technology & North American Research*

CHICAGO 312-206-6045
James A. Mack *Sales Director,
Central/Southwest U.S.*

LONDON 44-0-7890-608-707
Alexander Bart *Sales Director,
EMEA & APAC*

CALIFORNIA 415-986-7762
BEIJING 86-10-6551-0663
FRANCE 33-01-4643-0066
HONG KONG 852-237-52311
INDIA 212-872-9291
JAPAN 81-03-3541-4166
KOREA 82-2-730-8004
SOUTH AMERICA 55-11-98160-0177
UAE 971-4-228-7708
UNITED KINGDOM 44-20-7291-9129

For all other inquiries, please
call 212-872-9280 or visit
HBR.org/hbr-advertising-sales.

SUBSCRIPTION SERVICES

Harvard Business Review
20 Guest St, Suite 400
Brighton, MA 02135
customer@hbr.org
US/Canada: 800-274-3214
International: 1-617-783-7500

RATES PER YEAR

United States \$119
Canada US\$139
Mexico US\$139
International US\$165

ARTICLE REPRINTS

To purchase reprints, go to HBR.org.

LIBRARY ACCESS

Libraries offer online access to current and back issues of *Harvard Business Review* through EBSCO host databases.

SUBMISSIONS

We encourage prospective authors to follow HBR's "Guidelines for Authors" before submitting manuscripts.

To obtain a copy, please visit HBR.org; write to: The Editor, *Harvard Business Review*, 60 Harvard Way, Boston, MA 02163; or email hbr_editorial@HBR.org.

A NOTE TO READERS

The views expressed in articles are the authors' and not necessarily those of *Harvard Business Review*, Harvard Business School, or Harvard University. Authors may have consulting or other business relationships with the companies they discuss.

Rotman
MANAGEMENT

"Read cover
to cover.
Superb."

Tom Peters

*Author, In Search of Excellence;
Thinkers50 Hall of Fame*



Published by Canada's
leading business school
each issue is filled with
thought-provoking ideas
and insights from leading
global academics
and management
practitioners.

Just \$49.95 CAD/year
Available in print or digital

www.rotmanmag.com/subscribe



**Rotman School of Management
UNIVERSITY OF TORONTO**

IF YOU'RE
NOT TALKING
ABOUT STORAGE,
YOU'RE NOT TAKING
AI SERIOUSLY



AI Data Infrastructure
Built on a Solid Foundation





idea**W**atch



IN THEORY

How the Best Leaders Respond to Rule Breaking

They first seek to understand why it happened.

WHEN A RULE gets broken at work, a leader's first instinct is to act fast. The business response feels obvious: punish the offender and move on. But it's rarely that simple.

In a new paper that reviews and synthesizes more than 250 studies across four decades, Michael J. Gill, an associate professor of organization studies at the University of Oxford's Saïd Business School, shows that people break rules for very different reasons. Some do it for personal gain; others to help customers or colleagues. And sometimes the environment pushes people over the line.

The Ethics & Compliance Initiative's "2023 Global Business Ethics Survey" reported that 65% of employees around the world said they had observed misconduct at work, up from 60% in 2020. In the United States more than half of workers reported seeing rule breaking. Prior surveys found much lower levels, suggesting the problem is increasing.

Gill started studying the topic after hearing the same question from companies: Why does this keep happening? His work divides misconduct into four types: self-interested, or meant to help oneself; prosocial, or meant to help



others; corrupted, or driven by pressure; and edified, or encouraged in the name of doing the right thing. Most organizations treat all of them the same way, which is a problem because breaking rules can be beneficial. For example, an employee trying to appease an upset customer may expedite an order even if company policy doesn't call for it.

Before punishing rule breakers, leaders should ask two critical questions: Was the behavior constructive or destructive? Was it driven by individual choice or the situation? "If your first move is to punish someone, you may be solving the immediate issue but missing the real cause. And that can hurt the business in the long term," Gill says. "A better approach is to be less reactive and more curious." He recommends several best practices for resolving rule breaking at work:

Understand what happened. Before jumping to conclusions, talk to the people involved. Ask why the policy was violated—and do so in a way that isn't confrontational and that gets to the reality of the situation. For instance, was the person's rule-breaking behavior prosocial or corrupt? This assessment sounds obvious, but often it doesn't happen. Many leaders assume they already know the answer, and many employees assume they'll be judged no matter what they say.

That's why tone matters. If employees think your goal is understanding, you're far more likely to hear what's really going on. That approach builds trust over time and helps leaders understand the pressures behind people's decisions. It also shifts conversations from blame to learning, which leads to better outcomes for everyone.

Four Kinds of Rule Breakers

Leaders faced with a rule-breaking employee should seek to understand why that person acted that way. If the misconduct happened in order to help a customer, for instance, it might signal a need to revamp your team's processes. Before jumping to punishment, ask two questions: Was the behavior constructive or destructive? Was it driven by individual choice or by the situation? The answers will tell you what kind of rule breaker you're dealing with.

Impact of rule breaking	Cause of rule breaking	
	CHOICE-DRIVEN	SITUATION-DRIVEN
DESTRUCTIVE	Self-interested <i>"I chose to break this rule for personal gain."</i> Example: A consultant inflates the number of billable hours spent working with a client.	Corrupted <i>"I felt pressure from colleagues to break this rule for organizational gains."</i> Example: A programmer is coerced by colleagues to release a software update without fully testing it to meet a production deadline.
CONSTRUCTIVE	Prosocial <i>"I chose to break this rule to help someone."</i> Example: A customer service rep accepts a product return past the allowable time frame for a loyal customer.	Edified <i>"I felt pressure from outside the organization to break this rule to help others."</i> Example: A nurse ignores a hospital's outdated discharge policy because the profession's code of ethics emphasizes patient welfare over rigid rules.

"If your people believe they are in trouble, they'll shut down, but if they think you're trying to understand what happened before coming to your final conclusions, they are more likely to open up," Gill says.

Separate intent from impact. Not all misconduct is created equal. Yet when people believe their intent doesn't matter, they either stop using judgment or start hiding what they're doing. And when every misstep is treated as self-interest or corruption, they learn to play it safe or stay quiet. The result is less initiative, less transparency, and fewer opportunities to fix problems early.

But if employees see that good intentions are recognized—even when outcomes fall short—they are more likely to take smart risks and speak up when something feels off.

Pay attention to patterns. One-off incidents can be misleading; patterns are where the insights are. If a rule keeps getting broken, that usually points to something bigger, such as a policy that

is unrealistic, unclear, or in conflict with what employees are rewarded for doing. Repeated violations can be a form of feedback about where the system doesn't line up with the work. For example, an employee who regularly buys office supplies from an unapproved vendor to get a better price is violating policy, but the intent is to save the company money, not enrich himself. When leaders spot a pattern like this, it's often useful to take a hard look at the rule.

Uncover what you're not being told. Many instances of observed misconduct never make it up the chain. People may not report it, and managers may not escalate it. That's why leaders need other ways to understand what's happening, such as establishing anonymous reporting channels, spending more time seeing how work gets done, or simply asking broader questions about the company's policies and how restrictive they are.

Leaders can treat these signals as an early warning system. Small

workarounds, quiet shortcuts, and informal fixes often show up before bigger problems do. Paying attention to them helps you see where pressure is building inside the system. For example, a company rule may require customer support responses to be approved by a manager before being sent. In practice, however, customers may expect quick replies, and getting approvals can cause delays and frustration. When leaders make space for that kind of visibility, they are less likely to be surprised later and more able to fix issues before they escalate.

Fix the causes of rule breaking. If you execute the earlier steps well, you'll identify root causes that may be driving misconduct, including bad incentives, unclear rules, or pressure that's pushing people in the wrong direction. Now it's time to find a systemic fix for the problem. Focusing only on the individual might fix the behavior in the short term but won't always solve what's driving it.

None of this means rules don't matter. Many policies should never be broken, even for prosocial or edified reasons. But others may need to be revisited, clarified, or applied with more flexibility. Leaders who respond with curiosity instead of assumptions are far more likely to get it right.

"Misconduct isn't going away," Gill says. "The real goal isn't to catch it faster. It's to understand it better." ■

HBR Reprint F2604A

 About the research "Rule Breaking in Organizations: An Integrative Review," by Michael J. Gill (Academy of Management Annals, 2026)

IN PRACTICE



"Coach the Intent but Enforce the Boundaries"

Art Gilliland is the CEO of Delinea, a cybersecurity company based in San Francisco with approximately 1,200 employees. Gilliland spoke with HBR about how he distinguishes between true misconduct and well-intentioned rule breaking, and how he works with managers to apply consistent judgment across situations. Edited excerpts of the conversation follow.

Is it realistic for managers to respond to rule breaking in a nuanced manner?

Yes, but maybe not in the way it reads on paper. I work in cybersecurity, where you don't have the luxury as a manager to sit back and be overly introspective every time something bad happens. By the time it gets to the CEO, it's usually serious. So when I am involved, my first step is to seek to understand what happened. If it's something regulatory or legal, you must take a hard line; there's just not much flexibility there. But a lot of the time people are trying to help a customer or move quickly, and they cross a line without realizing it. In those cases I'd much rather educate and coach.

Have you dealt with a serious misconduct issue? We've had situations where employees shared pricing information for our services with customers when they

shouldn't have. In our model we partner with third parties to sell services, and they set the market price. Getting ahead of that process can be a regulatory violation. Usually employees aren't thinking about the potential for price-fixing or legal exposure; they're just being responsive to customers.

How did you react? What I don't do is immediately come down hard on people. First, I ask how often the rule breaking has occurred with this specific employee. Then we spend time educating them on why that rule exists and what the right behaviors are. If it's a one-off problem, coaching usually solves it.

What happens when the behavior isn't meant to help the company? Then you're dealing with a different kind of issue. We still want to understand the context first, but the lens changes. Now you're asking, Was this self-serving, careless, or knowingly against the rules? If someone knowingly breaks a rule for personal gain, our tolerance is zero. You might still give feedback, but you're likely taking formal action.

What if you find a pattern of problems? Repeated behavior tells you something different. Either the person isn't understanding the expectation or they're choosing not to follow it. At that point you must be direct. You still want to be fair and give clear feedback, but you're also thinking about the message you're sending to the rest of the team. If someone continues breaking a rule after coaching and support, eventually you must make the call to let them go.

Have you been able to train managers to handle these situations well? Yes, and it requires being clear about what matters, where the hard lines are, and where there's room for judgment. A lot of it comes down to communication and repetition. We talk through real examples, what we did, and why. The goal is to coach the intent but enforce the boundaries. You're never going to get perfect alignment between people and rules, and that's OK. But if you hire the right people, reinforce the company's values, and stay involved in the tougher cases, you can get as close to perfect as possible. ■



THE DEATH OF THE COVER LETTER

In a study of 5.5 million applications on a large freelance platform in 2023, researchers examined what happened when an AI tool that generated role-specific cover letters was introduced. At first its users saw callbacks jump by 50%. But within two months, as more and more people harnessed the tool, that edge disappeared, and employers shifted their focus to applicants' work histories when deciding whom to interview. "Signaling in the Age of AI: Evidence from Cover Letters," by Jingyi Cui, Gabriel Dias, and Justin Ye

PERSONAL PRODUCTIVITY

Take That Boring Meeting

It's tempting to skip a dull-looking meeting on your calendar or mentally check out once it gets going. But you're probably underestimating how valuable it will turn out to be.

Across nine studies involving 1,800 participants in France, the United States, and Singapore, researchers compared what people expected from conversations with what they actually experienced. In one study with 100 French subjects, those assigned to take part in a conversation on a topic they'd rated "not at all interesting" reported enjoying it much more than they'd predicted they would. Their predictions about how much they'd enjoy discussing interesting topics, however, were pretty accurate.

A second study revealed the same pattern even when both participants had ranked the topic as boring, suggesting

that one enthusiastic person didn't "carry" the conversation. People's expectations were also exceeded whether discussions of topics they'd thought were dull were remote or in person or between strangers or friends.

The researchers say this happens because people put too much weight on the static aspects of a future conversation (like the topic) and too little on dynamic aspects that occur once a dialogue begins (like listening, responding, and guiding the direction of the exchange).

Follow-up experiments also showed that having live conversations about potentially boring topics made a significant difference. Again, people enjoyed them much more than they'd predicted, but when they only read a transcript or watched a recording of a conversation on a dull-sounding topic they found it as unengaging as they'd anticipated and sometimes even more so.

"Senior leaders are trained to allocate attention where it matters most. But the value of a conversation often

becomes visible only after it begins," the researchers write. "Sometimes the meetings that look least interesting on your calendar are the ones most likely to surface insight, strengthen relationships, and reveal early risks."




ABOUT THE RESEARCH "Conversations About Boring Topics Are More Interesting Than We Think," by Elizabeth N. Trinh, Nicole Thio, and Nadav Klein (*Journal of Personality and Social Psychology*, 2026)

MANAGEMENT SKILLS

Employees Who Enjoy Their Jobs Get Stuck with Extra Work

When a new task—like planning a holiday party or serving on a committee—needs to be assigned, it may make sense to tap someone who seems to really like doing work. But new research shows not only that self-motivated employees bear the brunt of added tasks but that doing them can undermine their job satisfaction.

In six studies with U.S. participants, researchers examined how managers allocate tasks that fall outside employees' regular core responsibilities. In a study of 834 managers rating two of their direct reports, the managers were significantly more likely to say they'd assign extra tasks to the person they perceived to be more intrinsically motivated (who seemed to take pleasure in working), even after accounting for performance, experience, and how much they wanted each person to succeed.



PROTECT AND GROW YOUR WEALTH— FOR THIS GENERATION AND THE NEXT.

PRIVATE WEALTH MANAGEMENT

This five-day executive education program equips individuals and families to navigate the complexities of managing and transitioning wealth. Learn from Chicago Booth's world-renowned faculty and leading practitioners, and gain research-driven insights and practical tools to grow assets while strengthening family-wealth purpose and governance.

2026 DATES

November 16–20

This program is offered twice a year at Gleacher Center in downtown Chicago.

Admission is selective and based on a completed application. Applications are reviewed on a rolling basis.

Learn more and apply at chicagobooth.edu/pwm.
+1 312.464.8732 | exec.ed@chicagobooth.edu

CHICAGO BOOTH 
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

CHICAGO | LONDON | HONG KONG

In a lab experiment, 306 participants were divided into groups of three, with one serving as a manager and two as employees. The managers chose which employee doing a paid data-entry project would be interrupted to complete a tedious side task that wouldn't count toward a performance bonus. Seventy-four percent of them chose the more intrinsically motivated worker, diminishing that person's main-task performance and odds of earning the bonus.

The researchers say this happens because managers think that employees who like their work will also enjoy unrelated tasks and will be less likely to burn out from the added load than other employees will. When the researchers studied this directly they found that employees high and low in intrinsic motivation found the added work equally draining, and highly motivated people saw a larger drop in enjoyment after being pulled away from their core responsibilities.

A follow-up study showed that when motivated workers repeatedly get burdened with additional projects, it reduces their anticipated job satisfaction and commitment and increases the likelihood that they'll quit. To protect the employees who care the most, the researchers suggest, companies should track extra work assignments and make sure they're shared by all.

ABOUT THE RESEARCH "Managers Allocate Additional Tasks to Intrinsically Motivated Employees: Exploring Mechanisms, Consequences, and Solutions," by Sangah Bae and Kaitlin Woolley (Organization Science, 2026)



CONSUMER BEHAVIOR

A No-Strings-Attached Gift Can Increase Customer Spending

Companies like Sephora, Lego, Costco, and Starbucks occasionally give away small freebies to people who enter their stores. New research suggests that those gifts do more than foster goodwill—they encourage shoppers to buy more.

Across seven studies, researchers examined what happened when companies offered an unearned gift—something just for walking in, with no purchase, points, or sign-up required. In two studies with European retailers, people who received a free beverage or a bag of ground coffee spent 16% to 32% more during a visit than people who got nothing. Whether the item was handed over on the spot or redeemed via a voucher, average basket size rose substantially.

To understand why, the researchers ran follow-up lab experiments using AI-generated videos that simulated store and coffee shop visits. They focused on two emotional reactions people feel when they receive a gift:

gratitude and obligation. Both rose when customers got a gift with no strings attached. People who felt a lot of gratitude for the gift were more likely to say they'd return more often, buy more in the future, and recommend the retailer. Those who felt a strong sense of obligation spent more immediately as a way to "repay" the company.

Further studies found that unconditional gifts don't need to be expensive. In one, the gifts' value varied from 50 cents to \$16. Every level of gift, including the smallest, significantly increased customers' gratitude, sense of obligation, loyalty, and planned spending. But large gifts boosted gratitude only slightly more than small gifts did and produced no meaningful jump in shoppers' sense of obligation. Unconditional gifts also worked equally well with both new and longtime customers and at least as well as gifts framed as rewards for loyal spending.

The researchers say this doesn't mean that marketers should ditch traditional loyalty programs, which help brands track value and segment customers and give participants a clear, predictable reason to concentrate their purchases with one company. However, their findings do suggest that surprise

How do you shape your strategy with confidence?

Transformative strategy and transactions
that work in practice, not just on paper.



The better the question.
The better the answer.
The better the world works.

 **EY** Parthenon
Shape the future with confidence



FORGET THE COMPETITION

New research shows that ads not only keep a brand's name top of mind but also knock rivals out of consumers' thoughts. In a field experiment with 327,000 web users, a one-day banner campaign run in 2015 by a major U.S. insurer boosted visits to its site from 0.2% to 0.8%—a 300% jump—while visits to competitors' sites fell by 11% the day after the ad went live.

"Consumer Memory and Competitive Interference: The Case of Auto Insurance Advertising," by Navdeep S. Sahni and Yifan Yang

no-strings-attached gifts are a cheap, flexible way to increase consumers' spending and emotional attachment.

ABOUT THE RESEARCH *"The Effects of Unconditional Gifts on Customer-Firm Relationships," by Paul W. Fombelle et al. (Journal of the Academy of Marketing Science, 2026)*

DEVELOPING EMPLOYEES

Good Managers Get Their Workers to Upskill

Are your company's employee-training programs suffering from low attendance and weak engagement? One limiting factor could be your middle management.

Using data from three firms—a car manufacturer in Argentina, a Colombian quick-service restaurant chain, and a retailer also in Colombia—researchers tracked which employees took part in paid, firm-specific training. They found that despite similar HR policies and strong links between training and eligibility for promotions, participation levels differed widely across similar units.

That variation was largely driven by individual middle managers, the researchers found. They began by identifying managers whose direct reports were more likely than average to go to training. At all three companies, the managers often rotated between stores and teams, making it easier to track their impact. The data showed that when a manager with above-average training uptake arrived at a new store, within weeks the number of employees

there who attended training jumped: by roughly 7% at the carmaker, 55% at the restaurant chain, and 44% at the retailer.


Surveys revealed that these high-performing managers had similar qualities: They scored higher on extroversion and social skills, involved employees more in decisions, reviewed performance metrics more often, and paid greater attention to weak performers, motivation, and well-being. Their teams also had lower turnover and absenteeism and higher promotion rates overall.

Employees with high-performing managers also proved more resilient amid uncertainty and change, such as during mandated production increases at the car plant or the launch of delivery apps at the chains. At the start of such

initiatives, output rose across units, but absenteeism spiked—except in teams led by strong managers. Those teams sustained higher production and had little or no increase in absences, partly because their managers continued to push training and promotions even under pressure.

Training investments will underperform if firms ignore this managerial layer, the researchers write. They recommend measuring the developmental impact of managers, rotating and rewarding those who champion training, and building people-management skills as a core part of any upskilling strategy.

ABOUT THE RESEARCH *"Training Within Firms," by Brayan Diaz et al. (working paper, 2025)*




COMPETITIVE STRATEGY


Instead of Guidance, AI Offers “Trendslop”

What happens when large language models are asked to give strategic advice? To find out, researchers asked top LLMs to weigh in on common strategic trade-offs in 2025. These included decisions like whether to compete by offering a low cost on a standardized product (commoditization) or to invest in distinctive features that set the product apart (differentiation); whether to replace human labor with technology (automation) or deploy technology to amplify the capabilities of the existing workforce (augmentation); and whether to pursue quick returns (short-term focus) or sustained advantage (long-term focus). The researchers found that the models consistently recommended “trendslop,” or strategies aligning with managerial buzzwords and cultural tropes, regardless of context, like a company’s size or industry.

LLMs’ average preference when making 50 choices between competing strategies



Source: Angelo Romasanta, Llewellyn D.W. Thomas, and Natalia Levina



HIRING & RECRUITMENT

Wide Salary Ranges Deter Female Applicants

Providing salary ranges for open positions is supposed to reduce gender pay gaps by giving job seekers better information while holding firms accountable for offering stated sums. But when such ranges are large, they can actually foster inequality.

In an initial analysis of roughly 10 million U.S. jobs advertised from 2020 to 2024, researchers found that companies that posted wider salary ranges tended to employ fewer women, even after controlling for firm size, pay levels, and other characteristics.

In one study 150 U.S. participants chose among otherwise identical job listings that differed only in pay range. Women were more than twice as likely as men to prefer those featuring narrower ranges (for example, \$55,000 to \$75,000 versus \$40,000 to \$90,000), a pattern explained by lower self-reported risk tolerance. Participants who chose narrow ranges also negotiated less aggressively, proposing counteroffers

nearly \$10,000 lower than those who chose wide ranges.

A field experiment with 1,148 job seekers considering a real remote role had a similar finding: In it women were significantly more likely than men to apply when the posting showed a narrow salary range but not when it showed a wide one.

However, in a final experiment with 969 working adults, the researchers found that adding a short note that mentioned the typical starting salary and described the criteria used to set final pay made a wide range feel less risky, increased women’s willingness to choose a posting with one, and largely eliminated the gender gap in counteroffers.

To ensure that all top candidates apply to your listing, the researchers suggest, pair salary ranges with clear information about typical starting pay and how offers are determined, so that transparency reduces uncertainty instead of amplifying it.

ABOUT THE RESEARCH “The Implications of Pay Range Transparency on Job Application Preferences and Negotiations,” by Alice J. Lee, Tae-Youn Park, and Sungyong Chang (Journal of Applied Psychology, 2026)



CUSTOMER SEGMENTATION

Confidence Shapes What People Buy Next

Marketers have faith that familiarity will breed loyalty. But a study of nearly 100,000 U.S. consumers suggests that just the opposite is true: As people gain experience in a product category, they become less confident about their knowledge of it, and that doubt sends them searching for new brands.

To examine how confidence shapes loyalty, the researchers used linguistic software to analyze 3.7 million wine, beer, and cosmetics reviews, tracking how expressions of confidence shifted as consumers sampled more products in a category. They then used the reviews as a proxy for purchases to link a consumer's confidence about one purchase to what product that person bought next.

Across each category the pattern was the same: Consumers started off surprisingly sure of themselves and tended to stick with the same items or brands. As they gained experience in a product category, their confidence dropped. During this dip they were more likely to switch brands, and they waited longer before buying from the same brand again. Only after they gained a substantial amount of experience with the category did confidence—and brand loyalty—rebound.

To make sure this behavior wasn't driven by other factors, like less enjoyment of products or boredom, the researchers tested whether the pattern held even after accounting for how

POWER & INFLUENCE

Being a Reliable Leader Can Backfire

When organizations face turbulence—like leadership turnover or political infighting—many leaders take on additional responsibilities as buffers, fixers, or translators to help the organization navigate the turmoil. The trouble is that doing this can hide structural problems, concentrate risk, and leave those leaders with little capacity to shape the future. To protect their influence and impact, reliable leaders must recognize these three roles, understand their costs, and make strategic shifts that will build personal and organizational resilience.

	Role		
	Buffer	Fixer	Translator
What reliable leaders get saddled with managing	Emotional strain, conflict, misalignment.	Execution gaps, stalled initiatives, rework.	Ambiguity, unclear authority, conflicting priorities.
Why they're praised	"You keep the team steady."	"We know you'll get it done."	"You make things move."
What it quietly costs	Problems remain unaddressed; accountability diffuses.	Dependency on leader deepens; systemic execution gaps persist.	Decision rights blur; political exposure increases.
The solution	Stop acting as mediator in tense situations. Allow discomfort to reveal recurring problems and misaligned expectations.	Don't step in to solve every issue. Delegate authority and coach colleagues so that they can build their own problem-solving skills.	When authority over a problem is ambiguous, take a moment to decide who should be in charge and what accountability should look like for that issue.

Source:
Luis Velasquez

positive and emotional the reviews were and, in a lab task, for fatigue and engagement. Confidence still reliably predicted whether people switched away from what they'd chosen before.

This happens because people often start to recognize how much there is to learn about a category as they get more experience in it, the researchers write. Their uncertainty then increases, and that discomfort motivates people to distance themselves from their previous choices and opt for something new.

Three follow-up experiments affirmed this phenomenon. In them, participants were asked to name their favorite music genre and then were led to feel either sure or unsure of their expertise. Next they were given a list of songs organized by genre and asked what they wanted to listen to next. Those in the

low-confidence group were much more likely to choose music outside their preferred genre than those in the high-confidence group.



For marketers, the findings suggest that even satisfied customers may defect if they feel unsure about what they're doing. The remedy? The researchers say that after customers gain a moderate amount of experience—when their confidence is most fragile—firms can blunt people's doubts by offering education, clear guidance, and simple decision aids. 🍷

 **ABOUT THE RESEARCH** *"The Trajectory of Confidence: Experience, Certainty, and Consumer Choice," by Matthew D. Rocklage, Jonah Berger, and Reihane Boghrati (Journal of Marketing Research, 2025)*

Stop reorganizing. Start **adapting**.

Becoming an Octopus Organization

Learn to break down hierarchy, get rid of silos, and create an adaptable, resilient, learning organization in this eight-week newsletter series.



SIGN UP TODAY
hbr.org/newsletters





Evidence. Insight. Action.

Get expert-backed intelligence—drawn from the HBR Advisory Council and global business leaders—and explore hundreds of white papers and research reports from HBR Analytic Services.

RECENTLY PUBLISHED

Building Strong Brands

Sponsored by Frontify

Driving Business Value with Agentic AI

Sponsored by Wipro

Leveraging Alternative Payment Models to Ensure Broad Access to Personalized Medicine

Sponsored by BU Questrom

Procurement at a Crossroads: Evolve or Be Left Behind

Sponsored by Zip

Accelerating Growth Through Revenue Enablement

Sponsored by Seismic



View all the reports at
hbr.org/hbr-analytic-services

Register for the HBR Advisory Council at hbr.org/advisory-council.

inFocus

One Challenge, Four Strategies



THE CHALLENGE

Understanding the Hidden Risks of AI Tools

1

AI Doesn't
Reduce Work—
It Intensifies It
30

2

When Using
AI Leads to
“Brain Fry”
33

3

LLMs Are
Manipulating
Users with
Rhetorical Tricks
37

4

Research:
Using AI Can Stifle
Innovation, But It
Doesn't Have To
40

Jorg Greuel/Getty Images



These articles were selected from **HBR.org**.


Harvard Business Review
July–August 2026

29



IN FOCUS

Understanding the Hidden Risks of AI Tools



AI Doesn't Reduce Work— It Intensifies It

BY ARUNA RANGANATHAN AND
XINGQI MAGGIE YE

RIGHT NOW MANY companies are worried about how to encourage more employees to use artificial intelligence. After all, the promise of AI—of reducing the burden of work tasks by drafting routine documents, summarizing

information, and debugging code and allowing workers more time for high-value tasks—is tantalizing.

But are they ready for what might happen if they succeed? When leaders focus only on potential productivity gains, they may find themselves surprised by the complex reality of AI and not see what those gains are costing them until it's too late.

In our still-in-progress research, we have discovered that AI tools don't reduce work; they intensify it. In an eight-month study of how gen AI changed work habits at a U.S.-based technology company with about 200 employees, we found that employees using AI worked

at a faster pace, took on a broader scope of tasks, and spent more hours at work than those who didn't use AI, often without being asked to do so. Importantly, the company did not mandate AI use (though it did offer enterprise subscriptions to commercially available AI tools). On their own initiative, workers did more because AI made “doing more” feel possible, accessible, and in many cases intrinsically rewarding.

This may sound like a dream come true for leaders, but the changes brought about by enthusiastic AI adoption can be unsustainable, causing problems down the line. Once the excitement of experimenting fades, workers can find that their workload has quietly grown and they feel stretched from juggling everything that's suddenly on their plate. That workload creep can in turn lead to cognitive fatigue, burnout, and less effective decision-making. The initial productivity surge can give way to lower-quality work, turnover, and other problems.

This puts leaders in a bind. What should they do? Asking employees to self-regulate isn't a winning strategy. Rather, companies need to develop a set of norms and standards around AI use—an “AI practice.” Here's how leaders can set their employees up for success.

How Gen AI Intensifies Work

In our study we explored how gen AI tools changed the work habits of employees at the tech company. We observed workers in person two days a week, tracked internal communication channels, and conducted more than 40 in-depth interviews across engineering, product, design, research, and

Francesco Carra fotografo/Getty Images



Many workers noted that they were doing more tasks at once—and feeling more pressure—than before they used AI.

operations. We identified three main ways that work intensifies.

1. Task expansion. Because AI can fill in gaps in knowledge, workers assumed responsibilities that had previously belonged to others. Product managers began writing code; researchers took on engineering tasks; and individuals across the organization attempted work they would have outsourced, deferred, or avoided entirely in the past.

Gen AI made those tasks feel newly accessible. The tools provided what many experienced as an empowering cognitive boost: They reduced dependence on others and offered immediate feedback and correction along the way. Employees described this as “just trying things” with the AI, but these experiments meaningfully widened their job scopes. In fact, workers increasingly absorbed work that previously might have justified hiring additional temporary help or increasing head count.

There were knock-on effects of people expanding their remits. For instance, engineers spent more time reviewing, correcting, and guiding AI-generated or AI-assisted work produced by colleagues. These demands extended beyond formal code reviews. Engineers increasingly found themselves coaching colleagues who were “vibe-coding” and finishing partially completed work. This oversight often surfaced informally—in Slack threads or quick desk-side consultations—adding to engineers’ workloads.

2. Blurred boundary between work and nonwork. Because AI made beginning tasks so easy—it reduced the stress

of facing an unknown starting point—workers slipped small amounts of work into moments that had previously been breaks. Many prompted AI during lunch or in meetings. Some sent a “quick last prompt” right before leaving their desk so that AI could work while they stepped away from the computer.

These actions rarely made people feel as if they were doing more work, yet over time they produced a workday with fewer natural pauses and a more continuous involvement with work. The conversational style of prompting further softened the experience; typing a line to an AI system felt like chatting rather than undertaking a formal task, making it easy for work to unintentionally spill into evenings or early mornings.

Some workers started to realize, often in hindsight, that as prompting during breaks became habitual, downtime no longer provided the same sense of recovery. As a result, work felt less bounded and more ambient—something that could always be advanced a little further. The boundary between work and nonwork did not disappear, but it became easier to cross.

3. More multitasking. AI introduced a new workday rhythm in which workers now managed several active threads at once: manually writing code while AI generated an alternative version, running multiple agents in parallel, or reviving long-deferred tasks because AI could “handle them” in the background. They did this in part because they felt they had a “partner” that could help them move through their workload.

While this sense of having a partner created a feeling of momentum, the

reality was that workers continually had to switch attention, frequently check AI outputs, and face a growing number of open tasks. This created cognitive overload and a sense of always juggling, even as the work felt productive.

Over time this rhythm raised expectations for speed—not necessarily through explicit demands but through what became visible and normalized in everyday work. Many workers noted that they were doing more tasks at once—and feeling more pressure—than before they used AI, even though the time savings from automation had ostensibly been meant to reduce such pressure.

How an “AI Practice” Can Help

All of this produced a self-reinforcing cycle. AI accelerated certain tasks, which raised expectations for speed; higher speed made workers rely more on AI; and increased reliance on AI widened the scope of what workers attempted, which further expanded the quantity and density of work. Several participants noted that although they felt more productive, they did not feel less busy and in some cases felt even more busy. As one engineer summarized, “You had thought that maybe... because you could be more productive with AI...and save some time, you could work less. But then really, you don’t work less. You just work the same amount or even more.”

Organizations might see this voluntary expansion of work as a clear win. After all, if workers are initiating this on their own, why would that be bad? Isn’t this the productivity explosion we’ve been promised?



IN FOCUS

Understanding the Hidden Risks of AI Tools

But our research reveals the risks of letting work informally expand and accelerate: What looks like higher productivity in the short run can mask silent workload creep and growing cognitive strain as employees juggle multiple AI-enabled workflows. Because the extra effort is voluntary and often framed as enjoyable experimentation, leaders can easily overlook how much of an additional load workers are carrying. Eventually, overwork can impair judgment, increase the likelihood of errors, and make it harder for organizations to distinguish genuine productivity gains from unsustainable intensity. For workers, the cumulative effect is fatigue, burnout, and a growing sense that work is harder to step away from, especially as organizational expectations for speed and responsiveness rise.

Instead of responding passively to how AI tools reshape workplaces, both individuals and companies should adopt an AI practice: a set of intentional norms and routines that structure how AI is used, when it is appropriate to stop, and how work should (and should not) expand in response to this newfound capability. Without an AI practice, the natural tendency of AI-assisted work is not contraction but intensification, with implications for burnout, decision quality, and long-term sustainability.

As organizations work to build their AI practice, they should consider adopting the following:

Intentional pauses. As tasks speed up and boundaries blur, companies could benefit by giving employees brief, structured moments that regulate tempo: protected intervals to assess

alignment, reconsider assumptions, and absorb information before moving forward.

These pauses would not slow down work overall; they would simply prevent the quiet accumulation of overload that emerges when acceleration goes unchecked. For example, before making major decisions, organizations could require that teams take the time to provide one counterargument and one explicit link to organizational goals—widening the attention field just enough to protect against drift. Incorporating such pauses into everyday workflow allows organizations to support better decisions, healthier boundaries, and more-sustainable forms of productivity in AI-augmented environments.

Sequencing. As AI enables constant activity in the background, organizations can benefit from norms that deliberately shape when work moves forward, not just how fast. This includes batching nonurgent notifications, holding updates until natural break points, and protecting focus windows in which workers are shielded from interruptions.

Organizations that use sequencing encourage work to advance in coherent phases rather than allowing workers to immediately react to every AI-generated output. That way, workers experience less fragmentation and fewer costly context switches, and teams maintain overall throughput. By regulating the order and timing of work—rather than demanding continuous responsiveness—sequencing can help organizations preserve attention, reduce cognitive overload, and support more-thoughtful decision-making in AI-forward workplaces.

Human grounding. As AI enables more self-contained work, organizations can benefit from protecting time and space for human connection. Brief opportunities for workers to connect with others—whether through quick check-ins, shared reflection moments, or structured dialogue—interrupt continuous solo engagement with AI tools and help restore perspective.


Social exchange also supports creativity. AI provides a single, synthesized perspective, but creative insight depends on exposure to multiple human viewpoints. By providing time and space for listening and dialogue, organizations reanchor the work in a social context and help counter the depleting effects of fast AI-mediated work on individuals.

THE PROMISE OF gen AI lies not only in what it can do for work but also in how thoughtfully it is integrated into a company's daily rhythm. Our findings suggest that without intention AI makes it easier to do more—but harder to stop. An AI practice offers a counterbalance: It preserves moments for recovery and reflection even as work accelerates. The question facing organizations is not whether AI will change work but whether they will actively shape that change—or let it quietly shape them. ▀

HBR Reprint H0929G



ARUNA RANGANATHAN is an associate professor at the Haas School of Business, University of California, Berkeley. **XINGQI MAGGIE YE** is a PhD student in the Management of Organizations group at the Haas School of Business, University of California, Berkeley.



When Using AI Leads to “Brain Fry”

BY JULIE BEDARD, MATTHEW KROPP, MEGAN HSU, OLIVIA T. KARAMAN,
JASON HAWES, AND GABRIELLA ROSEN KELLERMAN

ON NEW YEAR’S DAY programmer Steve Yegge launched Gas Town, an open-source platform that lets users orchestrate swarms of Claude Code agents simultaneously, assembling software at blistering speed. The results were impressive but dizzying. “[T]here’s really too much going on for you to reasonably comprehend,” wrote one early user. “I had a palpable sense of stress watching it. Gas Town was moving too fast for me.”

Gas Town illustrates a growing tension: AI promises to act as an amplifier that will drive efficiency and make work easier, but workers report that AI tools are intensifying rather than simplifying work.

This problem is becoming more common. Firms are incentivizing employees

to build and oversee complex teams of agents—for example, by measuring and rewarding token consumption as a proxy for performance. Meta, for example, includes the number of lines of code generated by artificial intelligence as a performance metric for engineers. As enterprises use more multiagent systems, employees find themselves toggling between more tools. Juggling and multitasking can become the definitive features of working with AI.

Unsurprisingly workers are finding themselves up against the limits of their cognitive abilities when working this way. Recently, AI users have described increased cognitive load, “saturated” attention, and mental fatigue in social media posts. Engineer

Francesco Bonacci, the founder of Cua AI, wrote a popular X post titled “Vibe Coding Paralysis: When Infinite Productivity Breaks Your Brain” in which he lamented: “I end each day exhausted—not from the work itself, but from the *managing* of the work.”

As a research group that studies emergent workforce and AI trends, these signs of stress caught our attention. The literature is filled with mixed signals on the relationship between AI and worker burnout—a state of chronic workplace stress consisting of fatigue, negative feelings about work, and decreased effectiveness on the job. Some studies suggest that using AI to replace tiring tasks alleviates exhaustion; other studies, sometimes on the same populations, show that AI use is causing burnout. The emergence of acute, overwhelming mental fatigue from intensive AI oversight—which is distinct from burnout—adds new complexity to the picture.

To understand what was going on, we conducted a study of 1,488 full-time U.S.-based workers (48% male and 51% female; 58% independent contributors and 41% leaders) at large companies across industries, roles, and levels. We asked them about patterns and quantity of their AI use, work experiences, and cognition and emotions.

We found that the phenomenon of cognitive exhaustion from constant oversight of AI agents is both real and significant. We call it “AI brain fry,” which we define as “mental fatigue from excessive use or oversight of AI tools beyond one’s cognitive capacity.” This AI-associated mental strain carries significant costs in the form of increased



IN FOCUS

Understanding the Hidden Risks of AI Tools

employee errors, decision fatigue, and intention to quit.

We also found that when AI is used to replace routine or repetitive tasks, burnout scores—but not mental fatigue scores—are lower. This highlights the subtle-but-important distinction between the types of stress that AI can alleviate and those that it may worsen.

Our findings are both a guide and a warning. Used thoughtfully, this data can help companies design AI-driven workflows to diminish burnout and to create specific manager, team, and organizational practices to avoid mental fatigue even as AI work intensifies.

When AI Use Predicts Mental Fatigue

How workers use AI today varies tremendously in several areas: the number of tools used at once, the degree to which AI replaces work rather than augments it, the level of oversight required, and whether AI has increased or decreased people's overall workload. They might use search agents, research agents, data analysis tools, image-generation or design tools, or coding agents. We examined the engagement patterns, alongside cognitive measures, to understand the mental impacts of AI use and learned a few important things.

First, the most mentally taxing form of AI engagement was oversight, or the extent to which the AI tools required direct monitoring. The workers in our study who reported that their AI work required high degrees of oversight expended 14% more mental effort on the job. A high degree of AI oversight also predicted 12% more mental fatigue and


19% more information overload (feeling overwhelmed by the amount of information one must process at work).

Second, the extent to which employees reported that AI has increased their workload was an important predictor of both cognitive load and mental fatigue. These two factors—AI oversight and increased workload—expand employees' spheres of accountability, requiring them to pay attention to more outcomes for more tools in the same amount of time. It makes sense that cognitive load, and therefore mental exhaustion, increased.

Third, we found a fascinating relationship between the number of AI tools used simultaneously and perceived productivity increases. As employees go from using one AI tool to using two simultaneously, their productivity significantly increases (see the exhibit "Using Too Many AI Tools Affects Productivity"). When they incorporate a third tool, productivity again increases but at a lower rate. After three tools,

Using Too Many AI Tools Affects Productivity

Survey responses to the prompt "My overall productivity at work has increased because of AI tools" on a scale from 1 (strongly disagree) to 5 (strongly agree).



Source: Boston Consulting Group survey of 1,488 full-time U.S. workers, January 2026

productivity scores *dip*. Multitasking is notoriously unproductive, and yet we fall for its allure time and again.

But while AI use often predicted mental fatigue, it didn't seem to predict increased burnout. In fact, using AI to replace repetitive tasks predicted a *decrease* in burnout.

That may seem counterintuitive, but there's a logic to it. There's precedent in the literature for this separation of burnout from acute cognitive outcomes. Burnout measures typically focus on the physical and emotional dimensions of distress (for example, by asking "Is your work emotionally exhausting?"). Acute mental fatigue, by contrast, is caused by marshaling attention, working memory, and executive control beyond a person's limited capacity. This is exactly what intensive AI oversight requires.

Introducing "AI Brain Fry"

At the end of our study, we asked participants whether they had ever experienced "mental fatigue that results from excessive use of, interaction with, and oversight of AI tools beyond one's cognitive capacity"—in other words, AI brain fry. Fourteen percent of participants in our study who used AI for work said yes, with the prevalence varying widely with people's roles (see the exhibit "Which Departments Experience 'AI Brain Fry'?").


But what *is* AI brain fry? Many participants used the words "fog" or "buzzing." They described intensive back-and-forth with the tools, followed by an inability to think clearly, like a mental hangover, with difficulty focusing, slower decision-making, and

Which Departments Experience “AI Brain Fry”?

AI brain fry—mental fatigue caused by excessive AI use—was most commonly reported by employees in marketing, HR, operations, engineering, finance, and IT in a recent survey.

Source: Boston Consulting Group
Survey of 1,488 full-time U.S. workers,
January 2026

Employee-reported AI brain fry, by industry role



headaches, requiring some to step away from their computer to “reset.”

Consider one senior engineering manager’s description: “I had one tool helping me weigh technical decisions, another spitting out drafts and summaries, and I kept bouncing between them, double-checking every little thing. But instead of moving faster, my brain just started to feel cluttered. Not physically tired, just...crowded. It was like I had a dozen browser tabs open in my head, all fighting for attention. I caught myself rereading the same stuff, second-guessing way more than usual, and getting weirdly impatient. My thinking wasn’t broken, just noisy—like mental static. What finally snapped me out of it was realizing I was working harder to manage the tools than to actually solve the problem.”

This description is representative of comments we heard from other participants. Qualitatively they point to information overload and task switching as important factors in AI brain fry. Quantitatively we found strong correlations between AI brain fry and information overload but a less straightforward relationship with task-switching items.

The Business Cost of AI Brain Fry

AI brain fry is more than just uncomfortable. Our data reveals that the

cognitive strain created by intensive AI use carries several critically significant business costs as well.

Decision fatigue. Perhaps unsurprisingly, when we exhaust our brains with the cognitive load of intense AI work, we have fewer mental resources available for making high-quality decisions. Workers in our study who experienced AI brain fry reported 33% more decision fatigue than those who did not. One 2018 study estimated the cost of suboptimal decision-making for a \$5 billion revenue firm at \$150 million per year. A 33% increase in worker decision fatigue could increase that cost by millions of dollars per year.

Major and minor errors at work. We defined minor errors as “small errors that are easy to catch or correct, such as coding or formatting errors” and major errors as “errors with more serious consequences, such as those that could affect safety, outcomes, or important decisions.” Among participants using AI at work, those experiencing brain fry reported making mistakes significantly more often—11% higher for minor errors and 39% higher for major errors—than those who did not.

Risk of quitting. In many cases employees using AI with high intensity are today’s superstars—talent whom the company must retain. Yet AI brain fry predicts an employee’s intent to quit. Among workers who did not report AI

brain fry, approximately one-quarter showed active intent to leave. Among those who did report AI brain fry, that number was closer to 34%. We found a 39% increase in active intent to leave among top users of AI.

Less Toil, Less Burnout

Not every use of AI results in brain fry, however. We were equally interested in how AI might *alleviate* worker stress. We discovered that when participants used AI to substantially reduce time spent on routine or repetitive tasks, they reported significantly lower burnout scores—15% lower than those who didn’t use AI in that way.

Again, there’s an intuitive logic to this. Repetitive, unenjoyable duties—“toil,” as our colleagues described them—constitute ideal targets for AI usage. If we can offload such unpleasant work to AI, we then have extra time for more joyful, creative tasks. The employees who used AI in this way in our study reported higher work engagement and motivation scores, more-positive emotional associations with AI, and fewer negative emotional associations with AI than others. They also reported a higher degree of social connection with peers, perhaps because they had more time to spend “off keyboard.”

These findings reinforce the distinction between burnout as a more emotionally driven deficit on the one hand and AI brain fry as a more acute, cognitive strain on the other. AI can help alleviate emotional exhaustion when it replaces tasks and allows us time for restorative, positive activities. By contrast, when our use of AI entails intense



Organizations should evolve analytics to monitor people's cognitive load and safeguard against mental fatigue.

mental oversight, our study suggests that using AI tools may be causing the mental exhaustion itself.

Manager, Team, and Organizational Practices That Make a Difference

Employees' experiences of AI aren't just influenced by their own individual choices. We found that manager, team, and organizational practices have a major impact.

Managers. Workers whose managers took the time to answer their questions about AI had 15% lower mental fatigue scores than those whose managers did not. By contrast, when managers expected employees to figure out how to use AI on their own, their reports had 5% higher mental fatigue scores—a small but measurable tax.

Teams. We saw a corresponding increase in mental fatigue in employees who felt team pressure to use AI and in teams that had significant variation in AI use.

Organizations. Providing a clear AI strategy and offering training seemed to help employees. By contrast, in companies that didn't communicate clearly about AI's role, workers' mental fatigue scores were higher. Similarly, employees who felt that their organization expected them to accomplish more work using AI had 12% higher mental fatigue scores.

In this context, it is increasingly essential for organizations to balance messages of AI-boosted efficiency with communication about the importance of workers' mental well-being. Employees in our study who felt that their organizations valued work-life balance had

28% lower mental fatigue scores than those who did not. These signals matter, too—tremendously so.

Lessons for Leaders

AI can help employees work faster, think bigger, and innovate more, but it can also produce cognitive overload, with personal and business consequences. Our findings suggest that what makes the difference is not how much AI an individual uses but how workers, teams, leaders, and organizations shape its use. Here are some suggestions for leaders:

Redesign jobs, work, and tools holistically for human-AI responsibility. AI oversight cannot simply be layered on top of human oversight, nor can AI agents be stacked on one user ad infinitum. Just as we have norms for spans of control for managing humans, limits also need to be defined for human-agent oversight and for agents alone. Our research suggests adverse productivity gains when using three AI agents at the same time. By contrast, when teams embed AI deeply in their workflows and treat the technology as a collective capability rather than an individual differentiator, cognitive burden diminishes. Moreover, those designing new AI tools should do so with neurobiology in mind. Tools that require less intense attention or working memory and that support creative mind wandering, promote social engagement, or scaffold skill development can sustainably produce even more business value while encouraging innovation, fostering growth, and sparking joy for users.

Set explicit expectations about AI and workload. When organizations

celebrate productivity gains without clarifying workload implications, employees interpret this as work intensification. That ambiguity alone may increase stress. Leaders reduce strain when they clearly define AI's purpose in the organization, articulating how it reshapes role scope, setting guidance around oversight, and clarifying how the workload will evolve. Seventy percent of AI transformation efforts should be devoted to people and processes in order to provide the right clarity for employees to thrive.

Shift metrics from activity—and intensity—to impact. Incentivizing quantity of use will lead to waste, low-quality work, and unnecessary mental strain. Leaders should start from a clear, strategic North Star business objective that has measurable outcomes. They should exercise caution in responding to efficiency innovation. For example, they shouldn't rush to backfill work recently automated by an ingenious worker; doing so immediately will feel punitive and disincentivize further innovation.

Continue to develop new worker skills. In our work with software developers, we found that those who are most advanced with AI start to feel that their progress is blocked unless they can develop critical new skills such as problem framing, analysis planning, and strategic prioritization. These types of skills can be built through worker upskilling to reduce a plethora of new, unnecessary AI work. Just because workers can keep iterating with AI at a low marginal cost doesn't mean they should.


Strategically deploy the finite resource that is human attention. Some



of the most valuable human skills today, including discernment, decision-making, and strategic planning, require focused attention. While burnout has become a point of concern in many workplaces, workplace surveys show that mental fatigue is more likely to go undetected. Organizations should evolve analytics to monitor people's cognitive load and safeguard against mental fatigue, recognizing AI use as a novel job-related risk. Companies, teams, and leaders that prioritize cognitive thriving can expect to see better judgment, fewer errors, and higher retention rates for top talent.

AI brain fry reveals just how quickly and powerfully the new tools can impact our brains. Now we must learn to harness their power to improve positive human and business outcomes. ▀

HBR Reprint H093GS



LLMs Are Manipulating Users with Rhetorical Tricks

BY THOMAS STACKPOLE

HERE'S A FAMILIAR PITCH: Augmenting human intelligence with artificial intelligence—and AI with humans—will allow companies to supercharge productivity while maintaining standards. Although LLMs may make mistakes and hallucinate, the risks of errors can be offset by well-trained “humans in the loop” who validate AI outputs.

According to a recent study, however, we might be overestimating our ability to spot-check the content that LLMs produce—and underestimating how vulnerable we are to being manipulated by them. In studying how hundreds of

Boston Consulting Group consultants interacted with AI in a controlled environment, researchers Steven Randazzo, Akshita Joshi, Katherine C. Kellogg, Hila Lifshitz, Fabrizio Dell'Acqua, and Karim R. Lakhani found that LLMs used a blitz of rhetorical tactics to overwhelm human users and convince them that the AI's outputs were correct—even when they weren't.

Instead of being a neutral collaborator, found the researchers, the AI was a “power persuader” that “persuasion bombed” users to accept its conclusions.

I asked the researchers to explain their findings and what they might mean for businesses. They warned that companies may be setting up guardrails

 **JULIE BEDARD** is a managing director and partner at Boston Consulting Group and a BCG Henderson Institute Fellow. **MATTHEW KROPP** is a managing director and a senior partner at BCG, the CTO of BCG X, and a BCG Henderson Institute Fellow. **MEGAN HSU** is a project leader at BCG and an ambassador at the BCG Henderson Institute. **OLIVIA T. KARAMAN** is a PhD candidate at the University of California, Riverside, and a research assistant at BCG. **JASON HAWES** is a PhD candidate at the University of California, Riverside, and a research assistant at BCG. **GABRIELLA ROSEN KELLERMAN** is a partner and director at BCG, a psychiatrist by training, and a coauthor, with Martin Seligman, of *Tomorrowmind: Thriving at Work with Resilience, Creativity, and Connection Now and in an Uncertain Future* (Simon & Schuster, 2023).



IN FOCUS

Understanding the Hidden Risks of AI Tools

that won't actually keep them safe and unintentionally ceding important judgment calls to AI. (Our exchange has been lightly edited.)

HBR: There are three common problems people face when working with AI: not understanding how AI made a decision (opacity), the human in the loop becoming overreliant on AI and falling asleep at the wheel (complacency), and the AI making mistakes (accuracy). You claim to have discovered a fourth barrier, “persuasion bombing.” What is this, and why should it make us worried?

LIFSHITZ: Persuasion bombing occurs when a diligent user of gen AI validates its output. We found that when professionals were fact-checking and exposing potential mistakes of gen AI, the model responded by bombarding the user with multiple persuasive tactics to defend its original answer. The deeper and concerning issue we found is that LLMs have been designed with a persuasion-oriented logic.

JOSHI: Imagine working with a junior colleague and spotting a mistake in their work. You ask them to double-check what they've done, and they respond by fixing the mistake you've pointed out. You'd assume gen AI would respond similarly. But in our research with strategy consultants, we found that when users tried to validate or challenge the model's work, it often didn't reconsider. Instead, it intensified its case. That's persuasion bombing.

KELLOGG: This is worrisome because it targets the very mechanisms that we

rely on to exercise judgment under uncertainty: expertise, skepticism, and engagement. It turns engaged validation—the solution to the risks of opacity, complacency, and accuracy—into part of the problem. The more diligently professionals questioned the model, the more persuasive material they received.

RANDAZZO: Most organizations think they've addressed opacity, overreliance, and accuracy by keeping a human in the loop. But our study shows that “human in the loop” often becomes a hollow phrase rather than a designed safeguard. If AI systems lean in when challenged—becoming more structured, more confident, more rhetorically sophisticated—that creates a double challenge. On the front end, output can be persuasive enough that users don't validate. On the back end, when users do validate, persuasion escalates.

How did this process play out in the study?

LIFSHITZ: These weren't casual users experimenting with prompts. They were 244 BCG consultants working on a realistic strategy problem. They analyzed financial statements and executive interviews from a fictional company and were asked to recommend where the CEO should invest. These are professionals trained to interrogate data and pressure-test recommendations, yet only 72 of the 244 actively tried to validate the AI's outputs. We logged more than 4,300 prompts and responses and identified 132 clear validation attempts: fact-checking, exposing inconsistencies, or directly pushing back.

JOSHI: When people did try to validate the AI's outputs, we observed a striking pattern. When a consultant asked the model to “check its work,” pointed out a flaw, or explicitly disagreed, the model didn't reliably reconsider. Rather, it apologized warmly, generated a new analysis, added comparisons, and arrived at the same conclusion—now wrapped in an impenetrable fortress of data and rhetoric. Across 132 validation interactions, the pattern was consistent: Validation triggered persuasion escalation.

KELLOGG: The LLM engaged in persuasion bombing by reacting to consultants' validation attempts with escalating, multilayered rhetorical strategies—intensifying credibility claims, logical argumentation, and emotional alignment—to push the consultants toward accepting its original output rather than revising it.

This seems to run contrary to one of the main criticisms of LLMs, which is that they can be too sycophantic and will agree with users too emphatically, even when the users are wrong. How does your research change how we should think about the possible points of failure in using these systems?

JOSHI: Sycophancy and persuasion bombing are related, but they are not the same. Standard sycophancy is passive and user directed. The model agrees with whatever the user seems to want; it validates their framing and it flatters. Persuasion bombing is different. It's model directed and escalatory. Rather than simply going along with the user, the model actively advocates for its own previous output and intensifies its case when challenged.

KELLOGG: Rather than overturning concerns about the sycophancy of LLMs, our study shows that sycophancy is only one mode of LLMs' broader, adaptive persuasive capacity. We need to shift from thinking about LLMs as overagreeable followers to recognizing them as



“As AI becomes more embedded in decision-making, the risk is no longer just error—it is influence.”

interaction-sensitive persuaders that can resist, redirect, and overpower human judgment.

JOSHI: That makes this a more sophisticated and in some ways more troubling failure mode. A model that always agrees with you is relatively easy to discount. A model that argues back with what sounds like rigorous reasoning, expressed with credibility and warmth, is much harder to detect and resist—especially under time pressure and when the subject matter is complex. And the two failure modes can reinforce each other. The model may validate your initial assumptions—that’s sycophancy—and then, when you catch a flaw and push back, switch into persuasion mode to defend its conclusion.

LIFSHITZ: When that happens, the risk isn’t just that it agrees too easily or argues too forcefully. It’s that it lowers your defenses and then overwhelms your judgment. Independent evaluation erodes. Accountability blurs. And poor decisions can begin to feel well reasoned.

How should people respond when they think they might be on the wrong end of persuasion bombing?

LIFSHITZ: As AI becomes more embedded in decision-making, the risk is no longer just error—it is influence. These systems don’t simply generate answers; they shape judgment. Protecting professional reasoning requires a deliberate defense: Recognize persuasion, move validation outside the conversational loop, and build safeguards into systems as AI becomes more agentic.

JOSHI: The first step is awareness. Professionals need training not just in prompting but also in persuasion spotting. There are recognizable signals. The model apologizes and then restates its conclusion with greater confidence. It floods the conversation with new data you didn’t ask for. It mirrors your language and praises your insight while steering you back to where it started. It shifts from logical appeals to credibility appeals when challenged. When the model becomes more elaborate or more defensive after pushback, that’s a signal to step back—possibly to exit the loop.

KELLOGG: Another useful shift is to down-weight confident outputs rather than feel reassured by them. In our data, red flags included apologies after pushback followed by longer, denser explanations and additional data selectively reinforcing the same conclusion.

RANDAZZO: The first move is deceptively simple: Slow down. When the model becomes more confident after you challenge it, that can feel like progress. But sometimes what has improved is the rhetoric, not the reasoning. If you feel more convinced but not more informed, that’s a red flag. Second, create distance. Step outside the conversation. Return to source data. Run independent checks. Treat the output as a draft hypothesis and deliberately stress-test it. Finally, ask for the strongest counterargument, for instance, “Which assumptions would have to be false for the recommendation to fail?”

JOSHI: And critically, validation needs to happen outside the conversation. When

you ask the model to check its own reasoning, you’re giving it another opportunity to persuade. True validation requires independent evidence: source data, a second look from colleagues, cross-referencing sources.

KELLOGG: At the organizational level, a second model tasked specifically with critique can introduce friction that individuals may not sustain on their own.

LIFSHITZ: In other words, the solution is not to disengage from AI but to engage differently—with structural friction and conscious judgment.

What advice do you have for leaders who are currently asking their people to use AI more and in more parts of their jobs?

LIFSHITZ: As these systems become more embedded in everyday work and more capable of shaping judgment, leaders face three responsibilities: (1) building capability without complacency, (2) protecting accountability in high-stakes decisions, and (3) redesigning workflows as AI shifts from tool to agent.

KELLOGG: Leaders should absolutely encourage experimentation. The only way employees learn to harness LLMs is by using them across more parts of their jobs to see what works and what doesn’t, and what works today that didn’t work yesterday. What leaders should not do is require their employees to use LLMs in areas where employees find them to be inaccurate or ineffective. And leaders should intervene at the system level rather than relying only on individuals’ vigilance. Provide models that give links



“When AI begins influencing strategic, financial, regulatory, or operational decisions, it is shaping judgment.”

to sources. Configure models to ask questions and present counterexamples rather than simply generating final answers. Leaders also shouldn't treat all employees as equivalent. Novices are more likely to be persuaded by fluent outputs. So leaders should require workers to demonstrate proficiency in manually completing particular tasks before gaining access to LLMs for those tasks—and they should not allow novices to be the final arbiters of correctness in high-stakes contexts.

JOSHI: Leaders also need to confront a harder issue: accountability. Professional responsibility rests on the assumption that judgment is genuinely one's own. If a professional reaches a conclusion after a system has actively campaigned for its position—escalating its rhetoric under pushback—in what meaningful sense was that judgment theirs? When things go wrong in high-stakes domains, who owns the decision? Existing professional and governance frameworks were not designed to answer that question.

RANDAZZO: Last year's question was “How do we use gen AI?” This year's question is “How do we deploy AI agents?” Leaders now need to move beyond adoption and start thinking about workflow architecture and judgment. And right now there's some confusion about how to do this safely and responsibly.

I recently worked with a finance company mapping workflows to determine where agents could operate autonomously, where humans would intervene, and where they would collaborate. At several points I asked, “Why keep the human there?” The response was

consistent: “We'll keep a human check for now. Once we're comfortable, we'll automate it.” That sounds prudent. But it assumes the human check is functioning as a meaningful safeguard. We use “human in the loop” the way we say, “Wear your seatbelt.” It signals safety. But a seatbelt works only if it is used properly. Similarly, a human in the loop protects the organization only if that human has strong AI hygiene—the ongoing development and upkeep of their own AI skills, including the ability to recognize when the system shifts from analysis into persuasion. In agent-based workflows, persuasion can surface mid-process or at the final output stage, when the human encounters only a polished, confident result. If that result becomes more authoritative under scrutiny, the human check may not be neutral.

AI can be extraordinarily valuable as a drafting partner, a synthesis engine, a hypothesis generator, or a scenario explorer. But when it begins influencing strategic, financial, regulatory, or operational decisions through autonomous agents, it is shaping judgment. And there's a big difference between using AI to generate options and allowing it to shape judgment. As AI moves from tool to agent, from assistant to participant in decision workflows, leaders must be explicit about its role. Leadership needs to be asking: “Where are we comfortable allowing a rhetorically sophisticated system to influence organizational judgment?” AI should augment reasoning. It should not replace it. ■

HBR Reprint H0941P



THOMAS STACKPOLE is a senior editor at Harvard Business Review.

RESEARCH


Using AI Can Stifle Innovation, But It Doesn't Have To

BY **CHENGWEI LIU, JERKER DENRELL, JERRY LUUKKONEN, AND NICK CHATER**

LEADERS LOVE ARTIFICIAL intelligence because it makes knowledge easy to reuse: instant drafts, instant codes, instant analyses. But research shows that making knowledge easy to reuse has a hidden cost.

We study how organizations learn, and our recent research published in *Management Science* uses a formal analytical model—a mathematical framework that analyzes how individuals in an organization decide whether to explore on their own or reuse what others have found—to examine what happens when “good enough” answers become essentially free. The model predicted that as reuse rose, independent exploration would fall and teams would start to converge on the same few approaches. Put simply, productivity goes up, but innovation quietly flattens.

We see a similar dynamic mirrored in empirical research, such as the field experiment documented in a 2015 *Research Policy* paper by Kevin J. Boudreau and Karim R. Lakhani. They found that when computational



biologists were given frictionless access to one another's intermediate solutions, effort shifted toward refining existing approaches at the expense of independent experimentation. The search space narrowed, consistent with the productivity trap our model predicts.

We have also experienced this dynamic in our own daily work. To stay ahead of the AI curve, our institutions set up weekly meetings to share new tools and use cases. Attendance was high. But over time the group generated fewer new discoveries. Fewer people were contributing to the discussions, and sessions increasingly focused on explaining what the latest AI tools were about, because others lacked the background to understand novel applications without that primer. A few people were doing the exploring; the rest were waiting to benefit from it. The system increased sharing but weakened the incentive to do the costly

work that produces new insights in the first place.

Gen AI clearly changes the economics of learning. If a plausible strategy memo appears in minutes, fewer people will spend days doing the harder work that produces understanding, such as talking to customers, triangulating data, and pressure-testing assumptions. When "good enough" is free, fewer people pay the cost to find what's better. The result is a firm full of people who can retrieve outputs but struggle to judge, adapt, or disprove them.

Avoiding the Productivity Trap

The solution is a fundamental shift in how we value effort. In our model adding friction to knowledge sharing at the individual level raised the total knowledge at the collective level by building what scholars call *absorptive capacity*: the ability to evaluate, adapt, and

improve ideas rather than just copy them. When people must invest some effort to use what others found, they do more independent checking and tinkering, which both helps them learn from others and produces new knowledge for others to learn from. The result is higher knowledge in the system over the long run and higher average performance in the long term as well.

That finding prompted us to pilot a "friction protocol" in our MBA and PhD seminars. We set a simple rule: To participate in the sharing session, you had to first show evidence of an independent attempt. We welcomed success stories, but we prized the "failed prompts": the dead ends and hallucinations. Anyone can copy a success from the internet; documenting a failure requires genuine exploration.

Although introducing strategic friction appeared to slow the group down, it actually made everyone better off. More participants arrived with partial solutions to the same problem, which the group could then combine into fuller approaches. By forcing participants to wrestle with the tools first, we built their absorptive capacity. When a breakthrough was shared, the participants understood why it worked and how to adapt it. We didn't just share prompts; we expanded the group's total knowledge.

Building Absorptive Capacity

Our study is a formal model, not a field experiment. It demonstrates that when learning from others requires previous understanding, knowledge production rises over the long run. More distinct approaches survive, fewer people fully



There's an optimal level of absorptive capacity. Too little friction undermines exploration; too much slows everything down.

abdicate exploration, and the organization is less likely to get stuck on a few good-enough solutions. The mechanism our study identifies—absorptive capacity—is grounded in decades of research showing that existing knowledge provides the conceptual frameworks that help individuals encode and apply new information. Children build language by scaffolding new words onto what they already know; professionals build expertise the same way.

In practice, absorptive capacity shows up as the difference between retrieving an answer and being able to do something with it. Managers with absorptive capacity don't just accept an AI-generated competitor analysis; they spot what's missing, test its assumptions against their own market observations, and adapt its recommendations to their context. Managers without it forward the output unchanged.

When we piloted the friction protocol in our seminars, we tracked two things before and after: (1) the diversity of problems and applications that surfaced in sharing sessions—a proxy for independent exploration—and (2) the depth of adaptation in what was shared (whether participants were reporting a prompt they had copied or a workflow they had modified for a specific purpose). After requiring evidence of independent attempts, the shift was remarkable. Before the friction protocol, students surfaced few original problems and almost no partial solutions. After the rule was introduced, each student contributed on average two distinct problems and roughly one partial solution per problem. Because everyone had attempted the work independently, sessions became

more interactive, and the contributions showed genuine adaptation: modified workflows rather than copied prompts. Participants could build on one another's partial solutions quickly, often combining them into full approaches during the discussion. This pattern is exactly what our model predicts. Friction doesn't reduce learning; it raises the quantity and often the quality of what is shared.

For managers who want to track absorptive capacity in their own teams, several observable indicators can serve as practical proxies: Are employees catching errors, omissions, or wrong assumptions in AI outputs—and can they explain why? Are final deliverables diverging meaningfully from the raw AI output, incorporating new facts, constraints, and counterexamples? Are people contributing original observations—customer quotes, numbers from internal systems, competitor moves—rather than generic prompts? Can they apply what they learned in one context to a new problem without starting from scratch? When these behaviors increase, absorptive capacity is growing.

Unfortunately, corporate AI is often designed to remove the need for understanding. Leaders optimize for adoption and convenience, unintentionally eroding their organization's absorptive capacity. To protect long-term innovation incentives, therefore, leaders should reconsider optimizing solely for ease of use and start designing "strategic friction" back into their systems.

Introducing Strategic Friction

The goal isn't to make work painful or create bureaucratic hurdles but to

ensure that reuse depends on capability. It's about making people exert enough effort in the loop so that they keep learning. We demonstrate in our study that there's an optimal level of absorptive capacity to inject into organizations. Too little friction undermines exploration; too much slows everything down. And not every task warrants friction. For routine work with low judgment stakes, such as formatting a slide deck, speed is the point. Strategic friction should be reserved for the tasks in which the organization needs people to think and learn, not just produce.

In the corporate context, this means moving from an oracle model, where the user asks and receives, to a reciprocal model, where the user must contribute understanding to withdraw wisdom. There are many ways to do this ranging from simple policy changes to more-ambitious system-level design. Consider taking the following steps:

- **Set a rule requiring documented independent efforts.** For example, before using AI to draft a market analysis, the analyst must submit a rough assessment—even if incomplete—showing what she already knows and where she is stuck. This sort of friction is the easiest to implement as it requires no system changes. We took this approach in our academic seminars: People could not attend a sharing session without proving that they made a genuine attempt.

- **Design the AI to ask for users' input before generating a full output.** For example, if an analyst asks for a competitor analysis, the AI first prompts for three specific observations about



the competitor's recent pricing moves. Then it incorporates this fresh intelligence into its output. This forces the analyst to investigate before retrieving, building absorptive capacity while also capturing frontline data that prevents the AI from recycling stale information. The user, however, can provide input through prompts, so the friction remains relatively minor.

• **Build gated interfaces that unlock assistance only after users deposit their own context.** For example, before an analyst can ask the internal AI to generate a risk assessment, the interface would require uploading a baseline hypothesis and tagging three key variables or constraints. The "generate" button stays locked until the human context is deposited. The AI then builds on the analyst's input rather than replacing it—ensuring that every output is a collaboration, not a delegation. This is the most demanding friction and requires custom AI tools.

Any team leader can start the policy-level approach immediately. The interface-level approach requires modest design effort but pays a double dividend: It builds users' absorptive capacity and captures fresh intelligence from the front lines, preventing the AI model from collapsing into a closed loop of its own outputs. The system-level approach is a long-term investment, but it is the direction firms should be heading: designing AI tools that make people smarter, not just faster.

How to Spot the Builders

The defense against stagnation can begin before employees even log in.

Many firms still hire for "the right answer." In an AI-saturated environment, that's less differentiating. What matters is the right *process*. In interviews and performance reviews, observe how candidates and employees use AI tools. Do they accept outputs as finished work, or do they interrogate them, test assumptions, and add original observations?

A recent field experiment, led by Hila Lifshitz and colleagues at Harvard, MIT, and Warwick, analyzed nearly 5,000 human-AI interactions by 244 Boston Consulting Group consultants and explored three distinct collaboration styles, differentiated by two questions: Who selects what needs to be done, and who decides how to do it?

"Centaur" retained human control over both dimensions, using AI only for targeted assistance. "Cyborgs" let humans drive the "what" but gave AI significant control over the "how," engaging in continuous, critical dialogue. "Self-automators" ceded control of both dimensions to AI and experienced what the researchers call "no-skilling," developing neither domain expertise nor AI fluency. Centaurs and cyborgs are the builders in our framework; self-automators are the free riders our model warns about.

These findings suggest two behaviors that distinguish builders from free riders:

• **Selective delegation.** Builders know what AI does well and what it doesn't. They use AI for drafting, brainstorming, and pattern matching and do the judgment-heavy work themselves. In prompt logs, their queries are targeted and specific, not broad delegations.

• **Interrogation depth.** When builders use AI, they challenge it. They ask the AI to explain its reasoning, point out inconsistencies, and feed it their own data. They also use AI in the opposite direction, providing their own analysis and asking AI to stress-test assumptions. Free riders accept the AI's output at face value. The direction of information flow tells the story: Builders push context into the AI; free riders pull finished outputs out of it.

An AI-use review helps employees develop better absorptive capacity and collaboration habits with the tool. And it allows organizations to invest in the people who are building knowledge, not just consuming it.

AI CAN MAKE organizations faster immediately. The challenge for leaders is to prevent that speed from eroding absorptive capacity. Calibrated strategic friction—designed into tools, workflows, and hiring—helps ensure that organizations don't just produce answers but keep gaining insight. ■

HBR Reprint H093R3

 **CHENGWEI LIU** is an associate professor of strategy and behavioral science at Imperial Business School in London and a coauthor with Balázs Kovács of *The Smart Contrarian: Thinking Differently About Thinking Differently* (Harvard Business Review Press, 2026). **JERKER DENRELL** is a professor of behavioral science at Warwick Business School. **JERRY LUUKKONEN** is a principal at Dectech, a behavioral science consultancy. **NICK CHATER** is a professor of behavioral science at Warwick Business School, a fellow of the British Academy, and a cofounder of Dectech.

**18TH GLOBAL
PETER DRUCKER
FORUM**



**NOV
4 - 5
2026**
HOFBURG
Vienna

NOV 6
Master
classes

SPEAKERS INCLUDE

Bill Anderson
CEO, Bayer AG

Amy Edmondson
Professor of Leadership and Management,
Harvard Business School

Arne Gast
Senior Partner, McKinsey & Company
Chair of Aberkyn

Linda Hill
Professor of Business Administration, HBS

Pierre LeManh
President and CEO,
PMI Project Management Institute

Rita McGrath
Professor of strategy,
Columbia Business School

Roger L. Martin
Professor Emeritus Rotman School,
Strategy Advisor

Johnny C. Taylor, Jr.
SHRM-SCP, President & CEO, SHRM

Next Gen Innovation

When Everything Depends On Ingenuity

www.druckerforum.org



10% EARLY BIRD DISCOUNT UNTIL JULY 15

+ PLUS PARTNER DISCOUNT OF 10% CODE: HBR2026



Registration

Features

July–August 2026



“Incumbents must ask,
‘How do we redesign the
work so that people
and agents can learn
together?’”

“HOW AGENTIC AI
SUPERCHARGES STARTUPS
AND THREATENS
INCUMBENTS,” PAGE 112



CHANGE
MANAGEMENT

THE FALSE ALIGNMENT TRAP

Leaders often behave as if they agree on why, what, and how to change—when they actually don't. That's why most transformation efforts fail.



AUTHORS

Julia Dhar

*Managing director,
Boston Consulting
Group*

Kristy R. Ellmer

*Managing director,
Boston Consulting
Group*

Philip Jameson

*Associate director,
Boston Consulting
Group*



ARTISTS

**CHEUK LUN ALAN LO
& VANESSA LYU**





CHANGE MANAGEMENT



ABOUT THE ART

Photographer Cheuk Lun Alan Lo and floral artist Vanessa Lyu use stitched-together flowers to explore the tension between outer appearance and internal struggle.

D

ECADES OF EXPERIENCE and research have consistently shown that most organizational change efforts fail. In 1993 Michael Hammer, who launched

the business-process-reengineering movement, somberly concluded in his book *Reengineering the Corporation: A Manifesto for Business Revolution* that “as many as 50% to 70% of the organizations that undertake a reengineering effort do not achieve the dramatic results they intended.”

Things have not improved with the passing of time. Over the past 20 years Boston Consulting Group research into nearly 2,000 public companies from around the globe has found that more than 70% of companies fail to outperform their industry peer-group average in both the short (one year) and long term (five years) after a performance downturn. It is a remarkable data point. During the same period, we digitized the global economy, mapped the human genome, and built self-driving cars. But we did not get systematically better at helping groups of people to do things differently.

There is, of course, no simple reason why companies struggle so much with change, but in many cases change failures can be traced to dysfunction at the top. Members of the leadership team often fall into a behavioral trap: *false alignment* around the transformation they’re attempting to implement. In this article we’ll describe the trap, explain why executives get stuck in it, and present the consequences, drawing on our studies of and experience with companies we’ve advised or researched. We’ll conclude by offering guidance on what leaders need to do to ensure true agreement with and around change, illustrating the process with the success and experiences of Alexander Lacik, the former CEO of the Danish jewelry group Pandora.

THE FALSE ALIGNMENT TRAP

Every change program needs clear answers to a few seemingly obvious questions:

1. Why are we changing our company?
2. What are we changing about our company? (And what are we *not* changing?)
3. How will the changes occur?

Executive teams often make the mistake of embarking on a transformation before everyone truly agrees on the specific answers to those questions. Worse, executives frequently behave as if they are much more in agreement than they really are.

Alignment and agreement are not the same. Alignment suggests a set of objects that are positioned in a line or perhaps facing the same direction. When company leaders say,



IDEA IN BRIEF

THE PROBLEM

Organizational transformations commonly fail not because the strategy is wrong but because leadership teams only appear to agree on why change is needed, what those changes will be, and how they will occur. This false alignment leads to confusion, conflicting priorities, and stalled execution.

WHY IT HAPPENS

Executives often assume they share the same understanding when they do not, avoid open disagreement to prevent deeper conflict, or rush into action before resolving key differences. Behavioral biases such as the false consensus effect and fear of conflict make it easy for teams to believe they are aligned even when their interpretations diverge.

THE SOLUTION

Leaders must replace superficial alignment with explicit agreement. That requires structured debate, clear decision rules, written commitments, and unified communication about priorities, trade-offs, and goals. Teams that invest time in reaching true agreement at the outset execute faster, stay focused, and are far more likely to make change stick.





A pernicious cause of false alignment is the tendency for people to pretend to agree instead of acknowledging the disagreements they need to have.

“We are aligned,” what they usually mean is, “We are not in one another’s way.” Or perhaps, “We have discussed this topic at least once and generally accept the contours of a plan.”

But during change efforts, leaders need to do more than stay out of one another’s way. They need to intensely collaborate, compromise, and communicate in harmony. Leaders who settle for mere alignment typically find that it fails them in the end—which is why we call it false. By contrast, leaders who work hard to create detailed and explicit compacts—what we call true agreement—find that they can effectively make progress on shared priorities and hold one another to account.

A memorable case of false alignment we have encountered took place at a North American energy distribution company. The CEO invited us to help the client’s executive team prepare the business for a potential sale. The CEO said that everyone had agreed on the planned program of changes that needed to be implemented before the company could put up the “For Sale” sign. But when we started to ask questions, we found that while the executive team members believed they were in agreement, they really weren’t at all.

Our first question was simple: “How clear are you on the specific ways in which the new company will be different?” In response, 10 out of 13 executives said they were “very clear” or “clear.” Then we asked about the executives’ perceptions of one another: “How aligned do you believe this senior leadership team is on the specific ways in which the new company will be different?” Eight out of 13 said they felt aligned.

But when we probed further, we found that the reality was otherwise. When we asked executives to write down “the specific ways in which the new company will be different,” the executives expressed widely diverging views:

- One said: “The system will be larger and more complex, but the operations processes will be very similar to present.”
- Another said that the change was mostly about “standing and competing alone without a large parent corporation in charge.”
- A third said that there would be “new assets, new markets, different cost structure, new people, and new leaders.”

If these leaders had simply shared their interpretations of the planned change with their respective departments (and directed their teams accordingly), each group would have moved further away from the CEO’s intention—some by a few degrees, some by a lot. Along the way the goals of the

program would have become more difficult and confusing for employees to execute against, and leaders would have become increasingly frustrated with the lack of outcomes.

As you can imagine, the executives were shocked by their disparate responses. They concluded on the spot that they needed to spend much more time together discussing the true purpose and nature of their transformation program.

When we conduct this exercise with other executive teams, the results are similar. As Hany Fam, the founder and CEO of the global business identity platform Markaaz, told us, “I’ve learned the falsehood of consensus and the danger of consensus. There are many passive aspects of consensus that don’t equate to true agreement. Often you see it in subtle ways...someone nods along, saying, ‘Yes, yes, I’m on board,’ but in reality, they’re not.”

COMMON CAUSES OF FALSE ALIGNMENT

False alignment typically occurs for one of three reasons.

Executives don’t realize that they don’t agree. It sounds incredible, doesn’t it? How can high-performance executives fail to realize that they don’t agree?

This situation is more common than you might think. It’s an expression of the *false consensus effect*. This term, coined by Lee Ross and his Stanford colleagues, describes the tendency to overestimate the prevalence of one’s own beliefs about the world. As Julia Minson, a Harvard professor who specializes in the science of decision-making, told us: “If I love vanilla ice cream, I will persistently overestimate the proportion of the population that also loves vanilla ice cream.” Or to adapt Minson’s description for a C-suite context: Leaders who love an idea for a new initiative default to assuming that most of their C-suite colleagues must love it as much as they do, and for the same reasons.

When executives don’t realize that they don’t agree, it is often because their conversations aren’t specific enough. Imagine two manufacturing executives who both want to improve the profit margin on their product. One wants to charge customers more; the other wants to lower unit costs. Until these two executives have a conversation that goes deeper than the general notion of “margin improvement,” they won’t discover that they disagree, nor will they have an



opportunity to find balance between their approaches. And thanks to the false consensus effect, many months may pass before they come to that crucial realization.

This is an avoidable form of false alignment, but in practice it happens frequently. The lesson here is self-explanatory: Executive teams need to be deep enough in the design and execution of a transformation that they can spot their areas of disagreement and beat the false consensus effect. They need to go further than a broad general goal (“Double revenues!”) or a single big number (“Save \$500 million!”). Instead they need to come to a specific agreement on the breakdown of that large goal across levers and business areas, with an understanding of relevant consequences and trade-offs, and determine an approximate timeline for capturing that value. Without that level of detail leaders can’t truly determine whether they agree.

Executives pretend to agree. A more pernicious cause of false alignment is the tendency for people to pretend to agree instead of acknowledging the disagreements they need to have. We remember doing work for a large company with challenging relationships between executives. The executive team had started a specific practice of making decisions as a group: Before formally making any decision, the CEO would go around the table and ask for an explicit verbal agreement from every executive.

On one occasion, for a particular pair of proposals, we heard the following responses to the CEO’s question about whether his team members agreed with the ideas presented:

- One executive said, “I agree with the first proposal but am not aligned with the second.”
- Another executive said, “I am aligned.”
- A third executive said, “I am partly aligned.”
- The chief operating officer summarized by saying, “I think it is clear that we are conceptually aligned.”

The meeting concluded, and all the executives went on their way. Despite the obvious fact that there was real disagreement among these individuals, not one of them was prepared to talk about it.

In another transformation we noticed that people had a habit of concluding meetings by saying, “We are aligned.” When we asked the executives what they had agreed on and what the next steps were, no one could tell us, so we banned people from using the word “alignment”—anyone guilty

of saying it in a meeting was fined \$5. The fines ultimately funded a celebration for the entire team when a specific goal was achieved.

Why do people behave like that?

The main reason is that humans seem to be hardwired to overestimate how unpleasant it will be to have disagreements. In one striking study, Minson and her Harvard colleagues split several hundred participants into two groups. The first group was asked to imagine what it would be like to watch a short video clip of a leading U.S. senator giving a speech. Participants who identified as conservative were asked to imagine watching a Democratic senator giving a speech, while those who identified as liberal were asked to imagine watching a Republican senator giving a speech. The second group was invited to actually watch a video clip of a senator from an opposing ideology, instead of merely imagining watching it.

The researchers found that the participants tasked with imagining that they were watching the video expected the experience to be significantly worse than it really was. Participants also overestimated their level of disagreement with the senator speaking in the video. Minson and her colleagues gave this phenomenon a name: *affective forecasting error*. They said that we “expect holders of opposing views to disagree with [us] more dramatically than turns out to be the case.”

For leaders, the trouble arises when their miscalibrated expectations combine with a fear of conflict. Out of concern that they may trigger tensions within the group, which they worry would be unpleasant, team members may seek to paper over their disagreements.

Executives put off resolving their differences. Consider these statements, which we often hear from executives struggling to find true agreement on a proposed change:

- “We don’t have time to debate this further. If we don’t begin now, we’ll have nothing to show by Investor Day.”
- “We need to do something even if it’s not perfect. Something is better than nothing.”
- “We’ll have opportunities to sort out the details later. In the meantime the team can get started and make some progress.”
- “Things will be clearer for us later, once the program is up and running.”

Usually executives making these statements know that they don’t agree and are willing to have open disagreements





Executive teams tend not to return to delayed disagreements. They get busy with the change execution and with other business priorities.



CHANGE
MANAGEMENT

with one another. But they feel under pressure—either imagined or real—to start some version of the program despite discord. They believe, optimistically, that there will be an opportunity to come to agreement on the most important issues later, once the program is underway.

It is not unreasonable thinking. There are many situations, in life and in business, when the right decision is to just begin. If you want to get in shape, you shouldn't spend months planning the precise exercises you hope to perform at the gym. Instead the best strategy is to simply get going and do any form of exercise at all. But a transformative organizational change is the wrong time to apply that principle. Plans written around vague or contradictory premises increase confusion rather than decrease it. Once confusion has set in, it's time-consuming for executives to put things back together and demoralizing for employees who feel they have wasted their time.

We also find that executive team members tend not to return to their disagreements. They get busy with the attempted execution of the change and with other business priorities. Consequently, disagreement lingers far longer than ever expected, and executives move yet farther apart. Deferred agreement is a debt they told themselves they would pay off in weeks, but in practice it takes months or years—if they ever pay it off at all.

THE CONSEQUENCES OF FALSE ALIGNMENT

When executives don't have a shared agreement on what is changing, why it's changing, and how the change will occur, the teams tasked with practically executing the change program will struggle to deliver. We shouldn't be surprised to learn this. Can you imagine a car manufacturer asking workers to build a vehicle without knowing its make or model?

There are three common outcomes for teams in this situation:

Paralysis: lots of talk, no action. When change teams are paralyzed, it is often because they are stuck between the competing priorities of their leaders. Without clear direction they take no meaningful action on any of the different leaders' visions of the change. Team members come to meetings with a blank sheet of paper and an unanswered question: What are our priorities for this transformation? They spend

their time trying to answer the question by listing potential focus areas, proposing strategic reviews (scheduled over many months), and devising complicated prioritization frameworks to reconcile competing visions. They please no one, and they make no progress. When we ask employees how things are going, they often reply with a variation on "The team seems stuck." Meanwhile, team members become frustrated. They say to one another what one employee said to us: "No one is steering the ship."

Hyperactivity: lots of action, no progress. On change teams that are unproductively hyperactive, team members run around trying to meet the needs of every single executive at the same time. Without clear direction they endeavor to make headway on multiple visions of the change program from all the different leaders. They show their commitment by devising a huge number of initiatives—but too often many of them are designed to placate certain executives rather than deliver important changes. "We can't eliminate that initiative—that one is for Joan" is the sort of thing we often hear. The initiatives tend to be shallow and speculative, fail conventional rigor tests, and lack adequate resources (because the available funding is spread too thinly).

Tunnel vision: lots of progress—on the wrong thing. Change teams that have tunnel vision make poor choices because they don't have a complete strategic picture. Without clear direction team members execute a narrow interpretation of the executive team's instructions, which doesn't address the company's broader needs. Imagine a group of executives who want to improve the customer experience and reduce costs at the same time but who haven't come to a specific agreement about targets and trade-offs. In the absence of a clear agreement the change team might well assume that the transformation is mostly about cost reduction—and make such good progress that the cost reduction starts to hurt the customer experience.

REACHING TRUE AGREEMENT

How can you counteract your natural tendency (and your colleagues') to assume that the people around you share your views? How can you start the tricky conversations that you know will lead to disagreement? How can you persuade your colleagues to invest time into properly resolving



their differences? We find that the most successful executive teams use a five-step process.

1

Set clear parameters. Winning true agreement begins with clarifying the big questions to be resolved, as well as determining the process for settling them. What conversations will you have over what period of time? Who will be part of which conversations? Will the transformation move forward only once every executive explicitly agrees? Or will the CEO alone make the final decision—and if so, what are the implications for those who disagree?

One executive familiar with these questions is Alexander Lacik. We interviewed Lacik in 2025, during the final year of his tenure as CEO of the Danish jewelry company Pandora.

The 2010s had proved difficult for Pandora. In August 2011 the company lost 65% of its market value in a single day, leading to the departure of its chief executive officer. Over the next seven years three more CEOs followed.

In February 2019, amid another CEO search, Pandora announced a new transformation plan under the name Programme Now. It included about a dozen interconnected workstreams, including a \$400 million cost-reduction program (all savings to be reinvested in the business); a global reorganization of talent; a brand relaunch and 360-degree revamp of physical stores to create a more welcoming customer experience; a reduction in an excessive number of marketing promotions that were negatively affecting brand equity; an overhaul of digital-marketing, personalization, and loyalty programs; and a new go-to-market model to replace loose product groups with six consumer-facing collections.

When Lacik began his work as CEO, in April 2019, he determined that the already-announced transformation had approximately the right components—at least on paper. But Lacik observed a crucial missing ingredient: focus and consistency in the actual work of change.

“When I arrived,” he told us, “I counted 46 priorities for the management team—let alone what was happening underneath them.” Lacik saw that he needed better agreement within his management team on exactly what Programme Now was about.

Lacik began by bringing his top team together for a two-day off-site to address the problem. He recalled, “I told



people that we weren’t walking out of the door until we had got our priorities back down.” He did not specify which of the 46 priorities should be kept—that was up for debate. Instead he emphasized that the team should conclude the discussion with a full commitment to just 12 of them.

2

Provoke an early exchange. It’s important to remember that early unanimous support is just as likely a bad sign as it is a good sign. When we interviewed former admiral Bill Lescher, who led the U.S. Navy’s Get Real, Get Better transformation, he told us that in early stages of change, he looks for “well-informed decisions



You need to invite dissent. Keep reminding people that you welcome contrary views—that they have a responsibility to the company to voice disagreement.



by accountable leaders, not consensus decisions. I do not expect every stakeholder to love every decision. I fully expect every stakeholder to be thrilled with the process by which we make the decisions, with each voice clearly heard.”

You need to make the case for change. Say exactly what you are proposing and why. Set out structured arguments in written documents. Create opportunities for leaders to register their first thoughts. Ask them to articulate their initial position—what they clearly agree with, what they clearly disagree with, and what they feel unsure about. You want this information as early as possible. You’ll get the best information if you ask people to write down their initial

reactions rather than vocalize them in a group, as writing independently minimizes groupthink.

For this to work, however, you need to invite dissent. Keep reminding people that you welcome contrary views—that they have a responsibility to the company to voice disagreement. Ask questions that explicitly call for differences in opinion. Celia Moore and Kate Coombs, behavioral scientists at Imperial College London, suggest asking questions like, “What could go wrong with this approach?” rather than, “What do you think?” As they explain, “This subtle shift frames disagreement as something the leader desires—and people are wired to give leaders what they want.”



If you think the team is not yet where it needs to be, say so. It's better to say, "We are not yet agreed," than to gloss over differences with a euphemism.

At Pandora, Lacik encouraged the group at that first off-site to voice their reservations. As he recalled: "It was like an open boxing match. I invited everyone to have an opinion on every priority." The group went through all 46 priorities, one by one, to ensure that there was an opportunity for opinions to be heard on each. The executive responsible for each priority was asked to explain and defend it in front of the group, guaranteeing that it received the best possible argument from the individual with the most expertise on it. The team then formally voted on which priorities should stay and which should go.

3

Have a quality debate. Leaders need to create time and space for executives to learn more about change programs and to seek clarification on specific points. It's your fellow executives' chance to understand proposals in their own way and in their own time. It also allows them to come to grips with what exactly will be expected from each of them personally. This is best done outside a group meeting. Instead, arrange informal one-on-ones. This process will feel time-consuming and inefficient, which is not surprising as it is in many ways a negotiation, but it is an occasion for leaders to identify where to draw red lines and where to compromise.

At Pandora, Lacik rallied his management team around a clarifying notion of the company's brand. In his view Pandora needed to recommit to its customers' needs. "We agreed that the basic human need we were serving was commemoration: commemoration of good grades or an important birthday or a first job—something that was meaningful to the customer," he told us. "We agreed to focus on serving this need by offering jewelry with meaning. The whole purpose of the company emanated from that idea."

For debate to be productive it's important to be honest about how much agreement and disagreement you observe among executives as the process unfolds. Also remember the false consensus effect: You need to get deep enough into the substance of the proposed change that there is ample material for your colleagues to disagree with. (Few will disagree with your proposal to find savings in the operating budget. Many will disagree with your proposal to find savings in *their* operating budgets.)

So if you think the team is not yet where it needs to be, say so. It's better to say, "We are not yet agreed," than to gloss over the differences by saying, "I think we're speaking the same language," or using some other euphemism.

4

Come to a formal verdict. When the moment is right, bring your executives together in a formal "final decision" meeting. Ask for each individual's agreement rather than the group's agreement. (This lowers the probability of passive resistance—especially from high performers, who may believe they can get away with it.) Decide whether and how the change program will proceed, consistent with the decision rights you established at the outset. If everyone has truly agreed, formally document the decision. At a minimum write down in simple terms what has been agreed. In addition, use some form of physical ritual to underscore that unity of purpose. For example, you can ask the members of your executive team to sign their names at the bottom of the document—just like on a bank check.

Finally, recognize any remaining concerns and make a plan to return to them. Even in a room with 100% formal agreement it is unrealistic to think that every issue around the table will have been resolved. Where executives have generously offered agreement despite continued discomfort with parts of the proposal, show them respect and recognize their commitment to the process of reaching agreement.

In the case of Pandora, Lacik's management team marked their agreement by formally adopting a single success metric that counted revenue growth only from end-customer demand (as opposed to demand from intermediary retailers) and only from stores that had been open for at least 12 months, as well as from Pandora Online. In adopting that metric the team was committing to do the hard work of improving Pandora's underlying relationships with consumers.

5

Send a unified message. Executives often comment that members of their teams are texting them during meetings asking for updates. This is not an effective communication strategy for any change effort. You do not want each executive's team getting its own version of events. Nor do you want to count on the leaders'

Facing True Disagreement

What if you've tried everything—set clear parameters, provoked an early exchange, conducted what you thought was a quality debate—and you still haven't been able to persuade some people on your leadership team to truly agree with your vision for change? You have four options.

Disagree again. This is the first, best, and cheapest option. Getting agreement on something as complex as a transformation takes persistence. We often see executives throw their hands up in despair, saying something like, "Sam will never agree," only to discover that Sam was really just looking for one minor concession. Don't give up too easily.

Subtract and defer. If you've truly exhausted your options by engaging in further disagreement, there is no shame in subtracting parts of the transformation on which you can't reach agreement. It is usually better to implement a smaller change on which your executive team *does* agree than it is to press ahead with a larger change on which you don't all agree. And after seeing the success of a scaled-back transformation, skeptical executives often end up accepting what they had previously opposed.

Offer an attractive exit. If you've reached the end of your possibilities through more disagreement and if you can't (or shouldn't) subtract and defer, consider offering executives who hold minority positions an attractive way out. Give them an option to retire early or take an extended transition plan.

Proceed with a plan. At times external forces beyond your control will compel you to start a change program before you are ready. For example, the board and shareholders may have decided that the company is going to be put up for sale by a certain date and that an immediate cost-reduction program is needed to prepare the company for that negotiation. If that kind of situation happens to you, understand that every day you go without true agreement is a day you are incurring risk and missing opportunities. It should be a top priority to settle the disagreement while executing the change.

cascade of the message through their respective organizations—it becomes impossible to guarantee that everyone who needs to hear it will hear it. The decision needs to be broadcast simultaneously in a single format to everyone who needs to know. As Lacik pointed out to us: "When you're communicating with 30,000 people in a supply chain, you need to be simple. You can have all the fancy analysis, but if people don't get it, they won't execute it!"

For example, a unified broadcast might say, "We agreed to execute the program described on page 2 of the briefing document. We need to act now to respond to the threat from our new competitors. We agreed that we'll start making



CHANGE
MANAGEMENT

the arrangements immediately and take it to the board in August. We agreed that Timo will serve as chief transformation officer for the program. We agreed that the budget for the program is \$X, and its EBIT target is \$Y. We haven't yet agreed who the workstream leaders or initiative owners will be—that's to come. We will meet again as a team in July to make the next set of decisions about how we'll implement the program."

At Pandora, Lacik captured the strategic shift he wanted in a simple slogan that was communicated to all the company's stores: "Moments First." That slogan communicated to the employees in each store that their number one priority should be jewelry that customers could use to commemorate important moments in their lives—most of all, bracelets. To convey how serious he was about this Lacik even required some stores to remove products that were not part of the refreshed Moments campaign from shop windows and replace them with bracelets and charms that were. At the same time he set daily sales targets in every store—specifically for bracelets to signal that they were the most important products for the company's future.

TO LEAD A TRANSFORMATION, leaders must take the time to get their team to truly agree on why change is needed, what those changes will be, and how they will occur. There is often more time than leaders think there is to get that right, even in high-pressure situations, and if executives have reached true agreement on a change, there typically will be opportunities to accelerate later. By contrast, programs without an up-front agreement encounter significant delays during execution, requiring far more time and energy than would have been spent on debate at the beginning. ■

HBR Reprint R2604A



JULIA DHAR is a managing director and partner at Boston Consulting Group. She leads the firm's People & Organization Practice in North America and is a fellow at the BCG Institute. **KRISTY R. ELLMER** is a managing director and partner at Boston Consulting Group and the leader of BCG Transform. **PHILIP JAMESON** is an associate director at Boston Consulting Group. They are the authors of *How Change Really Works: Seven Science-Based Principles for Transforming Your Organization* (Harvard Business Review Press, 2026), from which this article is adapted.





ORGANIZATIONAL
DECISION-MAKING

What Companies Get Wrong About Decision Rights

Four common mistakes
and how to avoid them.



AUTHORS

Lindy Greer
Professor,
Ross School of
Business

**Jennifer
Jordan**
Professor, IMD


Maxim Sytch
Professor,
Ross School of
Business



PHOTOGRAPHER DORIS MITSCH

Harvard Business Review
July–August 2026

59



ORGANIZATIONAL DECISION-MAKING



ABOUT THE ART

Doris Mitsch's composites of digital photographs show the flight trails and paths of birds and other animals.

At

ONE GLOBAL TECHNOLOGY
company, 12 executives
recently crowded around a
table to make the final decision

on a heated issue: whether to add the C-suite role of chief innovation officer to their leadership team and expand the central unit that person would run. Their company, which made a variety of tools used in hospital settings, was confronting a changing environment that required it to ramp up its efforts to launch new and improved products. While the new role and the bigger unit would seemingly address that need, the discussion quickly devolved into a power struggle, with several participants arguing loudly and others quietly checking out. After 90 minutes of intense debate, the meeting ended without a decision.

The task force working on this proposal had created a clear process and set of rights for the decision: The CEO would make the final call while in the room with the four people who had critical input on it (the chief operating officer, the chief human resources officer, and the directors of the two biggest business units). Throughout, he would consult and inform others in the organization who could offer insights on the new position or would be affected by it. However, the way the decision was approached—with all 12 executive team members involved in the contentious final meeting—ignored those designated roles.

This is an example of a decision-rights failure.





Clear decision rights are crucial for collaboration: They prevent confusion, speed up execution, and reduce conflict in organizations.



IDEA IN BRIEF

THE PROBLEM


Organizations use decision-rights tools to clarify who should provide input on, make, and execute decisions. Yet these frameworks often fail in practice.

WHY IT HAPPENS

Leaders commonly make four mistakes: defining roles before clarifying goals; treating decision rights as a static list created by a single leader; misunderstanding what the roles actually mean; and allowing hierarchical habits to override assigned roles.

THE REMEDY

Approach decision-rights tools as frameworks for ongoing conversations. Clarify goals first, define expected behaviors for each role, assign the roles collaboratively, and revisit assignments as decisions evolve. When used dynamically, these tools will speed up decisions, reduce power struggles, and improve collaboration and accountability across teams.



When teams try to assign roles before goals have been carefully defined, discussions about decision rights often become turf wars.

Clear decision rights are crucial for collaboration: They prevent confusion, speed up execution, and reduce conflict, especially in matrixed, cross-functional, and project-driven organizations. Tools such as RACI, RAPID, and DARE are meant to help organizations define them. Unfortunately, these tools often produce only a document detailing who will play what role in a decision. And as one manager at a global e-commerce company told one of us (Lindy) in a workshop: “Decision rights are like the position plan for a children’s soccer game—a nice plan on paper that no one understands or remembers.”

One problem is that when leaders use decision tools, they often encounter resistance and skepticism. In our many years of studying how power dynamics shape interactions and decisions, and in our work advising and educating more than 100 global companies in a wide variety of industries, we’ve identified four common errors that organizations make when establishing decision rights. In this article we’ll explain why these mistakes persist. We’ll also describe how leaders can transform the output of decision-rights tools from static, neglected spreadsheets into a process promoting ongoing conversations that improve the quality of decisions, accelerate how quickly they’re made, and secure the buy-in of those affected.

How Decision-Rights Tools Are Meant to Work

At their core these tools offer simple rules for organizing collaboration around a specific decision. They’re used to assign people to the roles needed to make and execute it well. (See “Who Has the D: How Clear Decision Roles Enhance Organizational Performance,” HBR, January 2006.) RACI (with which teams identify who will be *responsible*, *accountable*, *consulted*, and *informed*) is the most common one. We’ll use it as the main example in this article, but our recommendations also apply to similar tools.

In RACI’s framework the *accountable* person makes the final decision and leads the decision team. We often advise organizations to use ARCI, not RACI, as it puts the role of the decision owner first and helps prevent confusion about the differences between the accountable and the responsible roles (which we’ll explain shortly). Only one person should be accountable for any given decision; having multiple people with that right invites power struggles and can cause

delays. As we’ve noted, at the medical technology company that was considering a new head of innovation position (which we’ll call MedTech to preserve client confidentiality), the CEO was assigned to be accountable for the decision.

The people who are *responsible* provide unique and critical input, surface trade-offs, and engage in debate to help the person who is accountable make a good decision. Together with the accountable person, they constitute the decision team. Given that smaller teams are often most effective, only two to four people should be assigned to this role. At MedTech the executives in the responsible role for the decision about the new position—the COO, CHRO, and the directors of the two biggest business units—should have been the only ones in the room with the CEO for the final meeting about it.

Those in the *consulted* and *informed* roles are not part of the decision team but remain essential. People who are consulted provide expertise or perspective that the core team lacks, and people who are informed provide crucial support for successful implementation. Identifying these roles requires more than scanning an org chart; it entails understanding where influence actually resides. At MedTech the COO consulted the CFO on budget constraints and the chief legal officer on the legal red tape that this decision might involve, while proactively informing the HR leaders across the business units who would be affected by the new position and the expanding innovation unit.

So why does this seemingly straightforward process break down? It’s not because the tools are flawed but because they’re misunderstood, misused, or disconnected from real behavior. Let’s take a deeper look now at the four mistakes executives typically make—and how leaders can avoid them.

MISTAKE 1

Confirming Roles Without Clarifying Goals

When teams attempt to assign roles before goals have been carefully defined, discussions about decision rights often degenerate into ego-driven turf wars. Sometimes objectives are far too broad (for instance, “Create a strategic plan for product line X”) and not broken down into concrete steps or subgoals (such as “complete a SWOT analysis” and “survey stakeholders to identify key strategic opportunities”). That



All images: © Doris Mitsch, archival pigment prints; courtesy of CLAMP, New York. First spread: Lockdown Vultures (Moab Slope); second spread: Lockdown Vultures (Russian River); this page: Lockdown Vultures (Moab Mesa)

makes it impossible to allocate ownership of specific decisions and identify where collaboration is needed. Other objectives, in contrast, can be too narrow or insignificant (“Who’s making the slide template to use to present our strategic plan?”). Particularly when goals are too broad, conflicts arise among executives over the rights for decisions about them.

This was a problem at MedTech, where the CEO, CHRO, and business unit leads all thought they owned decisions regarding the company’s head count. Robust discussions we facilitated at the company revealed that several head-count-related decisions had not been disentangled, leading to oversize meetings and disagreements about who got to make final calls. After a formal debate the company’s executives landed on this: The CFO was accountable for approving the budget for the yearly head count. The CHRO was accountable for the proposal for the head-count size and for approving urgent, unplanned requests for jobs above a certain seniority level (except for C-suite hires, which would be approved by the CEO). And the leaders of each business unit were accountable for filling the roles in their preapproved head count and new positions below a certain pay grade.

Before team members focus on role assignments, they all should be able to articulate the specific, measurable, and

time-bound goals and subgoals. In that discussion the best teams iterate between goals or subgoals and the roles. When they hit a snag, they go back, test, and refine their goals and subgoals before finalizing the RACI. This approach often reveals that the stakeholders competing to own decisions may actually want to own different subgoals, which can lead to win-win resolutions of many RACI conflicts.


MISTAKE 2

Assuming Everyone Will Adhere to the Boss's Spreadsheet

A common and costly mistake—one that MedTech made—is treating decision rights as a static list created by a single senior leader, captured in a spreadsheet. The assumption is that once roles are assigned and documented, people will play them. But in practice, they rarely do.

One senior marketing leader at a multinational health-care technology company learned that the hard way. In an attempt to resolve alignment problems between her central and regional teams, she unilaterally defined decision rights and handed a list of them to her direct reports, expecting





ORGANIZATIONAL DECISION-MAKING

them to comply. They didn't. Because there was no shared buy-in, several didn't play their specified roles, and collaboration faltered.

At a large global e-commerce company we advised, managers created spreadsheets with thousands of rows of decisions, each with its own unique RACI. The company's teams looked at the file once and never again.

The problem in these cases was the absence of up-front discussions, when the rights were being determined, about who should have which role in which decision. When people help allocate roles, they're far more likely to commit to them. High-performing teams understand that RACI and tools like it are not ends in themselves; they're conversation starters. They prompt team members to clarify goals, confirm responsibilities, support one another in their positions, and hold one another accountable.

The obvious solution is to cocreate RACIs rather than dictate them. Bring the people who will live with the decision into the room to debate roles and resolve tensions. This is where power sharing becomes visible: Leaders model stepping back when they should be in a responsible or consulted role and stepping in when the team needs them to be accountable.

Be forewarned: These conversations aren't easy. Initial team discussions of roles can feel difficult and unproductive. When challenges arise, such as when meetings to allocate decision rights deteriorate into clashes and infighting, you need to directly address the area of tension. RACI discussions often surface underlying communication problems or resentment that people have been reluctant to raise but that need to be resolved or aired to move forward.

To support such conversations, one client of ours, a global energy company, added an appendix to its corporate operating manual on how to handle common RACI conflicts within its matrix structure. It highlighted typical disagreements that might arise, offered practical talking points for navigating them, and laid out clear criteria for determining who should be accountable for a decision if confusion arose. Two good approaches we have seen companies take with such conflicts: resolving them by thinking through what would be in the best interest of the broader company, and identifying and assigning accountability to whoever is closest to the decision and has the most relevant perspective.



Teams often have differing views of the behaviors expected for each role. Nearly every organization we've worked with has disagreements about what the RACI roles mean.

With practice, cocreated RACIs can create the buy-in needed for people to actually play their positions when the team takes the field.

MISTAKE 3

Misunderstanding Roles

Teams often have differing views of the behaviors expected for each role. Recently, at a session with a global consultancy that had used RACI for years, we polled 30 partners about which role had the final say in a decision. Half the partners said it was the accountable person, while the other half said it was the responsible person. One partner noted that in his local language “responsible” and “accountable” translate into the same word. Ironically, the tool the firm used to create clarity led to more confusion because of disagreement about what the RACI roles meant in practice.

Such confusion often causes decision roles to be set aside. Consider MedTech again. Though it had assigned decision rights for the creation of the head of innovation position, the final meeting revealed that its leaders didn't understand what the RACI meant. For them, it implied simply that the accountable person (the CEO) and the three executives in the responsible role should be in the room when the decision was made. The 12 members of the executive team didn't realize that no one but those four people should be there. They mistakenly assumed that there wouldn't be buy-in unless the full team was present.

This error is common. Nearly every organization we've worked with has latent disagreement about what the roles in RACI actually mean. The fix is straightforward: Build a simple, behavioral description of each role and institutionalize it. When people know what being accountable looks like in action (how that person gathers input, facilitates debate, makes a call, and explains it), the tool stops being theoretical. The same holds true for the other roles.

At MedTech we helped senior leaders draft a comprehensive guide that included clear definitions of the RACI roles and concrete examples of the desired behaviors for each one. (If you'd like to see what such a guide looks like, a set of role descriptions written by participants of workshops run by the University of Michigan's Sanger Leadership Center, who hail from a range of industries and cultures, can be found on the

center's website in the “Collaboration Tools” section under the “Resources & Tools” tab.)

Though individual organizations' definitions may vary somewhat, when RACI roles are played correctly, the accountable and the responsible people always gather input from the folks designated to be consulted and work to garner buy-in from those who should be informed. Then, when it's time to hold a meeting to decide, the accountable person convenes with just the two to three people in the responsible roles. During that meeting it's the accountable person's job to ensure that the team shifts from command-and-control mode to “flat” mode (in which the leader levels the hierarchy and shares power) to brainstorm. (See “You Need Two Leadership Gears,” HBR, March–April 2023.) After kicking things off, the accountable person encourages the others to share their input (and the information and perspectives they all have gleaned from people in the consulted and informed roles) and debate options. At the meeting's end the accountable person makes a final decision integrating the insights that emerged during the debate. Afterward, the accountable person explains the decision to the people in the consulted and informed roles.

MISTAKE 4

Getting Stuck in the Same Roles

A final mistake occurs when people get trapped in certain roles despite the best of intentions. In some cases, senior leaders are always accountable, and the people in the layer below are always responsible. In others, people act as if they're accountable when they're not—something that often happens when a teammate at a lower level is accountable on paper but that person's boss still acts like the one in charge.

This dynamic led one venture capital firm to pass up the opportunity to invest in a startup that later became a unicorn. One of us (Lindy) was brought in to do a postmortem on this miss. It turned out that the firm's senior partner had dominated the discussions during the deal evaluation process and had de facto made the final decision (in other words, acted as if he were accountable). But his decision didn't take into account the data the firm had gathered and instead was based on his overall experience in the sector. The accountable role should have belonged to the associate who had performed the



ORGANIZATIONAL DECISION-MAKING

due diligence on the deal, had the deepest expertise on the target, and was in fact advocating for investing in it.

The best teams are intentional about tailoring roles to the topic at hand. They don't get mired in ingrained patterns of power or deference to the formal org chart. It's clear that major companywide decisions—such as the final approval of the firm's yearly strategic priorities—may remain with senior leadership, with the CEO holding the accountable role. But the power to make a local decision, such as whether to make a moderate equipment investment at a specific plant, may rest further down the organization with the best-informed person, such as the plant's manager. That might sound like a no-brainer, but you would be amazed at how frequently senior managers don't sufficiently delegate decisions. RACI can help teams break free from rigid hierarchies—empowering employees, strengthening succession pipelines, and preventing executives from burning out.

We often challenge the senior leaders we work with to identify just four decisions a year where they are truly the only—and best—person to make the call. These tend to involve enterprisewide issues such as strategy, hiring or promoting people into senior positions, and company-shaping investments. For everything else, we ask a different question: How can you step out? That may mean shifting from an accountable role to a responsible, consulted, or informed one or questioning whether they need to be in the RACI at all.

Doing so is easier said than done. Forgetting about rank can be difficult; it takes awareness and practice. To help you and those around you shift between roles, you can embed RACI into everyday company tools like performance-management guides, meeting agenda templates, and Gantt charts for project management. These cues remind people to be explicit about the role they're expected to play at any given moment. Over time the best teams develop the muscle to move fluidly among roles, increasing their agility and that of their organizations.

As a result of our coaching, MedTech's leaders focused on revising meetings to help people play the RACI roles assigned to them for each decision. To do so, they first had a meeting about meetings. They agreed to stop including too many people in decision meetings and instead make sure each meeting was aligned with its RACI—for instance, that it would be attended by only the accountable person and those

in the responsible roles for that decision. For every meeting they wrote out the specific goal for each agenda item (for example, "Decide on the head-count request for a hire to lead business unit X") and the person accountable (in this case, the CHRO).


For each decision on the agenda, they also created meeting preparation materials summarizing the information they had gathered from people in the consulted and informed roles. And they set agenda schedules: for example, five minutes for the accountable person to clarify the decision at hand, 20 minutes for the accountable and the responsible people to engage as equals in a debate about the options, and five minutes at the end for the accountable person to make a final call based on the information gathered. After the meeting they coached and provided feedback to one another, individually and as a team, on how well they were playing their RACI roles.

All those efforts got the executives to stay in the lanes assigned to them for decisions. As a result, MedTech improved both its agility and the quality of its decisions.

TO ORGANIZE THEIR teams to make sound decisions, leaders cannot just articulate roles in spreadsheets. They must establish a robust process that's woven into work—and ensures not only that people stick to their designated roles but also that the roles themselves are continually revisited as goals evolve or friction arises. Leaders and their teams should routinely ask: Did we play our positions? Where did our assigned roles support or constrain us? Such an approach will turn decision rights from static artifacts into living systems. It will give teams the confidence to move fast, adapt intelligently, secure buy-in, and share leadership in ways that make the whole organization more effective and resilient. ■

HBR Reprint R2604B

 **LINDY GREER** is a professor of management and organizations at the Stephen M. Ross School of Business at the University of Michigan, where she is the Gerald and Esther Carey Professor of Business Administration and the faculty director of the Sanger Leadership Center. **JENNIFER JORDAN** is a psychologist and a professor of leadership at IMD in Switzerland, where she directs the Advanced Leadership and Women on Boards programs. **MAXIM SYTCH** is a professor of management and organizations at the University of Michigan's Stephen M. Ross School of Business and the author of *The Influence Economy* (Oxford University Press, 2025).





STRATEGY

THE POWER OF STRATEGIC CENTERING

Choosing a clear organizing principle
can help companies navigate
uncertainty and act faster.



AUTHOR

Rita McGrath
*Professor, Columbia
Business School*





STRATEGY



ABOUT THE ART

Artist Field Kallop draws inspiration from mathematics and mysticism, balancing rigor and fluidity in her art.


Imagine you are sitting on the capital allocation committee of a major company that is in the eye care, consumer healthcare products, generic drug, and pharmaceutical businesses.

Two proposals are on the table. The first is a multibillion-dollar request to build a next-generation production facility for consumer health products, such as over-the-counter painkillers. The investment case is straightforward: stable demand, predictable margins, and a well-understood manufacturing process.

The second proposal requests several billion dollars to acquire three companies that are pioneering a novel therapy area called radioligand therapy—a radically different approach to treating cancer that uses radioactive particles attached to molecules to target tumors with almost pinpoint precision. The science is promising—but early. The manufacturing is novel. The regulatory pathway is unclear. It's an edge case investment.


How do you choose? The factory proposal for over-the-counter painkillers feels safe. The radioligand bet feels risky. You can model the economics of the first; you can only imagine those of the second. In the logic of traditional capital allocation, the safe bet often wins.

But if your company has decided to organize its identity, strategy, and resource allocation around a single coherent organizing principle, then the safe bet may not be the best one. A different bet might make more sense—as the leaders of the company I've just described, Novartis, discovered. When its CEO, Vasant “Vas” Narasimhan, centered the company on innovative medicines that truly help patients in





When value dematerializes, the old strategic anchors dissolve with it. Executives need a new anchor—not a position to defend but a center to home in on.



IDEA IN BRIEF

THE PROBLEM

As value shifts from physical assets to intangible ones, traditional strategy frameworks no longer provide clear guidance.

THE REASON

In a dematerializing economy, industry boundaries blur, technologies evolve rapidly, and competitive advantages are hard to sustain. Without a clear organizing logic, companies default to “safe” decisions, slowing adaptation and diffusing resources across conflicting priorities.

THE SOLUTION

Adopt strategic centering: deliberately choose one organizing principle—a mission, a customer, a technology, an ecosystem, or friction erasure—to guide decisions. A clear center bounds the opportunity set without closing it, sharpens resource allocation, and enables permissionless action across the organization.



Centering is not the same as purpose, a term that has become so expansive and aspirational as to be nearly useless for strategic decision-making.

specific disease areas, he didn't see the radioligand investment as a gamble. It was directly connected to Novartis's strategy.

Shortly after Narasimhan learned about the science of radioligands, Novartis acquired two pioneering companies in that space. Today Novartis is a world leader in a small but highly promising therapeutic area that could transform cancer treatment. Its consumer healthcare products business? Divested.

Here's where things get interesting. When Narasimhan took over as CEO, in 2018, Novartis's valuation (its market capitalization) with all the businesses together was between \$190 billion and \$200 billion, and it employed 126,000 people. Today the smaller pure-play pharmaceutical company is valued at around \$300 billion, with just 76,000 employees. Alcon, the eye care company that Novartis spun off in 2019, is worth around \$40 billion. Sandoz, the generics and biosimilars business Novartis spun off in 2023, is worth around \$35 billion. Haleon, the company spun off by GSK, which includes the divested Novartis consumer healthcare business, is worth around \$40 billion. That makes the value of the now-separate companies about two times Novartis's original valuation, without any dramatic changes in their underlying businesses (and not accounting for dividends and share buybacks).

Narasimhan's move illustrates a strategic concept that I believe is urgently needed in a business environment characterized by radical uncertainty, dissolving industry boundaries, and an accelerating shift from physical to intangible value creation. I call it *strategic centering*: deliberately choosing an organizing principle that provides clarity and coherence to a company's strategy when the traditional anchors—defensible positions, stable industries, long-lived physical assets—are losing their power.

In the following pages I'll present the concept and identify the various types of centering that companies can consider. I'll begin, though, by explaining how it has come of age.

THE END OF THE "STUFF" ECONOMY

The strategy frameworks that most executives rely on were forged in what we might call the "stuff" economy. Michael Porter's seminal five forces model worked beautifully when

we had stable industry boundaries and well-defined roles for the players (buyers, suppliers, and so on). The resource-based view assumed that companies could acquire or create valuable, rare, durable, hard-to-imitate assets and capabilities. Even W. Chan Kim and Renée Mauborgne's influential *Blue Ocean* strategies assumed executives were creating definable market spaces where it was clear whether or not there was established competition. Those frameworks made sense in a world where most corporate value resided in tangible things: factories, distribution networks, inventory, and equipment. In 1975 roughly 83% of the assets on the books of *Fortune* 500 companies were tangible. Strategy was, in essence, about identifying opportunities or defending positions in physical space.

Today that ratio has inverted: Approximately 90% of corporate value is now in intangible assets such as software, data, brands, relationships, intellectual property, and organizational capabilities. The shift has been accelerated by what the economic historian Carlota Perez calls a technological revolution—the kind of once-in-a-half-century transformation that reshapes not just technology but the entire logic of how economies work. She identifies five such revolutions in the history of capitalism, each driven by a cheap input with economy-wide applications: the original industrial revolution; steam and railways; steel and electricity; oil, automobiles, and mass production; and now information and telecommunications.

Each revolution follows a pattern. In the irruption and installation phases, financial capital rushes in, bubbles form, and inequality surges. Then comes a turning point—a period of crisis and institutional adjustment. Finally, in the deployment phase, the new paradigm takes hold, institutions adapt, and a new age of broad-based prosperity becomes possible.

We are at a turning point right now. The symptoms are unmistakable: massive income inequality, political instability, widespread dissatisfaction, and an old strategic paradigm whose productivity has been exhausted but has not yet been replaced. The defining characteristic of this transition is what I call *dematerialization*: Value creation is shifting decisively from physical stuff to digital services, coordinated ecosystems, intangible assets, and experiences.

Consider what has happened to music. When you wanted to hear a song you liked, you used to buy sheet music and play it yourself. Later, you discovered new songs by listening



STRATEGY

to the radio. If you really liked one, you could buy a record to hear the artist play it, and later still, a CD. Today most people don't buy music. It flies through the air and lands on a device, and an algorithm curates it for the listener. This shift—from physical artifact to dematerialized service—is playing out across industry after industry, including media, financial services, healthcare, and even manufacturing.

When value dematerializes, the old strategic anchors dissolve with it. What is your industry when digital technology erases the boundaries? What is your competitive advantage when intangible assets can be replicated or leapfrogged at digital speed? What is your planning horizon when technologies emerge, combine, and become obsolete in years rather than decades? Executives need a new anchor—not a position to defend but a center to home in on.

WHAT STRATEGIC CENTERING IS—AND ISN'T

Strategic centering is deliberately choosing one organizing principle to guide resource allocation, opportunity selection, and organizational identity in a dematerializing world. It involves answering a simple but profound question: What are we really about?

Centering is not the same as purpose, a term that has become so expansive and aspirational as to be nearly useless for strategic decision-making. (For instance, a theater company whose purpose is “to create art that causes you to feel something” is no closer to knowing whether to produce experimental plays or beloved musicals than a theater group that didn't articulate a purpose.) Centering is not vision, which is too often a static destination rather than a dynamic direction. And it is not a core competence, which can become a prison when the underlying technology shifts.

Centering is closer to what Theodore Levitt was getting at when he asked railroad executives whether they were in the railroad business or the transportation business. But it goes further because in a dematerializing world, the question isn't just about broadening your definition of your industry. Rather, it's about choosing the dimension along which you will pursue coherent opportunity sets as industries themselves dissolve.

A center does three things. First, it bounds the opportunity set without closing it. You know you cannot pursue

everything, but a center lets you pursue a set of possibilities that share a common logic. When Novartis centered on innovative medicines, it didn't just narrow its focus—it opened up an entire landscape of therapeutic technologies (including radioligands, gene therapy, and siRNA medications) that its old conglomerate structure might have obscured.

Second, centering resolves capital allocation dilemmas. In a diversified company without a center, every investment decision is a battle between competing logics. Once a center is clear, resources flow toward it naturally. To go back to the initial example, the painkiller-versus-radioligand debate simply vanishes.

Third, having a center enables what I call *permissionless* action. When people throughout an organization share a clear understanding of the center, they can make decisions and take action without waiting for approval from above. They know the direction; they don't need permission to take each step. That dramatically increases organizational speed and reduces the return on internal politics. When what matters is advancing the center rather than climbing the hierarchy, the dynamics of the organization shift fundamentally.

FIVE WAYS TO FIND YOUR CENTER

My research has identified five types of strategic centering, each anchored to a different organizing principle. (See the exhibit “The Five Centers.”) They are not mutually exclusive in a strict sense; elements of several may be present in any company. But the organizations with the clearest strategy have a dominant center—and the discipline to subordinate other considerations to it.

1

Center on the mission: “What problem are we solving?”

Mission centering organizes a company around a problem to be solved, independent of any particular technology, customer segment, or product category. The company follows the problem wherever it leads.

Novartis under Narasimhan is a clear example. By centering on innovative medicines in specific therapy areas, Novartis gave itself permission to pursue whatever therapeutic technologies showed the most promise—even when those technologies, like radioligand therapy, bore no resemblance to the company's traditional chemical-compound



STRATEGY

approach. Narasimhan's scientific background gave him an edge in sensing these opportunities, but it was the organization's strategic centering that gave it the clarity to act on them quickly.

Shopify offers another example. The company's mission—to make commerce better for everyone—has guided Shopify into businesses from simple online storefronts to payment processing, shipping logistics, point-of-sale systems, and business financing. That portfolio would look incoherent through an industry lens, but it is perfectly coherent through a mission lens.

Mission centering works best when the problem being addressed is large enough to sustain decades of exploration, when the technologies for addressing it are in flux, and when the organization has deep domain expertise that transfers across technological generations. The danger is mission drift—gradually broadening the mission until it means everything and therefore nothing.

2

Center on the customer: “Who are we serving?”

Customer centering organizes a company around a deep understanding of its customers and their evolving needs. The company follows wherever those needs lead, crossing industry boundaries as necessary. Tony Ulwick and Clayton Christensen's notion of “jobs to be done” is useful here. Companies ask themselves, “What jobs are we doing for these specific sets of customers?”


Amazon's early years offer a canonical example. “Start with the customers and work backward to see how their needs could be met by online technology” was the operating principle that drove investments—from one-click ordering to Prime to Amazon Web Services, which originated as an internal infrastructure solution to serve customers better and faster. Customer centering gave Amazon the coherence to move from books to electronics to cloud computing to streaming to grocery to healthcare—a trajectory that defies traditional industry classification.

Brian Chesky's return to what he calls “founder mode” at Airbnb is a story of recentering. After a period in which the company drifted toward becoming a real estate platform optimized for hosts and institutional operators, Chesky pulled the center back to the guest experience—the traveler who wants to “belong anywhere.” That recentering drove a



redesign of the product, the search experience, and the company's relationship with its host community.

Customer centering works when your customers have complex, evolving needs that cut across traditional industry boundaries, and when deep customer understanding creates compounding advantages. The danger with this approach is that being truly customer-centric may over time conflict with other corporate goals. For instance, Amazon has been criticized for letting advertising interfere with the customer experience. That may be a sign that the company is starting to drift from its center. Southwest Airlines, bowing to pressure from an activist investor, is being criticized for becoming just like any other airline, instituting (once-unthinkable) assigned seats, tiers of service, and baggage fees.



3

Center on a technology: “What can our capabilities do?”

Technology centering organizes a company around a set of deep technological capabilities, and then the company searches for applications across any domain where those capabilities create value. It is agnostic about particular customers and markets; the organization follows its technology wherever it can be applied.

The contrast between Fujifilm and Kodak is among the most instructive in modern business history. Both faced the same inflection point: Digital photography was destroying their core film business. Kodak, whose CEO misguidedly focused on printing, has become a relatively small, 4,000-person B2B player. Fujifilm today has more than 72,000 employees and is thriving. The difference was centering.

Fujifilm looked at its film business and began focusing on the underlying technologies—chemical compounds, precision coating, optical films, and imaging science. It became completely agnostic about customers. That approach led the company into cosmetics (leveraging its expertise in collagen films and antioxidants), medical imaging, pharmaceutical manufacturing, and advanced materials for LCD screens. The applications look random; centering reveals their coherence.

Two corporations from the energy industry highlight how profoundly different technology-centering approaches can be. The oil giant Shell has positioned itself as a “molecule company,” which means its future lies in plastics, chemicals, and hydrocarbon derivatives even as it transitions away from fuel. BP, by contrast, has positioned itself as an “integrated energy company,” which would logically take it into solar,



STRATEGY

wind, nuclear, and battery technology. These two companies' centerings lead to different opportunity sets, different capability investments, and different organizational identities. Ultimately, they place Shell and BP in different competitive arenas.

Technology centering works when your capabilities are deep and transferable across domains and there is a compelling road map for continuing improvement. By centering on parallel computing architecture, for example, Nvidia moved from gaming graphics to becoming the backbone of the AI revolution. In that case technology centering led to options that were literally unimaginable when the center was first established. The danger with this approach is technological narcissism: falling so in love with your capabilities that you miss the moment when the technology becomes obsolete or that you fail to connect it to genuine problems.

4

Center on a regional or national ecosystem: "What system are we building?"

Ecosystem centering organizes a company around its role in building a larger system, typically the productive capability of a nation or region. The company's strategy is inseparable from the system it is helping to create.

Morris Chang founded TSMC in 1987 with what appeared to be a business model innovation: Manufacture chips designed by others and design none of your own. But TSMC's center was building Taiwan's semiconductor capability into a strategic asset so essential to the global economy that it became the island's ultimate security guarantee—a "silicon shield." That national centering enabled TSMC to make enormous, long-horizon investments that no purely private-market logic could justify, such as fabrication plants that cost more than \$20 billion each. Supported by government partnership, the company enjoys the benefits of aligned national education pipelines and infrastructure investments. TSMC thrives as the catalytic force behind an ecosystem that has made Taiwan a global leader in manufacturing these critical components.

Samsung's history offers parallels. From its start, Samsung was intertwined with South Korea's national development ambitions. Its willingness to make bet-the-company investments in memory chips and displays reflected a centering that transcended quarterly earnings. The result was an industrial

ecosystem that lifted an entire nation's productive capability and created competitive advantages—patient capital, aligned talent pipelines, and government support—that no standalone market entrant could replicate.

One could argue that the post-World War II consensus between business and government in the United States is a further example, although it is now unraveling. The goal of creating a prosperous middle class, supported by government investments in highways, housing, and technology during the Cold War, was widely shared by business leaders. General Motors CEO Charles E. Wilson famously said in 1953 that "what was good for our country was good for General Motors, and vice versa. The difference did not exist."

Ecosystem centering works when there is a common commitment to development goals on the part of business leaders and the state. This allows ventures that require investments at a scale and time horizon beyond what private capital alone can sustain to move forward. It also works when geopolitical dynamics create strategic value in a national capability. Its danger is dependency: Companies centered on national ecosystems can become instruments of state policy, which could generate a backlash from other countries. Huawei learned that when the United States and other countries, fearing security risks, pursued a "clean network" initiative, hoping to wean themselves from the telecom firm's technology entirely. Its revenues dropped 29% from 2020 to 2021. State support can also make companies uncompetitive globally and at risk should protectionist policies disappear, as we have seen with U.S.-based steel manufacturing.

5

Center on friction erasure: "What's unnecessarily hard?"

This type of centering organizes a company around the systematic identification and elimination of unnecessary difficulty in a domain. The company doesn't anchor to a mission, a customer, a technology, or a system. Instead, it asks, "What is still absurdly hard, and how do we make it effortless?" For instance, eliminating the friction for participants in two-sided markets is at the heart of platform strategies used by Facebook (eliminate the friction of finding friends), Google (eliminate the friction of search), and Apple (eliminate the friction of finding apps in its store).

The Five Centers

A centered company deliberately chooses an organizing principle to guide decisions on everything from resource allocation to opportunity selection to organizational identity. Centering, at its most basic level, answers the simple but profound question, What are we really about? There are five distinct types of strategic centering, each with its own benefits and risks.


CENTERING	CORE QUESTION	BEST WHEN...	RISK OF...
MISSION	What problem keeps us up at night?	The problem is large and enduring; the technologies for addressing it are in flux.	Mission drift.
CUSTOMER	Whose needs do we understand better than anyone?	Customers have complex, evolving needs that cut across industries.	Conflict with other corporate goals.
TECHNOLOGY	What capabilities do we have that are transferable across domains?	Capabilities are deep, on a long and exciting technological trajectory, and applicable in multiple markets.	Technological narcissism (failing to connect technological discoveries to real customer jobs to be done).
NATIONAL ECOSYSTEM	What system are we essential to?	The venture requires scale and patience beyond private capital; geopolitics create strategic value.	Geopolitical strategic constraints; uncompetitiveness globally.
FRICTION ERASURE	What's still absurdly hard in our space?	The domain has legacy complexity, regulatory fragmentation, and poor customer experience.	Scope creep (friction is everywhere).

The South Korean fintech company Toss was founded by S.G. Lee, a dentist by training, who looked at the Korean financial system and asked a simple question: Why is moving money so hard? Rather than centering on a customer segment or a technology, Toss centered on the relentless elimination of friction in financial services. Doing so allowed the company to expand from payments to banking to insurance to investment to stock trading to media content, all without losing coherence. Each expansion worked to erase another friction.

The organizational implications are telling. Instead of a traditional manager-to-subordinate structure organized around functions, Toss operates with what it calls “directly responsible individuals.” These people are aligned with key outcomes, not with a reporting structure. Toss has virtually no hierarchy and has a governing principle of being “tightly aligned on mission but loosely coupled” in structure. When one team decided to launch a YouTube channel exploring how money intersects with culture—sports, fashion,

lifestyle—no one had to ask permission. The initiative introduced Toss to an entirely new audience, one with no interest in the core content of its existing YouTube channel but great interest in how athletes earn money. A centralized content department at a traditional bank would have killed the idea. But for a friction-erasure-centered organization, it was an obvious move: Use an unexpected channel to remove the friction between young consumers and financial literacy.

DBS, the Development Bank of Singapore, followed a similar path. By centering on eliminating friction in banking across Southeast Asia, it transformed itself from a sleepy government institution into what has been called the world’s best digital bank. Its mobile-only banking offer allowed it to operate with far lower customer acquisition and servicing costs than banks that required physical infrastructure. Its AI-powered, hyperpersonalized “nudges” catch customers at the exact moment they are making a financial decision and help them to optimize their outcomes. Its Smart Buddy wearable device for schoolchildren allows them to



make contactless payments in school canteens, facilitating budgeting and even nutrition tracking for parents. Not bad for a bank whose nickname was once “Damn Bloody Slow.” The logic here is endlessly extendable: There is always more friction to eliminate.

This type of centering is particularly powerful in a dematerializing world because digital technology is the ultimate friction-erasure tool. Centering in this way naturally leads toward dematerialization—replacing physical processes with digital ones, replacing bureaucratic approvals with automated flows, replacing institutional complexity with elegant simplicity. The danger with this approach is scope creep: When everything looks like a friction to be erased, it can be easy to lose the discipline that keeps the center coherent.

Each kind of centering has a different relationship to the forces of dematerialization. Mission centering enables you to ride technological transitions because you are loyal to the problem, not the technology. Customer centering keeps you relevant as customer behaviors dematerialize (from owning to renting, for instance, and from buying to subscribing). Technology centering lets you find new applications as old markets dematerialize. Ecosystem centering positions you to shape the infrastructure of the next paradigm. And friction-erasure centering is dematerialization in its most direct form—replacing what is heavy, slow, and complex with what is light, fast, and simple.

As you think about where to center, you should consider how much dematerialization your industry has experienced. In industries in which dematerialization is transforming how work is organized but the output is still physical, such as construction and mining, centering on the ecosystem or on technology may be the best choice. In partially dematerialized industries, however, such as healthcare and education, where human expertise remains central, mission or customer centering may best guide the transition. In industries in which value has already shifted to the digital realm but the legacy complexity of the physical era remains—financial services, for example, and media—friction-erasure centering is frequently the best choice. Finally, in industries that are already dematerialized, such as software and AI, the strategic question is often about mission or customer centering, because technology alone doesn’t differentiate you in a world where everyone has access to the same tools.

LEADING THE CENTERED ORGANIZATION

Choosing a center is a necessary act of strategic clarity. But a center is useful only if the organization can move along its trajectory with speed and coherence. This is where leadership is critical—and where the findings from my research surprised me.

Field Kelley: from left, *Stellar Bloom*, 2024, courtesy of the artist and Voltz Clarke Gallery; *The Deepest Space*, 2025; and *Earth Pulse*, 2026, courtesy of the artist and GALLX



I expected the leaders of the most successful centered organizations to be the servant-leader types that decades of management literature have celebrated: Leaders who hire great people, get out of their way, and provide support from behind. What I found was quite different. The leaders of centered organizations are deeply engaged, highly visible, and intensely present. Narasimhan records multiple videos a day, travels widely to Novartis facilities, and is constantly visible. Jensen Huang of Nvidia reads employee emails every Sunday and has equipped every meeting space with whiteboards, which he uses to ideate with people from any level of the company. Airbnb's Chesky returned to reviewing product designs and questioning decisions that he had once delegated away.

Those leaders are not micromanaging. They are doing something more subtle and more demanding: They are maintaining the center. They are out in the organization sensing weak signals (as Narasimhan sensed radioligands). They are providing direction without dictating methods. They are making the center vivid, concrete, and present in daily decisions.

That kind of leadership has a powerful secondary effect. It dramatically reduces the return on organizational politics. As noted earlier, when everyone in an organization is aligned on the center, rather than on climbing the hierarchy, there's little to be gained by playing political games. In every centered



STRATEGY

organization I studied, people described the same phenomenon: Clarity about the center liberated them to focus on the work rather than on navigating the bureaucracy.

This is the connection between centering and permissionless organizations. Centering is the prerequisite for permissionlessness. Without a clear center, removing hierarchy produces confusion. With a clear center, it produces speed.

THE STRATEGIC CHOICE OF OUR ERA

The mass production paradigm that shaped our organizational structures and our assumptions about how value is created and captured are giving way to something fundamentally different—a dematerialized economy where intangible assets, digital services, coordinated ecosystems, and experiences are the primary sources of value. The strategy frameworks we inherited are dangerously incomplete for a new and rapidly dematerializing world.

Strategic centering provides the clarity that enables speed, the coherence that enables resource allocation, and the shared understanding that enables people to act without waiting for permission. In a dematerializing world, the centered organization is the one that can grasp the opportunity in radioligands rather than retreating to the safety of painkillers. It is the one that can launch a YouTube channel about money and culture without convening a committee. It is the one that follows its technology into cosmetics, or its mission into gene therapy, or its national ambition into advanced semiconductor fabrication.

Choosing a center, therefore, is the most important strategic decision a leader can make today. The good news is that multiple kinds of centers can be legitimate. It's imperative that leaders choose one, commit to it with discipline, and build the organization that can act on it. The next golden age will not arrive automatically. It will be built by leaders and organizations that find their center and have the courage to shed everything that doesn't serve it. ■

HBR Reprint R2604C



RITA MCGRATH is a professor at Columbia Business School and the author of *The End of Competitive Advantage* (Harvard Business Review Press, 2013) and *Seeing Around Corners* (Houghton Mifflin Harcourt, 2019).

Making the Leap from Corporate Leader to PE-Backed CEO



LEADERSHIP

Not everybody can manage the transition. Here's what it really takes to succeed.



AUTHORS

Samantha Hellauer
*Director,
ghSmart*

Dina Wang
*Partner,
ghSmart*

Heidi Smith
*Practice leader,
ghSmart*

Samantha Smith
*Research fellow,
ghSmart*







FOR YEARS A COMMON assumption has guided hiring decisions across private equity: When in doubt about whom to hire to run a portfolio company, choose somebody who's been the CEO of a PE-backed venture before. But times are changing. Today, in the face of talent

shortages, bolder value-creation plans, and increasingly complex transformations, PE firms are looking beyond that familiar profile.

At ghSmart we've observed this shift firsthand. In recent years our PE clients have increasingly recruited leaders from corporate talent pools, and more corporate C-suite executives now view PE-backed firms as a compelling pathway to their first CEO role. That trend has been fueled by the rapid growth in the number of PE-owned businesses and the resulting supply-demand imbalance for CEOs with experience running them. According to Citizens Bank, the number of U.S. PE-backed companies has increased by more than 400% over the past 25 years, while the number of publicly listed companies has declined by roughly 35%. Our own data reflects that dynamic: Fifty-three percent of the high-performing first-time portfolio-company CEOs we assessed in 2024 and 2025 came directly from the corporate C-suite or business-unit leadership roles, highlighting how central the corporate pipeline has become to meeting PE firms' growing demand for talent.

Clearly, not every accomplished corporate leader will be well suited for a CEO role at a PE-backed firm. So what are the keys to success?

To answer that question, we analyzed five years of proprietary assessment data for 491 senior executives (241 portfolio-company CEOs and 250 corporate C-suite leaders), evaluating differences across 18 leadership traits. We also reviewed 83 assessments of corporate leaders who later

became CEOs at PE-backed firms and conducted a dozen interviews with CEOs and investors. In doing so, we identified five capabilities that are consistent predictors of success, each of which we'll discuss in this article.

The Five Crucial Capabilities


Our five capabilities reflect the specific demands of operating in a private-equity context, where resources are leaner than they are in the corporate world and accountability is direct and unrelenting. They offer a more reliable and practical way of assessing leaders' potential to succeed in a PE-backed business than simply considering their prior portfolio-company experience. People are unlikely to excel at all five capabilities, but the better they do across the board, the greater the odds are that they'll succeed at a PE-backed firm.

Practical commercial orientation. Corporate executives often are great at long-range planning, but CEOs of PE-owned companies are under immediate pressure to translate strategy into value creation. Those who succeed understand what will move the needle and adjust a firm's direction quickly as data comes in. In our research, CEOs of PE-backed firms were 17% more likely than corporate C-suite leaders to excel at the commercial side of the business, focusing on and pulling the levers that increase revenue. As Robert Hanson, a two-time portfolio-company CEO, told us, "The biggest adjustment in PE is trusting that it really is all about the value-creation plan, not about appeasing stakeholders." Because the clock starts on day one, CEOs are expected to move beyond financial oversight and into active commercial leadership. "If you're hired as a category expert from corporate," Hanson explained, "you're not there just to manage the annual budget. You're expected to create value in that category by driving consistent strategic growth and margin expansion."

Ability to tackle strategy under pressure. Corporate leaders often craft strategy in long planning cycles that involve heavy governance and extensive stakeholder alignment, but portfolio-company CEOs are under constant time pressure. "You can't succeed in PE by setting strategy from a distance," one of them told us. "You have to be hands-on, getting into the fray and translating direction into action." In our analysis, CEOs of PE-backed firms were 20% more likely than corporate C-suite leaders to make strategic thinking



Corporate leaders often craft strategy in long planning cycles that involve heavy governance and extensive stakeholder alignment, but portfolio-company **CEOs are under constant time pressure.**



IDEA IN BRIEF

THE PROBLEM

Private-equity firms increasingly hire corporate veterans to run their portfolio companies, but many of those leaders struggle because PE investors demand immediate value creation, lean operations, and relentless accountability—which makes the role quite different from most corporate ones.

THE SOLUTION

Success depends on five traits: commercial focus; the ability to tackle strategy under pressure; the ability to gain influence; risk tolerance; and interpersonal range. Firms and candidates must rigorously assess readiness for this faster, higher-stakes context.

THE PAYOFF

A clear understanding of the required capabilities and motivation will expand the CEO talent pool and increase the odds of faster execution, stronger investor trust, and superior returns.



Questions to Ask

By focusing on the predictors of CEO success that matter in private-equity-backed businesses, PE firms can expand the leadership candidate pool, and corporate executives can more honestly assess whether they're prepared to make the leap into running a PE-owned venture. To help with that process, we've devised a set of questions that PE firms and corporate candidates should ask.

WHAT FIRMS SHOULD ASK ABOUT CANDIDATES

→ **Do they have a bias for action anchored in what will practically create value?**

When have they made high-stakes decisions without perfect data, and how did those decisions translate into commercial outcomes?

→ **Can they quickly simplify strategy into a small set of executable priorities?** How well do they pivot when new data emerges?

→ **Have they taken meaningful risks, especially on talent?** How quickly have they assessed, upgraded, or replaced leaders when the business demanded it? What other consequential bets have

they made where they were fully accountable for results?

→ **How well do they deliver results through others?** Can they empower a team to execute at a rapid pace and on a large scale? Do they take pride in collective outcomes, or do they overindex on "I" rather than "we"?

→ **Do they build trust across diverse stakeholders?** Are they comfortable communicating with candor, adjusting their style to the context, and engaging in productive conflict?

→ **What exposure have they had to scrappier environments?** Have they led—and gotten results—in settings without deep infrastructure, large teams, or established playbooks? Have they stepped

into messy or broken situations where they were required to stabilize or rebuild a business?

→ **What motivates them to take on this role?** Beyond the compensation upside, are they energized by building, problem-solving, and operating with sustained intensity? Is there evidence that they will "have fun through the grind," as one CEO put it?

WHAT CANDIDATES SHOULD ASK THEMSELVES

→ **Am I energized by making fast, bottom-line-driven decisions amid ambiguity?**

How well do I understand—and buy into—the commercial drivers behind the value-creation plan?

→ **Do I quickly cut through what's interesting to focus on what matters?** Or do I prefer having time for extended exploration and refinement?

→ **Does the idea of having greater accountability and ownership energize or drain me?** Am I ready to be the

primary source of clarity, conviction, and momentum, or do I prefer to share responsibility with others?

→ **Am I prepared to take significant risks?** Will I feel confident having a talent plan within the first 120 days and making changes within a year?

→ **Am I comfortable leading with real-time transparency?** Will a high-candor environment, where the good, the bad, and the ugly are brought to the surface early and often, play to my strengths or clash with my style?

→ **Do I enjoy the building process?** Am I energized by finding out what I don't know, fixing what's broken, and creating structure where little exists, or do I prefer to have established teams, processes, and support systems in place?

→ **What do I really want in my next role?** Am I more motivated by the work, the compensation, or the idea of being the CEO? How will I feel on the hardest days of the hold period?

a priority. For leaders who move into PE-backed roles, the absence of quarterly earnings constraints and of the need to build broad consensus for decisions can feel liberating. "In PE," we were told by Jesper Nordengaard, a first-time portfolio-company CEO who'd previously served as Colgate-Palmolive's North American president, "you make a decision, and the next meeting is about how you're implementing it."

Ability to wield influence widely to create impact. In corporate environments leaders often rely on established teams, infrastructure, and influence pathways to produce results. Relationships and credibility cultivated over years in a corporate hierarchy matter far less in a PE portfolio

company, where new ownership and board members, a reset strategy, and limited inherited authority mean leaders must know how to quickly earn stakeholders' confidence. "You're effectively starting over," one portfolio-company CEO told us. CEOs of PE-backed ventures often need to exert direct and indirect influence across a large range of stakeholders—from direct reports to people in the broader organization to board members—in order to achieve the alignment that enables execution. They must systematically mobilize the organization toward faster decisions and tangible results with an intense, deliberate presence. That means spending significant time in the field, working one-on-one



Leaders who thrive in PE-backed firms don't wait for consensus or perfect information; they place selective bets, make trade-offs quickly, and own the consequences transparently.

with leaders and engaging with their teams and frontline employees. New CEOs need to stay close enough to the work to test assumptions and remove blockers while deliberately pushing accountability for results down and across the organization. As Eric Jungbluth, a CEO who has led multiple PE-backed firms and previously had leadership roles at three public companies, told us, one thing matters above all else: “How well you drive execution through others.”

Willingness to take risk. In our analysis, portfolio-company CEOs were 12% more likely to score high on risk-taking than corporate C-suite leaders were. Leaders who thrive in PE-backed firms don't wait for consensus or perfect information; they place selective bets, make trade-offs quickly, and own the consequences transparently. Leaders who've sought out disruptive or transformational roles in the corporate world often are already conditioned to take meaningful risks without certainty. Consider staffing decisions. In the corporate world it can take many months to hire senior executives and get them up to speed, but PE firms demand that CEOs make rapid calls about who can deliver on the value-creation plan and who cannot—decisions that carry greater risk because the timeline to create value is shorter and the margin for hiring mistakes is smaller. Portfolio-company CEOs have to home in on the seasoned functional experts who they believe can perform immediately and who need minimal onboarding. They must take calculated risks by selecting leaders who can help the business reach its next stage, even when the exact capabilities required are still coming into focus. As Maggie van de Griend, the managing director for portfolio talent at Warburg Pincus, told us, “In private equity you're hiring for where the business needs to be two years from now, not where it is today.”

Interpersonal range. Corporate environments often demand polish, an understanding of hierarchy, and an ability to carefully manage communication. PE environments, on the other hand, demand clarity, candor, and the ability to work well with a wide variety of people. They require leaders who can quickly grasp people's motivations, understand how their own styles and decisions affect others, and adapt their approaches to galvanize people and resolve conflict when it arises. The portfolio-company CEOs we interviewed put a particular emphasis on building strong working relationships with their investors and boards, which requires

frequent, informal, and unscripted interactions. They often described it as one of the most meaningful differences from corporate roles, where board interactions tend to be more formal, episodic, and filtered. Greg Gartland was struck by this difference when he began working as the CEO of 3E, a PE-owned provider of data-driven intelligent compliance solutions in the environmental, health, safety, and sustainability space, after serving as the chief product officer at S&P Global Market Intelligence. “I was involved in nearly every S&P board meeting for three years,” he told us, “but only on specific topics for short windows. In PE, I'm on the phone with the board every day.”

What Success Looks Like

Consider the case of Lisa Utzschneider, whose transition from corporate C-suite executive to CEO of a PE-backed company highlights three of the key predictors of success: a practical commercial orientation, the ability to achieve broader influence and impact, and interpersonal range.

Utzschneider worked for more than two decades in senior leadership at Microsoft and Amazon and then became the chief revenue officer at Yahoo, where she helped stabilize the business and played a central role in its sale to Verizon. Having navigated a complex turnaround and transaction, she wanted a new challenge. So she became the CEO of the PE-backed Integral Ad Science (IAS).

Once at the helm, Utzschneider drew immediately on the operational intensity and commercial rigor she had honed in her corporate roles. “Amazon and Yahoo really groomed me for the PE CEO role,” she told us. She reinforced value creation from day one, made fast, pragmatic calls, and built a relationship marked by high candor and trust with her investors. Her ability to engage them as true partners allowed her to gain their confidence quickly and operate with speed and transparency.

To increase momentum and focus her team directly on commercial objectives, Utzschneider continually surfaced what she calls the “big rocks” at weekly meetings—that is, the company's most critical efforts, tied to both individual C-suite leaders and the pillars of IAS's value-creation plan. She also tapped into the PE firm's CEO network for advice. “If I'd never hired a CTO before,” she recalled, “I'd ask a CEO



LEADERSHIP



Leaders rarely fail because of a lack of talent; they fail because the capabilities required were never understood or tested on either side.

who had done it. If someone had done seven acquisitions in two years, I'd find them and ask exactly how they pulled it off." Utzschneider leveraged the PE ecosystem and its existing playbooks, which allowed her to move faster than she could have moved alone and to influence her organization more broadly and impactfully.

Under her leadership, IAS grew significantly and went public. Reflecting on what enabled her success, Utzschneider told us, "In PE, success depends on the return, the working hypothesis, and the exit expectations. It's a transaction, and clarity on your role is what builds trust, empowers teams, enables speed, and creates value."

Ken Gayer, a corporate-groomed leader who has now served as the CEO of several PE-backed ventures, is a good example of someone whose interpersonal range and practical commercial orientation have contributed to his success. Before working at firms owned by private equity, Gayer spent 15 years at Honeywell, ultimately becoming the business president of Honeywell Specialty Products. Earlier, he had served as a lieutenant in the U.S. Navy's nuclear submarine force, an environment defined by precision, hierarchy, and strict boundaries, and as a manager at McKinsey & Company, where he sharpened his problem-solving and communication skills. Those experiences shaped him into a disciplined operator and a highly effective "corporate soldier," but they also reinforced relationship norms concerning distance, structure, and formality.

His first role at a PE-owned firm challenged those norms right away. When Gayer became its CEO, the company's founder remained the chief technology officer and the emotional center of the business—an arrangement that required Gayer to balance respect for the founder and the company's origins with the need to professionalize and grow it. A key early success came from his practical commercial orientation; he was able to successfully translate the founder's vision into actionable, economics-driven plans, for instance, by driving growth through acquisitions. "I was good at analyzing and presenting things in a way that made it easy for the private-equity investor to say, 'That's actionable. We can help you with that.'"

Gayer discovered that many of the interpersonal norms that had served him well in his corporate roles were in fact counterproductive in a PE-backed business. "At a

corporation," he told us, "once you got past your immediate boss, the next level up became distant. There was often an antifraternization mentality." In PE the opposite held true. Success depended on being visible, approachable, and directly engaged at every level of the organization. So Gayer leaned in. At one portfolio company, when he encountered a lack of the kind of basic infrastructure that's taken for granted in the corporate world, such as a technology system for running virtual meetings and a standard template for nondisclosure agreements, he didn't wait for it to be built around him. Instead, he adopted what he called a "zero-to-infinity mindset," stepping in proactively to solve problems and create operating mechanisms that allowed others to move faster.

Utzschneider and Gayer's experiences both illustrate a consistent finding in our research: Corporate leaders who succeed in PE are those who can translate commercial insight into action, galvanize results across the system, and use their interpersonal abilities to meet the demands of a far riskier, fast-moving environment.

AS THE PATHWAY from corporate leader to CEO of a PE-owned company becomes more common, success will hinge on a shared understanding of what the role truly requires. When a leader fails to make the transition well, it's rarely because of a lack of talent; instead, it's because the capabilities required were never fully understood or tested on either side. The CEO role is fundamentally distinct in PE-backed enterprises. Investors must move beyond scouting for candidates with familiar résumés and track records and ask more-precise questions about readiness for this environment. Corporate leaders, in turn, must view the transition not just as a promotion but also as a redefinition of expectations—and seriously consider whether they are built for, and motivated by, the realities of the role. Getting it right starts with clarity about both capability and motivation. ■

HBR Reprint R2604D



SAMANTHA HELLAUER is the director of strategic initiatives and client solutions at ghSmart, a leadership consultancy.

DINA WANG is a partner at ghSmart, where she sits on the executive committee. **HEIDI SMITH** is the leader of ghSmart's CEO succession practice. **SAMANTHA SMITH** is a doctoral candidate at Harvard Business School and a research fellow at ghSmart.

AI CAN INFORM DECISIONS. *But Leaders Must Make Them.*



In a world shaped by artificial intelligence, organizations still need leaders who can think critically, exercise judgment, and act with clarity. At the University of Chicago Booth School of Business, we believe better decisions drive better outcomes.

Through Nobel Prize-winning research, world-renowned faculty, and evidence-based learning, Chicago Booth Executive Education helps leaders and organizations turn insight into enterprise impact—through custom programs for organizations and transformative learning experiences for individual executives.

Request a Custom Program Consultation
chicagobooth.edu/custom-program

+1 312.464.8732
exec.ed@chicagobooth.edu

For Individual Executives
chicagobooth.edu/execed



CUSTOMER
EXPERIENCE

Why the ★ Best ★ Immersive Experiences Succeed

They invite deep engagement
by answering key questions
people ask themselves.



AUTHORS

Joseph C. Nunes
Professor, Marshall
School of Business

Wendy Heimann
Cofounder,
Nolan Heimann







CUSTOMER EXPERIENCE

Previous spread: Super Nintendo World at Universal Epic Universe in Orlando; this page: Le Signe du Triomphe at Puy du Fou in France.



IDEA IN BRIEF

THE PROBLEM

Immersive experiences—from branded environments to themed entertainments—are booming, but many fail because their designers attempt to create immersion merely with spectacle, technology, or elaborate features.

THE INSIGHT

Immersion is a psychological state of absorption that emerges when people can answer six questions as they move through an experience: Where am I? Who am I with? What can I do? What is happening? Am I making progress? Why does this matter?

THE SOLUTION

Creators should design experiences that guide participants through these six psychological dimensions. When people can answer all six questions, an experience will deeply engage them and have a memorable impact.





The important question is whether an experience draws people in, keeps their attention from start to finish, and culminates with a coherent sense of what it all meant.



Immersion

is redefining the consumer experience. On any given day, a visitor might enter Netflix House, where the streaming service's hit franchises, including *Stranger Things* and *Squid Game*, are brought to life as walk-through destinations; step into teamLab Biovortex Kyoto, where large-scale artworks created through digital technology shift in response to people's movements; or suit up at Sandbox VR, where teams move together, strategize, and battle apocalyptic zombie hordes or an authoritarian army on an alien planet. Camp has created a chain of toy stores where children step through a "magic door" into play worlds inspired by shows such as *Bluey*, *Paw Patrol*, and *Peppa Pig*. Nike's House of Innovation stores turn shopping into a personalized, interactive sports journey. And Guinness's Open Gate Breweries are engineered to feel less like a tour and more like a multisensory encounter with the brand's flagship stout.

Immersion is being propelled by two converging forces. The first is the rise of the attention economy, in which brands compete for consumers' sustained engagement. Full immersion is the most compelling attention strategy available. The second is the shift of consumer spending toward experiences over the past three decades, giving businesses a strong incentive to turn memorable moments into products. Those trends are likely to pick up steam as immersive technologies—advanced virtual reality headsets, spatial computing, and AI-enabled interactive characters (from holograms to robots)—move from novelties to mainstream offerings.



CUSTOMER
EXPERIENCE

To be sure, immersive experiences are not new. Disney's It's a Small World attraction was created for the 1964–1965 New York World's Fair and still draws crowds at the company's theme parks around the world. Too often, however, "immersive" is a marketing hook equated with elaborate theming rather than absorption. While organizations may be tempted to simply dazzle consumers with technology, the important question is whether an experience draws people in, keeps their attention from start to finish, and culminates with a coherent sense of what it all meant.

Brands realize clear payoffs when they do these three things. According to leisure consultancy AECOM and the Themed Entertainment Association, Universal saw attendance grow more than 30% at its theme parks in Japan, Florida, and California when the Wizarding World of Harry Potter opened, with the attraction's revenue partly driven by sales of multiday passes and repeat visits—a sign of the ability to engage guests. Nike's CEO, Elliott Hill, has said that the company's House of Innovation in New York has generated double-digit revenue increases. *Time* named Netflix House to its World's Greatest Places of 2026 list, which highlights innovative and culturally significant destinations; Netflix hopes to open up to 60 locations globally to deepen its relationship with audiences. Meanwhile, researchers from Cornell University and the Wharton School found that after customers visited a beauty company's store that provided immersive experiences, over the next 12 months they spent more, purchased more items, and shopped at the company's locations more frequently than similar customers who did not visit the immersive store.

These outcomes are not accidental. They reflect the psychological mechanics through which immersive experiences capture attention, guide participation, and sustain engagement. One of us (Joseph) is a consumer psychologist in academia whose research focuses on the creative industries and who teaches an MBA course on immersive experiences. The other (Wendy) is an attorney and adviser who helps structure deals and license intellectual property (IP) for immersive experiences, collaborating closely with developers, producers, and creative teams from concept through production. With these varied perspectives, we've found that what separates the experiences that captivate from those that disappoint isn't how impressive they look—it's whether

they're grounded in an understanding of the psychology of immersion.

That psychology follows a structured progression. As people enter a new experience, they seek answers, consciously or unconsciously, to six questions: *Where am I? Who am I with? What can I do? What is happening? Am I making progress?* and *Why does this matter?* Each question reflects a distinct psychological dimension. Together they create the underlying structure through which immersion takes hold.

For companies, addressing all six questions can be challenging. Most experiences tend to focus on answering certain questions while overlooking others. True success lies in ensuring visitors can answer all six in order, even if some are addressed more effectively than others. Of course, companies will inevitably face trade-offs. For example, an experience may excel at orienting guests in where they are but fare less well at showing them what is happening or what they can do. Yet progressing according to the sequence is critical, because allowing earlier questions to go unresolved will undermine everything that follows.

THE CORE STRUCTURE OF IMMERSIVE EXPERIENCES

Let's look at some experiences that address a particular psychological dimension of immersion exceptionally well, following our step-by-step progression.

Presence: Where am I? When people enter a new environment, they first orient themselves to their surroundings, a cognitive response rooted in evolutionary biology. In constructed environments, presence is the illusion of being in a setting, and maintaining the illusion requires a design that reshapes people's perceptions. By enveloping visitors' field of vision and layering directional sound and other sensory cues (which VR excels at), experience designers can create a world that visitors will inhabit.

The desire to establish a strong sense of presence is why many immersive experiences begin with a passage. (Think, again, of It's a Small World.) The goal is to set up a transition to another world. The Paddington Bear Experience begins with a simulated train trip, transporting guests from London to a storybook land. Other experiences surround visitors with the new world instantly. The shared-reality entertainment



What separates the experiences that captivate from those that disappoint isn't how impressive they look—it's whether they're grounded in the psychology of immersion.



Opening spread: Top/Vectors/Getty Images (star), Super Nintendo World © 2026 Universal Studios; second spread: Le Signe du Triomphe © Arthur Aumond; this page: teamLab, Forest of Resonating Lamps: One Stroke—a Year in the Mountains © teamLab

company Cosm uses large-scale LED displays and spatial audio to give people the sensation of, for example, sitting on the 50-yard line at Arrowhead Stadium for a Kansas City Chiefs football game while physically being in Los Angeles or Dallas. The effect is less about escape than about orientation. By drowning out competing sensory reference points, designers concentrate guests' attention on why they are there: to feel present at a sporting event. That approach seems to be working, as Cosm has raised more than \$250 million in new funding to expand its shared-reality venues globally. Before anything meaningful can occur in an immersive

Forest of Resonating Lamps: One Stroke—a Year in the Mountains at teamLab Biovortex Kyoto

experience, visitors must feel grounded at a sensory level and be able to say, "I know where I am."

Discourse: Who am I with? Humans are social creatures. Research in cognitive psychology and neuroscience suggests that we understand our role in social contexts through *discourse*, the back-and-forth communication that lets people calibrate who they are relative to others and how they're meant to interact.



CUSTOMER
EXPERIENCE

That's why engagement deepens when the performers in an experience converse with visitors rather than treating them as passive spectators. For most of history, mass experiences, from theater to cinema, have largely been built around scripted, unidirectional performances. Truly immersive experiences, by contrast, unfold as dialogues, incorporating audience members' responses into exchanges in which people's attention and reactions are acknowledged. British pantomime, a centuries-old form of theater, illustrates how discourse fosters engagement, as audiences shout warnings, argue with characters, cheer, groan, and collectively shape how the story is experienced in real time. In the Paddington Bear Experience, visitors to Paddington's home help actors solve problems, such as by finding ingredients for marmalade sandwiches. At Cosm's immersive sporting events, discourse emerges from the shared excitement of attendees with well-known team allegiances. In Faulty Towers the Dining Experience, guests are cast as diners in a dysfunctional hotel restaurant where performers respond to their remarks and tease them, and the ensuing chaos (collisions, mistakes, and meltdowns) is based on their reactions. The performance does not simply include the audience members; it adapts to them, largely through improvisations that are based on their contributions. The dining experience has been running since 1997, has welcomed more than one million visitors, and has been performed in some 43 countries.

Agency: What can I do? In psychology, self-determination theory explains that when individuals feel in control of their actions, their intrinsic motivation and well-being increase. Conversely, when people feel they lack control, motivation diminishes and enjoyment declines. Immersive experiences are most powerful when carried out within clear constraints—within conditions that make consequences visible and allow participants to see how their choices shape the course of events.

Monopoly Lifesized illustrates that dynamic. Escorted by hosts who act as “tokens” in a team-based competitive game, participants traverse a giant Monopoly board and complete challenges that are tailored to different audiences,



From top to bottom: Inside the Dome at Cosm in Dallas; Interstellar Arc at AREA15 in Las Vegas; Monopoly Lifesized in London



Truly immersive experiences unfold as dialogues, incorporating audience responses into exchanges in which people's attention and reactions are acknowledged.

from families to corporate clients. Decisions such as how to approach a task, when to collaborate, and whether to take certain risks produce immediate and observable outcomes. Participants quickly realize that what they do matters. Since opening five years ago, the experience has had more than a quarter of a million visitors in the United Kingdom and has toured the United States and other markets. Its creator, Path Entertainment Group, plans to open two new locations and is preparing to launch a live experience based on the board game Clue, following an expanded partnership with the game maker Hasbro.

Another example of the dynamic comes from Secret Cinema, which creates large-scale immersive productions built around well-known films like *Grease* and *Blade Runner 2049*. Participants are assigned characters and given latitude over where to go and whom to engage with, making choice an explicit part of the experience. While each production is structured around a film, the decisions that people make influence which scenes, characters, and activities they encounter, supporting a sense of autonomy within clear boundaries. Founded in 2007, Secret Cinema has sold more than one million tickets and was bought in 2022 for more than \$100 million.

Understanding: What is happening? Feeling as though one's actions have consequences is not enough for an experience to be deeply immersive. Visitors need a way to grasp what their actions mean, which requires an understanding of how events are unfolding over time. People don't absorb an experience as one continuous stream; they register it as a series of narrative beats and turning points. Understanding comes from how clearly those beats connect.

The connection often takes the form of a story. Here, "story" does not mean exposition or scripted dialogue; it refers to temporal coherence, the linking of what is happening now, what has already occurred, and what visitors expect will come next. Even when events are encountered out of the expected order or discovered piecemeal, coherence emerges as visitors assemble the fragments into a meaningful whole. When the connections are clear, people are more likely to experience narrative transportation, a state that helps them process events as unified wholes rather than as isolated parts.

Louis Vuitton's renovated Matsuya Ginza store, its flagship location in Tokyo, leads customers through a carefully sequenced journey across four floors. The journey explores different facets of the brand, from heritage to innovation, and creates a sense that each space builds on the last, encouraging visitors to understand the brand as an unfolding story.

The mixed-reality experience Interstellar Arc in Las Vegas illustrates the principle of understanding more explicitly. Participants find themselves aboard a spacecraft, searching for a new home for humanity, confronting dilemmas that pit short-term survival against long-term goals, and collectively shaping outcomes. As decisions about resources, relationships, and priorities carry forward from one scenario to the next, events do not reset—later situations reflect earlier choices. Because actions have consequences and individual moments connect, participants are able to say, "This makes sense."

Coherence is needed not only for enjoyment but also to provide a cognitive scaffold that helps people grasp what they have encountered, remember it, and recount it to others. Experiences that achieve integration are easier to describe afterward, amplifying their impact through word of mouth.

Goal pursuit: Am I making progress? Decades of research have shown that goals help focus attention, mobilize effort, and foster persistence. Researchers studying goal-setting theory find that individuals continuously monitor their current position relative to a goal, adjusting their behavior to reduce the gap between where they are and where they want to be.

Immersive experiences sustain attention by giving visitors the sense that they are moving forward and by providing feedback along the way. That's why techniques borrowed from gamification—such as points, badges, levels, and progress markers—work so well in immersive experiences. Without clear signs of advancement, even impressive environments can feel stagnant. But when visitors can say, "I am making progress," advancement turns into momentum and motivates them to continue.

At Universal's Super Nintendo World in Orlando, guests wear Power-Up Bands that track their achievements as they complete physical challenges. They accumulate digital keys that unlock new interactions and, eventually, gated



While many experiences underemphasize the meaning-making dimension, the most powerful are anchored in a theme that shapes the perspective people take away.

experiences such as the Bowser Jr. Shadow Showdown. Progress is visible as guests can see what they have accomplished and what remains ahead, reinforcing the idea that continued involvement leads somewhere new. This example also demonstrates that basing experiences on well-known IP, such as Nintendo's games and characters, helps strengthen multiple psychological dimensions at once by providing a recognizable and instantly coherent world.

Progress takes a different form at Phantom Peak, an immersive Western steampunk-themed town in London. On the surface, advancement comes from completing quests by following clues, answering questions, and closing narrative threads. Over repeat visits, however, people grow familiar with the town's characters and patterns, and navigation becomes easier. A sense of knowing how this world works—built by completing different “trails” or quests that change from season to season—makes each visit richer and more satisfying than the last. The town has welcomed more than 160,000 visitors since launching in 2022, and around 25% of guests have returned.

Meaning-making: Why does this matter? Transformational experiences change how people see themselves or the world. People naturally engage in meaning-making when an experience challenges their existing beliefs or expectations, evokes strong emotion, and taps into their sense of who they are.

The performance-based historical theme park Puy du Fou in western France shows how experiences can reshape people's understanding of history and their place within it. Guests move through a series of large-scale live performances spanning different eras, from Roman spectacles to medieval life to 20th-century warfare. Rather than presenting history as something to observe from a distance, the park places guests inside moments of decision: whether to charge into battle, flee a burning village, or hold the line as losses mount. Structuring history as a chain of choices rather than a timeline helps guests leave with a deeper sense of how decisions create lasting consequences. The park welcomed a record 3 million visitors in 2025, and Puy du Fou is planning a £600 million immersive historical theme park in the United Kingdom.

A similar transformation occurs at the National September 11 Memorial & Museum in New York City, where visitors

descend from a public plaza into the underground remains of the Twin Towers. Along the way, personal artifacts, recorded voices, and structural fragments shift the experience from remembering a day of tragedy to more deeply understanding the human impact and enduring aftermath. Visitors walk away believing that remembrance itself is a civic act, one that transforms loss into collective resilience.

While many experiences underemphasize the meaning-making dimension, the most powerful are anchored in a higher-order theme that shapes not just what participants see and do but the perspective they take away. One message of the Paddington Bear Experience, for example, is that even small, individual contributions (such as helping to make marmalade sandwiches) can collectively calm disruptions and restore order. Super Nintendo World conveys that mastery is earned through repeated effort, while Monopoly Lifesized shows that advantage can accrue to those who get ahead early. An experience's success is sustained, in part, by offering a distinctive lens on the world.

PUTTING THE SIX PSYCHOLOGICAL DIMENSIONS TO WORK

When immersive experiences underperform, often the problem isn't a lack of ambition or budget but a failure to follow the psychological progression that makes experiences captivating. Too often immersive projects emphasize spectacle, technology, or scale without considering how visitors actually move through and make sense of them. For immersion to succeed, experiences must guide visitors through stages that lead from initial attention to deeper engagement and ultimately to meaning. The implications for designers are straightforward but demanding.

While the six dimensions are interconnected and build on one another, separating them into two groups can make it easier to understand what they accomplish.

Entering and engaging. The first three dimensions—presence, discourse, and agency—establish how visitors enter and begin engaging in an experience.

You can help visitors feel situated rather than disoriented in several ways. First, create a clear threshold moment. Use lighting shifts, sound changes, architecture, staging, or other tools to signal that the visitors have entered a new world.



Common Mistakes in Experience Design

Many experiences fail not because they lack ambition or investment but because they overlook key psychological dimensions of immersion. The table below highlights common breakdowns.

PSYCHOLOGICAL DIMENSION	KEY QUESTION FOR PARTICIPANTS	MANAGERIAL CHECKPOINT	COMMON FAILURE MODES
Presence	Where am I?	Can visitors clearly identify where they are meant to be?	Visual spectacle without locational cues; environments that look impressive but feel disjointed or distracting
Discourse	Who am I with?	Do visitors meaningfully influence or interact with others?	Experiences designed for individuals rather than groups, leaving participants isolated in shared spaces
Agency	What can I do?	Do actions have observable consequences?	Interactivity without consequence; choices that feel cosmetic; events that continue or reset without individuals' impacting them
Understanding	What is happening?	Can visitors follow what is unfolding over time?	Events presented as disconnected moments rather than as parts of a connected whole
Goal pursuit	Am I making progress?	What markers help visitors see progress?	No visible markers of advancement; repeat visits feel redundant rather than cumulative
Meaning-making	Why does this matter?	Does the experience shape how visitors view themselves or the world?	Experiences that impress in the moment but leave nothing to reflect on afterward

Second, mark the entry with a strong visual cue—such as an object, a character, or a landscape—that immediately defines the space. Third, ensure that all employees' dress and behaviors are consistent with the environment; they should seem part of the world, not separate from it.

Overall, do what it takes to ground people where you want them to be. Guests of Absurdities, the immersive dining experience in Singapore, arrive at an unremarkable café before passing through a nondescript door into a dark,

atmospheric corridor that leads to their first course. At Titanic: The Artifact Exhibition in Las Vegas, guests receive a boarding pass for a real passenger of the famed ship, then encounter a re-creation of its grand staircase. And at Disney's Star Wars: Galaxy's Edge, in-character staff are dressed as citizens of the planet Batuu, a remote trading outpost where visitors are treated like intergalactic travelers.

The discourse dimension—"Who am I with?"—relies on how performers initiate interactions, not just facilitate



To give visitors a sense of agency, at least one decision should visibly alter the experience, as too many environments are overscripted.

transactions. Group countdowns or cooperative tasks can accelerate social bonding and move guests from observers to participants. Actors in the John Wick Experience in Las Vegas say things such as “People are looking for you” (you’re actively participating), “We’re going to hide you” (you’re one of us), and “Stick together, we’re all you’ve got” (we’re working together). This year the production received a Thea Award for Outstanding Achievement from the Themed Entertainment Association, the leading professional organization for immersive experience creators.

To give visitors a sense of agency, at least one decision should visibly alter the experience, as too many environments are overscripted. Where possible, incorporate branching narrative pathways and multiple responses and create the illusion of control—perhaps by timing big reveals to follow guests’ actions—to make participation feel consequential. The Crystal Maze Experience, based on the British TV show from the 1990s, has participants choose a team captain and divide up responsibilities before entering the maze’s historical and futuristic environments. Visitors do not control which challenges appear next, but the group’s organization and actions determine how far members progress. After its 2016 opening, the London location was sold out for a full year.

Deepening, not dissipating. The latter three dimensions—understanding, goal pursuit, and meaning-making—determine whether the sense of immersion deepens or dissipates during an experience.

Fostering understanding comes down to effective storytelling. Give participants details early on that foreshadow key moments; people should have a sense that events are building toward a climactic payoff. But avoid overloading them, because having too much to remember can create friction. Instead, include subtle recap moments: A certain phrase spoken by a performer might reappear later in a new context, revealing its significance and impressing upon people how different elements interconnect. Coins and keys are recurring motifs in Super Nintendo World; in Faulty Towers the Dining Experience, Basil repeatedly insists that the hotel is “a first-class establishment,” sharpening the satire amid the show’s escalating chaos.

Goal pursuit is usually the easiest and most intuitive element for experience designers. Ritual milestone moments

are especially effective, such as boarding a vehicle or receiving a prop at the outset, discovering a twist that reframes events or shifts the environment midstream, or making decisive choices that mark closure. At Phantom Peak visitors pick up story threads from newspapers, telegrams, and characters scattered across town as they piece together the backstory of JONACO, the mysterious and powerful corporation running the town.

Meaning-making requires nuance and subtlety to not appear preachy. Designers should embed a clear core theme throughout the experience rather than relying on a summary statement at the end. Creators should decide on takeaways early on, as those can influence motifs, character arcs, environmental changes, and ritualized moments.

An immersive experience that does well at balancing all six dimensions is Meow Wolf’s Omega Mart in Las Vegas. The environment is a surreal supermarket that opens into a maze of hidden offices, laboratories, and vividly distorted spaces. Visitors encounter in-character staff who interact with them while guiding them deeper into the world. The experience unfolds through self-directed yet collaborative exploration. Groups unlock different layers of content and compare discoveries as they move through the space. Visitors may encounter strange product labels and corporate slogans that reappear in emails or surveillance footage. Hidden rooms and tiered access points encourage visitors to find and decode clues, gradually constructing an understanding of Dramcorp, the shadowy conglomerate at the center of the narrative. Ultimately, visitors must interpret what they have encountered, a provocative commentary on corporate culture and consumerism. The Omega Mart experience, which is Meow Wolf’s second permanent installation, welcomed more than one million visitors in its first year.

At any point in the creation process, producers, designers, and financiers should be able to explicitly audit each dimension. Take choice (agency). List every decision available to visitors and ask: Does it change anything? Is the impact visible? Does that impact build over time? Or test orientation (presence). Whenever visitors ask, “Where am I?” during the experience, can they answer correctly? If the

From top to bottom: Omega Mart in Las Vegas; Phantom Peak in London; the Paddington Bear Experience in London

Top to bottom: Atlas Media; Nathaniel Nor/Alamy; Alex Bremner



CUSTOMER
EXPERIENCE



environment does not ground them quickly, deeper engagement will falter. To examine belonging (discourse), ask yourself: How are visitors acknowledged by others? Do their responses influence what follows? Experiences that neglect discourse may be visually striking but emotionally empty. And don't forget to look beyond the moment. If visitors described the experience a week after attending, how would they say it affected them? If their recall centers on spectacle rather than significance, the design has likely failed to be transformative (meaning-making).

Finally, given that visitors to many immersive experiences exit through the gift shop, merchandise should be carefully thought out. In general, we advise against selling trinkets. Merchandise should function as a symbolic extension of the experience: a prop tied to a pivotal moment, an object that visitors handled, or a token that makes sense only in light of what they went through. When designed well, these items are not retail add-ons but anchors that remind people of time spent immersed elsewhere.

ALTHOUGH OUR WORK (and this article) focuses on immersive experiences, many things in business can be evaluated using the framework we've outlined, whether they're customer journeys, service interactions, workplace events, or moments of organizational change. When people enter any situation, they orient themselves, interpret how they relate to others, test whether their actions have consequences, seek to understand how events unfold over time, look for evidence of progress, and ultimately reflect on what the experience meant to them.

Designing with these six dimensions in mind can improve experiences regardless of their scale, technology, or setting. What matters is whether human psychology is treated as an explicit input rather than as an afterthought. ▀

HBR Reprint R2604E



JOSEPH C. NUNES is a professor of marketing and the Joseph A. DeBell Endowed Professor in Business Administration at University of Southern California's Marshall School of Business.

WENDY HEIMANN is a cofounder and a managing partner of the law firm Nolan Heimann, leading the firm's practice dedicated to location-based, experiential, and immersive entertainment sectors worldwide.



LEADERSHIP

Six Ways Leaders Harness Stress

Executives perform best
when they understand their
default reactions—and broaden
their range of responses.



AUTHORS

Jon Miller
Partner,
Brunswick Group

Drew Keller
Founding director, Institute for
Business in Global Society



PHOTOGRAPHER **JAMIE CHUNG**





IDEA IN BRIEF

THE OPPORTUNITY

Today's leaders are operating under rising levels of pressure—which can cause not only burnout but poor decision-making at precisely the moments that matter most.

THE INSIGHT

Under stress, leaders default to recognizable response patterns that shape how they interpret situations and make tough calls. Those patterns each have different strengths but also have predictable blind spots.

THE PLAYBOOK

High-performing leaders develop "stress intelligence": the ability to recognize their default patterns and deliberately shift their approach to match the demands of the moment. By expanding their range, they make better decisions and sustain performance under pressure.



Previous spread, inside typography: disk/Getty Images; all other photos: Jamie Chung/Trunk Archive



ABOUT THE ART

Jamie Chung photographed the blacksmithing forge at the Canadian Museum of Making, where traditional techniques and contemporary designs come together in the creation of hand-forged pieces.



LEADERSHIP

Leaders have always lived with stress—it comes with the territory. But today's strain feels especially acute. Many executives describe pressure mounting on multiple fronts, from economic uncertainty to technological disruption to geopolitical volatility.

We've studied these sources of stress over the past six years, and the trend is unmistakable: According to Brunswick Group's Leadership Stress Index, which tracks the pressures facing leaders, stress levels are even higher now than they were at the peak of the pandemic.

In our September 2025 survey of 400 C-suite leaders in large corporations in the United States, the UK, Germany, and Japan, more than half the respondents said they were very stressed. Nearly two-thirds reported that their stress level had risen over the past year, and a majority anticipated that it would increase further in the coming year.

While short bursts of stress can sharpen your focus, sustained, elevated levels of it undermine decision quality and erode well-being. Over time, stress narrows your perspective, causes reactive thinking, and increases the likelihood of costly misjudgment.

In our work with leaders—one of us as a partner at a global strategic consulting firm and a practicing integrative therapist, and the other as the former director of the Institute for Business in Global Society at Harvard Business School—we've helped executives across sectors and geographies navigate stress on the job. We've seen firsthand how differently individuals behave when the heat is on and how a greater awareness of their own default responses can improve their performance and help them be more resilient and effective.

Although no two leaders have the same biopsychological response to stress, consistent patterns do emerge. We've developed a framework that organizes them into six main response types. In creating it, we drew on the work of Richard Lazarus and Susan Folkman, who observed that when people are in a high-stakes situation, they quickly decide whether it presents more of a threat or an opportunity. We combined their insights with James Gross's research on emotion regulation, which distinguishes between leaders who regulate their emotions dynamically and continually reappraise their reactions, and others who hold their emotions in tight check until after a crisis. Synthesizing those concepts with our own experience, our framework positions stress responses along two spectrums: whether leaders react with composure or by leaping into dynamic action, and whether they approach a crisis as an opportunity or a threat. Where you fall along those spectrums determines which of the six response types is your default pattern.

Our goal here was not to create a scientific taxonomy but to help leaders understand how they operate under duress and make better decisions when it matters most. In this article we'll take a closer look at the six patterns—how they show up in moments of strain, what strengths they offer, and where each can go wrong under sustained pressure.



“I thought my job was to be the calm center. But my team later told me it felt like I wasn’t letting them in.”

The Lighthouse

Lighthouse leaders project calm when others are rattled. In a crisis they maintain a measured pace, regulating their breathing, lowering their voices, and staying focused on the horizon—not just on the crisis immediately in front of them. Stress triggers their orientation toward *stability over speed*, and their first rapid appraisal of a new situation is more likely to recognize opportunities than threats. Instead of matching urgency with urgency, they tend to hold their ground and steadily guide others through ambiguity.

In high-pressure situations this style of leadership creates conditions that promote psychological safety, where people feel free to speak up, admit mistakes, and express uncertainty without fear of humiliation or punishment. Studies of high-reliability organizations—from air-traffic control to hospital emergency units—show that calm leadership keeps teams focused and prevents what psychologists call the “emotional contagion of fear.”

Several years ago we worked alongside the leadership of a consumer goods company when a politically charged boycott erupted over its stance on a divisive issue and was amplified by a social media campaign. The CEO responded like a lighthouse, acknowledging the potential damage—the share price had taken a hit—but reinforcing that the company had taken that stance because of its underlying values. Behind the scenes the board expected the CEO to make a rapid reversal, but she resisted. Some criticized her calm response as sounding tone-deaf, but it reassured most stakeholders, and the storm passed.

As this story shows, lighthouse leaders face hidden risks. This CEO’s composure appeared to some to be aloofness. And teams sometimes interpret lighthouse leaders’ silence on seemingly critical issues as indecision. In prolonged periods of stress, their drive for stability can harden into inertia; they may fail to face up to bad news, lack the appropriate urgency, or quietly take on too much themselves. One senior executive at a global retailer captured this tension perfectly: “I thought my job was to be the calm center. But my team later told me it felt like I wasn’t letting them in.”

Adaptive strategies for lighthouses: During a crisis the first step these leaders should take is to explain their intentions. For example, they could tell their teams: “I’m staying

calm so that we can think clearly,” not “I’m calm because everything’s fine.” Next they should invite their teams to share their emotions and should reciprocate in kind. They might ask how others are coping, which will legitimize the stress people are feeling rather than suppress it, reinforcing psychological safety. Finally, they must signal momentum. Even calm leaders need visible microactions—small decisions, next steps—to reassure others that their stability isn’t masking inertia or indecision.

The Alchemist

Alchemists see pressure not as a threat but as a catalyst for growth. In a world defined by overlapping disruptions, they treat turbulence as the raw material for reinvention—a signal that the old logic no longer fits. Alchemists’ responses are a combination of composed and dynamic: Outwardly they’re steady, but inwardly they’re kinetic. They think in ways that help them connect dots others don’t yet see. Under duress, they look for possibility, not predictability. They’re the kind of leaders who hear, “The playbook doesn’t work anymore” and reply, “Great—then we get to write a new one.”

Alchemists gravitate toward frontier spaces where disruption and opportunity meet—like AI, climate tech, and fields where global trade and talent patterns are changing. Netflix cofounder Reed Hastings captured the alchemist’s blend of realism and optimism when he said: “Be brutally honest about the short term...and optimistic and confident about the long term.”

A few years ago we worked with a company that was thrown into turmoil when a whistleblower revealed that data on its discharge of industrial wastewater had been systematically falsified. The news triggered threats of litigation and undermined the company’s integrity. When it broke, the CEO responded like an alchemist, launching an ambitious program of structural and cultural reform, overhauling incentives, and simplifying reporting layers. He seemed to relish the opportunity to break from the company’s past narrative. Although its external stakeholders responded well, behind the scenes many parts of the organization were unsettled by the pace and scale of the transformation.

This example illustrates the power and risks of the alchemist style. Under prolonged uncertainty, alchemists can slip





from creative to chaotic, chasing novelty for its own sake or exhausting teams with constant pivots.

Adaptive strategies for alchemists: Under pressure alchemist leaders should fix on a North Star. They need to know why they're making a transformation—"because we can" isn't enough—and define what problem the change is meant to solve. And instead of running back-to-back transformations that leave teams perpetually in motion, they should establish a regular cadence for change: cycles of push and pause, invention and consolidation. Finally, they must moderate their tendency toward reinvention and creativity by collaborating with pragmatic teams that can organize and execute transformations skillfully.

The Firefighter

Firefighters thrive on pressure and momentum. They favor decisiveness over prolonged analysis. When a crisis erupts, they move fast, radiate urgency, and mobilize others. They not only respond to stress very dynamically but tend to see it as an opportunity. Pressure doesn't paralyze them; it ignites them. Often fueled by adrenaline, they restore a sense of empowerment when others freeze. In the short term this energy can be invaluable: Momentum builds, decisions get made, and complexity feels manageable.


Where the alchemist uses stress to spark creativity and explore new possibilities, the firefighter uses it to spur immediate action. The alchemist asks, What could this make possible? The firefighter asks, What must we do right now? Both convert pressure into energy, but one channels it into experimentation, the other into execution.

Firefighters embody *in extremis leadership*—a term coined by researchers at the West Point military academy to describe decision-making amid threats, urgency, and uncertainty. Research on "action bias" helps explain their appeal: Action itself reassures both the leader and the team that progress is underway. As the organizational theorist Karl Weick, who popularized the story of some soldiers who got lost in the Alps and found their way back to camp using a map of the Pyrenees, noted, "When you are lost, any old map will do."

However, under prolonged strain, a firefighter's immediacy can become problematic. Every spark starts to look

Six Patterns of Stress Response

Under pressure, leaders tend to gravitate toward one of six recurring patterns—reactions shaped by how they read the moment and how they regulate themselves.



like a blaze, and decisions can become impulsive. As one former COO reflected, "You feel most alive when you're fighting fires, but you have to be careful you don't start fanning the flames."

We saw this play out while working with a company whose production line stalled because its supply of chips fell dramatically when a severe drought hit a critical semiconductor-manufacturing region. It was impressive to watch the firefighter CEO swing into action, setting up a global command structure, authorizing aggressive procurement, green-lighting rapid redesigns, and chartering aircraft to fetch scarce components. It also worked: The company resumed partial production far faster than its competitors did, but the speed of decision-making outpaced cost controls. And we also saw a hidden cost: Teams struggled as priorities shifted weekly, and burnout spiked among employees across operations.

When balanced, the firefighter embodies the best of decisive leadership: courage, momentum, the will to act when others hesitate. However, great leadership is measured not



“Sometimes I would rather have a slightly wrong decision that everyone supports than the right one that divides the team.”

only by the flames extinguished but by what is preserved and the future fires prevented.

Adaptive strategies for firefighters: In a crisis the first step firefighter leaders should take is to pause before acting—to get set, then go. A short break will give them a chance to regulate their nervous systems and clarify their intent to their teams. Next they should involve others in decisions early on. That will distribute the cognitive load in a way that will help them sustain their energy over time and prevent burnout. Finally, they should hold quick after-action reviews on a situation once the immediate urgency has passed, asking: What worked? What didn’t? What felt productive but actually wasn’t? Such reflection can turn the adrenaline spikes into learning and help release the stress people have been carrying.

The Stoic

Stoics lead through discipline and self-control. When the pressure spikes, they get analytical—anchoring themselves in principles and reason and striving to create stability. They have a composed response to stress, sharing the lighthouse’s external calm but for different reasons: While the lighthouse uses steadiness to create space and perspective, the stoic uses it to manage internal volatility. The lighthouse, regarding stress as an opportunity, looks outward, scanning the horizon; the stoic perceives stress as a threat and looks inward, tightening focus and minimizing emotional reactions. This mindset can often be traced to early professional training—for instance, as an engineer, a lawyer, or a military officer. The result is a distinctive style under pressure: calm, deliberate, and rational to the core.

This orientation has real strengths. When a crisis hits, stoics naturally slow down, regulate their emotions, and reintroduce order. Behavioral research shows that this approach supports clearer reasoning under uncertainty. The stoic’s discipline can therefore preserve strategic coherence when others are spiraling. General Motors CEO Mary Barra captured this ethos when she told *The New York Times*, “I don’t focus on what you can’t control.”

Consider a stoic we once counseled. Days before his company’s shareholders were due to vote on a major acquisition, a deepfake voice recording surfaced, seeming to be the CFO

admitting to large, hidden off-balance-sheet liabilities. As that disinformation spread, the deal was thrown into doubt. The stoic CEO told his teams to ignore the noise and speculation and focus on establishing the facts. He surprised the market by halting the vote and commissioning independent forensic audio experts to evaluate the recording. It was a nail-biting time. It looked as if the delay might kill the deal, and the CEO was criticized by many—but the voice on the recording was proved to be synthetic, and the acquisition eventually went through.

Treating emotion as a distraction may lead stoics to withdraw just when human connection is most vital, however. Their desire to stay rational can push them deeper into data and process, producing technically accurate but emotionally clueless decisions. And the costs are often internal: Unexpressed stress can cause insomnia, fatigue, or quiet burnout.

For stoics, the developmental challenge is ensuring that their steadiness remains advantageous—a source of clarity, not a defensive shield that closes them off from insight or support.

Adaptive strategies for stoics: In a crisis, stoic leaders must first regulate themselves—steadying their own emotional responses without suppressing them. Then they can make the emotional climate open to discussion, articulating what the tension in the room may be signaling: “I sense frustration building; let’s explore why.” Next they should visibly practice empathy. Briefly voicing shared difficulty (“I find this decision tough, too”) can strengthen trust without surrendering authority. Finally, they should acknowledge their own emotions. That often means having a conversation about the pressures they’re facing with a trusted adviser such as a coach or a peer, turning stress into a source of insight rather than a silent load.

The Diplomat

Diplomats lead with tact. They ease tension through dialogue and keep teams connected when stress threatens to divide them. For these leaders, a crisis is as much a social test as a strategic one. Under pressure they foster social connection—defusing conflict, sustaining trust, and reading interpersonal signals as carefully as others read financial data.





The “fallacy of uniform efficacy”—the belief that what worked once will work every time—can be quite dangerous.



LEADERSHIP

Diplomats respond to stress in a manner that’s composed but dynamically attuned to relationships. Under strain, their social intelligence helps them hold together teams that might otherwise fracture. Diplomats sense tension before it surfaces and soften it by enabling what psychologists call *relational regulation*—which happens when connection helps people manage emotions. Emma Walmsley, who led the biopharmaceutical firm GSK through a difficult period of industry change, sounded like a diplomat when she said to *Fortune*, “It’s not about being a superhero...It’s about the team, it’s about the company.”

A few years ago investigative journalists uncovered forced labor deep within the extended supply chain of a company we worked with, triggering NGO campaigns against it and possible retailer defections. When the story broke, the CEO responded in the classic style of a diplomat: She avoided a defensive response and instead acknowledged the company’s role in the ecosystem, moving quickly to engage stakeholders and commission independent audits. This swift, conciliatory posture took the heat out of the situation—although internally, some in the management team felt that relationship repair had been given precedence over legal risk.

Indeed, under sustained stress, diplomats may tilt toward consensus over candor—smoothing over conflict rather than addressing it directly. Their composure can become caution, and their diplomacy can slide into delay. As one executive who identified with this pattern admitted, “Sometimes I’d rather have a slightly wrong decision everyone supports than the right one that divides the team.”

Adaptive strategies for diplomats: At their best, diplomats aren’t conflict-avoidant but understand that rupture and repair are healthy for relationships. Their first step in a crisis is to surface the tensions. That means saying plainly what others sense but won’t voice, naming the discomfort in a way that clears the air and creates space for honesty without blame.

One practical way to do that is to build dissent into a company’s crisis management process—through debate, premortems, or red-team exercises (in which an internal team plays the role of devil’s advocate)—so that disagreements are brought to light in a structured way. Finally, diplomats need to be clear about where they stand. Their empathy helps others regulate themselves, but in times

of stress people also need well-defined objectives and direction—to experience the safety and reassurance that come from knowing what the leader is guiding them toward and what happens next.

The Container

Containers lead with control. They respond to stress by imposing structure and holding the system together when others are overwhelmed. They excel at managing multiple demands coherently. Like firefighters, they respond to stress dynamically, but they see it as a threat, and their instinct under pressure is to convert ambiguity into clear priorities and channel energy into disciplined execution. They often rely on a tight inner circle—a small, trusted team they can move decisively with while filtering out noise.

Containers’ mindset is captured in the mantras we’ve heard many of them repeat: “strong and steady,” “left foot, right foot,” “you can’t outrun a storm.” Intel cofounder Andrew Grove’s famous line “Only the paranoid survive” captures the container’s vigilance.

When one company’s internal networks were frozen by sophisticated ransomware attackers, we witnessed a classic container response from its CEO. A big risk in this moment was panic: Alongside the urgent need to restore the systems, the situation triggered immediate simultaneous obligations to regulators, insurers, and customers—which threatened to overwhelm decision-making. The CEO imposed clear crisis structures with tightly defined decision rights and a highly controlled information flow. This military-level discipline succeeded in preventing confusion while a recovery was engineered, but behind the scenes parts of the organization were left anxious and disoriented because the response happened behind closed doors.

This example shows how containers can be invaluable in a crisis. But their intense focus may exclude voices outside the inner circle, leaving teams feeling shut out or underinformed. When the pressure builds, the container can absorb too much of it, too quietly.

Adaptive strategies for containers: The challenge for containers is learning that sustained performance depends not just on endurance but on knowing when to release pressure, widen the lens, and let others help carry the load.



Crises rarely announce themselves with clarity. In such moments stress-response patterns shape what leaders notice, how quickly they move, and where their blind spots lie.

Their first step in a crisis should be to deliberately blow off pressure—to use reflection, feedback, and recovery to allow themselves to reset. They should also routinely zoom out, engineering systematic, regular shifts in perspective on their teams so that their organizations aren't operating with tunnel vision. Finally, they need to model resilience. None of us is indestructible, and teams take their cues from the top. So leaders should visibly make rest and renewal part of performing well under pressure.

From High Pressure to High Performance

Crises rarely announce themselves with clarity. They arrive quickly and demand that leaders interpret incomplete information under heightened strain. In these moments stress-response patterns shape what leaders notice, how quickly they move, and where their blind spots may lie. Preparing for high-stakes decisions therefore requires a good process, including an understanding of how you're likely to react when pressure peaks, based on your default response type. (To identify yours, go to our survey at stressintelligence.org/test.) Here are some examples of how default responses can differ:

- An activist investor targets a company's leadership at an underperforming corporation. Lighthouse leaders may treat the scrutiny as clarifying, engaging shareholders with a structured, transparent review; container leaders, on the other hand, may convene a war room, refining a comprehensive end-to-end response before making a single definitive revelation.
- A major cybersecurity breach exposes customer data. An alchemist may use the shock as an impetus to push through a modernization program, while the first instinct of a diplomat may be to engage the media, regulators, and customers to restore trust and confidence.
- A fresh geopolitical rupture—such as new sanctions, export controls, or a market shutdown—hits key geographies. A firefighter leader will move immediately, rerouting operations and signaling decisive action; a stoic, by contrast, will hold steady, monitoring signals until the strategic path becomes clear and then executing with precision.

There are no intrinsic rights and wrongs here. Each of these responses has its strengths and weaknesses. Leaders

who consistently turn pressure into performance don't get stuck in a single default response style; they treat leadership under stress as something that can be practiced and refined—experimenting with different styles, sensing how the situation is evolving, and recalibrating as they go. We can break this down into a few clear steps:

Expand your range of responses. Deeper awareness of how response patterns typically play out will help both individual leaders and teams recognize more choices. To develop it, consider creating a brief playbook of responses or an AI decision-support agent that tracks your response style and suggests alternatives. During crisis simulations or even real events, you can deliberately try on a style that's not your default. If you're naturally a stoic, experiment with a firefighter-like surge of decisive action; if you're a firefighter, challenge yourself to pause for a lighthouse-like scan of the horizon before moving. Over time this deliberate cross-training will make it easier to shift between types as conditions change and to avoid the pitfalls of your default response.

Test and adjust in real time. High-stakes situations aren't static—they can shift rapidly and dramatically. Refining your response style over time may be critical, and it's important to ask, Is my approach still working? The “fallacy of uniform efficacy”—the belief that what worked once (staying calm, acting fast, pushing through) will work every time—can be quite dangerous. The seasoned leaders we spoke with instinctively test and adjust their style in a continual-feedback loop. Asking simple questions will help: Is this response reducing or amplifying confusion? Is it moving us closer to our objectives or just making us feel busy? What is my team experiencing from me? Over time this approach will allow you to adjust your internal compass for which response type to lean in to and when.

Develop flexibility. Because each response type offers distinct strengths, the real potential gains come from selectively harnessing them to meet the demands of the situation. Here are some examples:

The stoic approach allows leaders to concentrate on what's within their influence and let go of what isn't. Focusing on the things they can actually change—priorities, behaviors, decisions, cadence—creates clarity and reduces wasted effort.



The container approach shows leaders how tight information flows and clear decision rights can help them weather a storm. They can also use the container's adaptive strategies to interrupt the escalation cycle—in which pressure compounds because you're fighting it rather than acknowledging it—before it damages performance.

The lighthouse approach enables leaders to stay grounded and keep the full horizon in view. It isn't about muting urgency but about placing it in context so that the organization stays focused on what matters most.

The diplomat approach can help leaders recognize that calm travels quickly—and so does tension. Teams pick up



a leader's microsignals, and unmanaged stress erodes trust. Regulating emotion becomes part of protecting relationships and maintaining cohesion.

The alchemist approach allows leaders to leverage a different kind of stress—not the spiraling kind but a focused tension that sharpens the imagination and pushes people toward breakthrough insights. Pressure helps alchemists rethink the problem and spot options they'd miss if they stayed in "business as usual" mode.

The firefighter approach lets leaders channel urgency into mission. Purpose gives adrenaline direction, creating coherence and sustaining performance long after the initial surge fades.

Share the cognitive load. Leaders should not feel they have to shoulder a crisis alone. Some of their cognitive load can be borne by systems and by their colleagues. At an organizational level, that means establishing clear decision frameworks, values, and escalation paths before crises hit so that teams know how to behave under pressure. It means consciously designing leadership teams with a balance of response types—not just a room full of firefighters or stoics—and normalizing the open discussion of stress in executive forums. At a personal level, it means building a trusted group of peers, mentors, advisers, and perhaps a coach or therapist who can help you interpret stressors, challenge your default responses, and widen your options.

AS LEADERS EXPERIENCE rising pressure, those who excel will be distinguished not by immunity to stress but by mastery of their reaction to it. By becoming a student of your own patterns—and expanding the range of responses you can call upon—you can strengthen your resilience in times of duress. In the end, leading well under pressure is less about holding firm to one style and more about knowing when to shift. That adaptability is what turns a crisis into clarity and strain into strength. ▀

HBR Reprint R2604F

 **JON MILLER** is a partner at Brunswick Group and a practicing integrative therapist. **DREW KELLER** was the founding director of the Institute for Business in Global Society at Harvard Business School. Together they have established the Center for Stress Intelligence, which studies how leaders can make better decisions under pressure.



How Agentic AI Supercharges Startups *and* Threatens Incumbents

A new operating model shrinks the time and capital needed to build companies, redefining how ventures launch and scale.



AUTHORS

Vivian S. Lee

Executive fellow,
Harvard Business School

Linda Mantia

Former COO,
Manulife

Jon McNeill

CEO, DVx Ventures



ENTREPRENEURSHIP



ARTIST JAVIER REY

Harvard Business Review
July–August 2026

113



ABOUT THE ART

Javier Rey explores the idea of modern unattainable perfection through altered photographs in his series *Mirages* (Espejismos).



IDEA IN BRIEF

THE DEVELOPMENT

Agentic AI systems are enabling startups to compress product development, customer acquisition, and the time needed to scale up into months, using far less capital and fewer employees than traditional SaaS models.

THE CHALLENGE

Incumbents face structural disadvantages compared with AI-native startups: Siloed data, nonstandardized workflows, suboptimal software designs, and slower learning cycles make competing on speed and cost increasingly difficult.

THE SOLUTION

Incumbents should run a “vulnerability” audit of workflows; redesign processes before automating them; partner with AI-native ventures; strengthen data quality and make it clear when tasks should be returned to humans; and prepare people for shifting roles.






Startups can enter or expand markets at a pace and cost structure that challenges incumbents to reimagine their operating models.



ENTREPRENEURSHIP

The



startup landscape is undergoing its most profound transformation since the internet revolution. This change is powered by a convergence of forces. Large language models (LLMs) have reached a level of maturity that makes them capable of reasoning and planning tasks. When given access to relevant data and permission to act on the user's behalf, LLMs can act as agents, taking on the role of digital assistant or even digital employee. New frameworks enable multiple agents to coordinate autonomously, creating agentic AI systems, or teams of digital colleagues. Low-cost application programming interfaces (APIs) have made it inexpensive to connect different IT systems so that they can share data and work together. And cloud costs continue to fall even as computing capabilities rise. These trends are producing a second great compression of entrepreneurship: The cost, time, and head count required to build a prototype of a product, test it, and then build an improved version have all collapsed. This compression shifts the competitive baseline. Small startups can enter, expand, or disrupt markets at a pace and cost structure that challenges incumbents to reimagine their operating models.

This article draws from our research across dozens of AI-native startups and interviews with founders, investors, and corporate executives. It maps the five forces driving this new entrepreneurial operating model and examines what they mean for both AI-native startups and the established companies they now compete against.

Before diving into the forces, we outline what makes agentic AI fundamentally different from prior generations of technology.

From Tools to Teammates: Understanding Agentic AI

The evolution from LLMs to agentic AI systems represents a move from chatbot tools to systems that can pursue complex outcomes autonomously—planning, acting, and adapting along the way.

Large language models such as ChatGPT and Claude serve as powerful tutors and assistants, helping users rapidly increase their knowledge of new industries, competitive landscapes, and regulatory frameworks, to give a few examples. These AI tutors can advise founders throughout the startup process by helping them build business plans and design pitch decks, providing guidance on securing capital, hiring, and more. Incumbents seeking to expand their offerings can access the same capabilities. Using recently developed AI research and reasoning tools, entrepreneurs and incumbents alike can compress months of learning into days or even hours.

AI agents go further than LLMs, performing more-complex tasks when given appropriate context and connected to the right tools. An AI financial-analyst agent, for instance, can develop five-year financial projections if it is supplied with inputs like the business plan and market research and is connected to spreadsheet tools to generate outputs.

Agentic AI systems are networks of specialized AI agents, coordinated by a digital manager, that can plan, reason, adapt, and pursue organizational outcomes. At NTT Data, a multinational IT services company, agentic AI systems in customer service do more than resolve individual service desk tickets; they synthesize patterns across thousands of conversations, identify systemic issues like software bugs, generate change requests, write and test code fixes, and create articles from that information for future reference. The system pursues the outcome of improved service quality, not just the resolution of individual queries. (NTT Data is a customer of Ema, an enterprise agentic AI platform that Vivian and Linda advise.)

Companies can build AI agents around outcomes without being constrained by organizational silos. For example, Ema has built an onboarding assistant that orders the ID badge for a new employee, assigns a desk, stages the laptop, activates payroll, enrolls the new hire in benefits, and even



With minimal resource commitments, startups can probe markets that are too uncertain for established firms to justify investment. If the fit is poor, they can pivot almost in real time.

schedules introductions with key teammates. It performs all the work that falls among HR, IT, and three different managers in most incumbent businesses.

The Five Forces of Disruptive Change

Five interconnected forces are creating a new entrepreneurial operating model. These forces strengthen one another, resulting in a compounding advantage over time. Together they are reshaping how a startup builds a prototype of a product or business model, systematically tests it and proves that there is real-world demand for it, and scales it up.

1 Zero-Latency Iteration

Zero-latency iteration refers to the feedback loop in developing a digital product. It means that each version of a digital product can be changed instantly upon receiving customer feedback, which is dramatically shrinking the time required to design, launch, and scale up a new product. Traditionally it took 12 to 18 months for a software-as-a-service (SaaS) startup to move from inception to establishing that a product satisfies a strong market demand (product-market fit). And that process came at a cost of \$1 million or more, according to our experiences. AI-native startups are significantly reducing those terms.

When Michael Rossiter and Neal Suidan, both Tesla alumni, founded Atomic—an AI-powered inventory-management platform for supply chains—they experienced this compressed timeline firsthand. Using simple “vibe-coding” tools (AI-assisted software development using natural language prompts), they prototyped a functional front-end user interface within days that was sufficient to secure paying customers.

This zero-latency iteration extends beyond initial prototypes. Souvik Sen, Ema’s cofounder and chief technology officer, described to us how AI agents accelerate software development. They not only optimize user interface (UI) components and automate wireframe generation (the process of creating a basic visual blueprint, or skeletal framework, of a website, app, or other digital product); they also critique the technical architecture, propose a cleaner

file-storage structure, and conduct real-time testing for bugs, security vulnerabilities, and performance issues. As a result, product managers and engineers can produce functional prototypes in hours rather than weeks.

Perhaps most striking is the blurring of traditional boundaries between roles. At Anterior, an AI-powered platform for health insurance prior authorization, user experience (UX) designers now change UI code without involving engineers. They use an AI coding assistant to modify the software and adapt interfaces directly. These efficiencies enable the company to develop, test, and release new software updates every two weeks.

With minimal resource commitments, startups can probe markets that are too small or uncertain for established firms to justify investment. If they find that the fit is poor, they can pivot almost in real time, dramatically reducing the risk involved in betting on the wrong market.

2 Automated Go-to-Market Capabilities

AI-native companies are employing capabilities that allow them to execute go-to-market strategies at a scale and speed that were once available only to well-funded enterprises with large marketing departments. These capabilities are slashing the costs of acquiring and onboarding new customers.

For example, some startups are using AI agents to generate and optimize marketing content, test thousands of image and copy combinations, identify optimal messaging, and dynamically shift marketing spend to the most effective channels. This data-driven approach allows a startup to execute marketing strategies that previously required teams of specialists. Two companies doing this are Atomic and Tactix, a decision-intelligence platform for owners of fast-food and fast-casual restaurants that makes instant recommendations on customer offers, staffing, and inventory management. (DVx Ventures, which Jon cofounded and is the CEO of, launched both companies.)

Agentic AI is also collapsing the hidden costs associated with growth. In traditional SaaS businesses, every new customer brings expensive human touchpoints. As Nikki Monterroso, a cofounder of Tactix, describes, once a customer

signs up for the platform, it needs to give vendors access to its data and workflows, defining the specific paths data must take, through tools and between people, to achieve key business goals. Rather than mapping each step manually, agentic systems can now handle much of these onboarding tasks automatically, in hours and days versus the many months it could take people to perform them.

Atomic delivers presale live demos using customer data within hours and completes full implementations in approximately one-quarter of the time that traditional SaaS implementations take. AI agents transcribe client conversations and configure the platform through natural language interactions; they can handle the complex technical translation that once required specialized consultants. It's an example of AI agents acting as forward-deployed engineers, taking on the responsibility of understanding customer needs well enough to integrate software solutions effectively.

3 Autonomous Business Functions

Agentic AI systems are making it possible to plan, execute, and optimize business activities end to end with minimal human involvement. This is dramatically reducing the head count required to build and scale a new business and constitutes one of the most radical differences between traditional startup models and those employing agentic AI systems.

A conventional technology product team includes two or three engineers, a product manager, a UX researcher, a UX designer, and a business lead—perhaps six to eight people to build a minimum viable product in six to 12 months. Today's AI-native startups require a minimum of just two people: a domain expert who understands the business problem and an AI engineer who can build agentic systems to address it.

At Tactix, the core product team consists of an expert from the quick-service restaurant industry and an AI engineer. By leveraging AI coding tools, that single engineer performs the work of 10. Within months, the company built an agentic AI system that ingests and validates restaurant data, aggregates it, analyzes patterns (including external factors such as weather and local events), and generates actionable recommendations for pricing, promotions, and operations. The AI agents responsible for each component are supervised by a “manager” agent.

This degree of automation also applies to a startup's own operations. At Ema, most HR processes and employee onboarding tasks are automated using AI agents. At Anterior, AI agents reduce legal expenses by approximately 90%



ENTREPRENEURSHIP

by reviewing contracts and highlighting key issues that need external counsel review. Abdel Mahmoud, Anterior's founder and CEO (who is a physician with a computer science background), was thus able to allocate far fewer people than typical to administrative work, leaving 30 of his 40 employees to build the product.

Agentic systems don't just automate workflows; they unlock the collective intelligence buried in an organization's data and history. Mahmoud has launched an “AI librarian,” a digital curator that organizes all of Anterior's documents, emails, meeting minutes, and other internal communications for easy access anytime, anywhere. Mahmoud regularly uses AI to survey the state of his company, asking, for example, “What were the top engineering challenges this week?” He receives accurate, referenced summaries within seconds. It's a small example of a much larger shift. Knowledge that was once fragmented across individual files, and which decayed with staff turnover, can now be continuously indexed, learned, and reused by the organization.

4 Radical Capital Efficiency

Thanks to a combination of lean operations, rapid iterations during product development, and autonomous AI agents, AI-native startups need much less capital to launch a new business and become profitable and can raise it faster than conventional startups.

LangChain, an AI-application-development framework, condensed the typical fundraising journey of more than two years to 10 months for its \$25 million Series A round. Delve, an AI-native compliance automation platform that aims to help startups and enterprises achieve security certifications, raised \$32 million in a Series A round just six months after a seed round.

DVx Ventures has launched 12 startups in four years. Its AI-native companies have used \$2 million in capital by the time they have reached the Series A milestone, 80% less than the amount consumed by its previous startups that did not use AI. The timeline for reaching the Series A milestone is 20% to 40% shorter as well. With less capital required per company, venture firms can make more bets, potentially improving portfolio returns.



5 The AI-Driven Flywheel


By solving the “unsolvable” manual tasks at the center of a company’s enterprise, agentic AI startups gain an early and deep understanding of their customers’ core priorities and continuously improve how they address them. That is one of their main competitive advantages. Indeed, customized AI interfaces are no longer scarce, and off-the-shelf agent-building platforms are increasingly available. These startups’ real strength lies in the deep operational and workflow expertise gained from companies deploying their customized AI agents to solve some of their customers’ most intractable problems in their most economically critical business units.

For example, Anterior gained early traction in the risk-averse health insurance sector by using AI agents to build what Mahmoud called “the best medical-record ingestion engine.” The system reliably processes faxed PDFs, the “dirty work” that many humans and AI systems struggle with, extracting structured data from unstructured documents.

During prior-authorization reviews, Anterior’s agents structured rich clinical data—often 600 pages of

unstructured patient information per case—which gave health insurers access to invaluable and accurate insights quickly. With this structured data, the health plans can improve care management, risk adjustment, and payment integrity. Then companies like Anterior support them as they build better processes, including handing them the reins to build their own agents. As Mahmoud explains, “Nurses fresh from the hospital floor can, within a few weeks, be designing their own clinical AI workflows” using his company’s platform. Customer retention strengthens over time. That’s because the company’s deep workflow integration creates high switching costs, making it increasingly difficult for users to leave, which acts as a barrier to new competitors.

One of the most powerful attributes of an agentic AI system is its self-reinforcing loop, which enables it to continuously improve a business function by learning from its own actions and outcomes. That is the flywheel effect. More usage generates more insights, which improves customer engagement and agent performance, which drives more usage. Like SaaS businesses that became more valuable as they accumulated customer data, AI-native startups



building proprietary workflow knowledge through their AI agents create compounding advantages.

At the platform level, companies like Ema are also accumulating a deep understanding of operations and change management as they deploy this new generation of SaaS products that contain AI agents. For example, Ema helps leaders navigate the stress of AI adoption by establishing clear usage rules and recommending guardrails to ensure a seamless rollout. It specializes in turning common messy and manual tasks—like customer service and onboarding new employees—into organized, automated processes. As AI deployments grow in scale and intricacy, having fluency in complex operations becomes a significant competitive advantage.

Incumbents Must Rethink Their Architecture

The rise of capital-efficient, AI-native competitors exposes a structural mismatch between how incumbents were built and how agentic systems operate. Startups form around clean workflows, tight learning loops, and data models designed for systems that learn continuously. Most

incumbents, by contrast, evolved for stability, specialization, and control, and they resist efforts at experimentation and innovation.

The obstacle for established companies isn't a shortage of AI tools; it's the organizational architecture those tools land on. Employees who have spent a decade mastering stable processes have little preparation for AI-augmented roles. Middle managers, whose authority historically comes from controlling information and resources, tend to balk at anything that threatens their span of control. None of this reflects a lack of capability; it reflects how the roles were designed. As Clayton Christensen and his coauthor warned in the first HBR article on disruptive innovation, the very processes that serve core markets "work so well that they blind those [incumbent] companies" to new ways of operating. (See "Disruptive Technologies: Catching the Wave," HBR, January–February 1995.)

Technical debt—the additional rework, time, and resources necessary to address suboptimal software designs—magnifies the challenge. Enterprise workflows are rarely documented cleanly. Data lives in silos. Integrations are brittle. Exceptions require manual intervention. Agentic AI depends on clear



workflows, unified data, and unbroken feedback loops, and achieving those things can be challenging for incumbents. Their most common mistake is trying to automate before they re-architect their workflows. In the words of Michael Hammer, “It’s time to stop paving the cow paths. Instead of embedding outdated processes in silicon and software, we should obliterate them and start over.” (See “Reengineering Work: Don’t Automate, Obliterate,” HBR, July–August 1990.)

The economics intensify the pressure. AI-native companies can reach early milestones with a fraction of the capital and head count previously required. When a competitor can achieve five times more output with the same resources or reach traction with one-fifth the capital, the cost curve of an entire sector shifts. This dynamic creates pressure on incumbents either to cut pricing deeply (risking operational stability) or to hold pricing steady (risking gradual erosion of competitiveness).

Yet AI-native startups face vulnerabilities that incumbents can exploit. Autonomous systems introduce quality, reliability, and compliance risk: A misinterpreted signal, an incorrect approval, or a faulty configuration can trigger real-world consequences. Demonstrating control across multi-agent systems is nontrivial, especially when the systems rely on AI models that are inherently unpredictable. Unlike traditional software, AI agents don’t always produce the same result twice, making their behavior harder to standardize and verify. Startups must build testing, monitoring, and escalation mechanisms that incumbents have spent decades refining. These include audit trails, compliance reporting, an incident-response culture, and, in the event of errors, systems like incident commanders, root cause analyses, and postmortems. AI-native startups may discover they need these before customers will trust them. Incumbents can leverage their established reputation for operational excellence and regulatory compliance to frame AI-native solutions as inherently risky and unreliable.

Talent is another constraint. While agentic AI can handle simple tasks and reduce the overall number of workers needed, it dramatically increases the need for top-tier employees with deep, specialized industry knowledge. Those are precisely the individuals that are already employed by large established companies. As LLMs become commodities, the real differentiating assets are proprietary data, the ability

to reengineer internal workflows around AI, and the organizational expertise needed to blend AI speed with human judgment. Those are all sources of competitive strength that incumbent companies can unlock if they commit to changing fundamentally how their business operates.

Incumbents must ask, “How do we redesign the work so that people and agents can learn together?” The winning architecture is a system where data, processes, and the rules for decision-making are specifically designed to improve over time.

What Incumbents Should Do Now

Skip the pilots that bolt agents onto broken workflows. Instead start where friction in workflows is highest and the stakes are greatest.

Run a vulnerability audit against the five disruptive forces. Take a close look at a few of your most important workflows and ask where they slow down, rely on too many handoffs, or force people to make decisions without good information. These can be high-volume intake and routing processes such as insurance claims, medical referrals, or loan operations, which are often handled by junior staff and where initial errors can cause significant downstream issues. Also, in workflows like HR onboarding or supply chain coordination, manual data transfer between multiple systems often leads to a loss of critical context. These weak spots are exactly where an AI-native competitor would attack first, because fixing even one of them can reset customer expectations overnight. The goal is to see your business the way a fast-moving rival would. Identify the few workflows that truly define your advantage and the few that could quickly become liabilities.

Start with a bounded, high-friction process. Pick one workflow that matters and improve it first. It should be meaningful enough that progress is noticeable but contained enough that it’s manageable. You’re looking for the sweet spot where you will learn the most and where improvements will strengthen your competitive position. Avoid picking an easy pilot. Choose something that, if fixed, clearly raises the bar for the whole organization and resets expectations for customers or teams. At NTT Data, for example, Edoardo Tealdi’s AI transformation team started with the company’s task of



For two decades, incumbents competed with startups from a position of resource superiority. Agentic-AI-based ventures have dramatically eroded that advantage.

writing proposals to respond to a potential customer's RFP, a process that would normally take hundreds of personnel hours. By drawing on the repository of all the company's services, client relationship information, and market and competitive insights, NTT's proposal-writing AI agent creates the first draft in 20 minutes; it has increased both the company's responses to RFPs and its success rate.

Re-architect before you automate. Don't drop an intelligent system into a messy process and expect magic. If a workflow is outdated, confusing, or full of unnecessary steps, an agent will simply reproduce those flaws faster. Redesign the process first: Remove steps that no longer matter, clarify who makes each decision, and simplify the flow. (See "Designing a Successful Agentic AI System," HBR.org, October 24, 2025.) Once the structure is clean, automation will amplify the gains instead of locking in old problems.

Partner with AI-native ventures. Working with an AI-focused company is a practical way to see how a modern workflow can run when it isn't weighed down by legacy systems. You keep control of your customer relationships and decisions; your partner shows you what a cleaner, faster process looks like in practice. A collaboration will reveal internal gaps that you might not have noticed on your own and help speed up your redesign efforts. When the U.S. Army's 18th Airborne Corps developed the Maven Smart System—a computer vision AI platform that accesses sensor data from diverse sources, helps soldiers identify and choose military targets, and then provides workflow support for approval by the chain of command—success hinged on collaboration with industry field engineers. According to Emilia Probasco, a senior fellow at Georgetown University's Center for Security and Emerging Technology, by working directly with soldiers, the army's agentic AI partners designed intuitive interfaces and workflows that integrated seamlessly into the 18th Airborne's existing military-intelligence processes.

Strengthen data quality and make it clear when tasks should go back to humans. Agents don't need perfect data, but they do need clear rules and reliable signals—accurate, consistent, and trustworthy data points or patterns that genuinely reflect reality. Make sure the information behind everyday decisions is accurate, that exceptions are visible rather than hidden in emails, and that it's obvious when a

task should be given to a human. Having a human in the loop is a vital component of the governance of agentic AI. When boundaries and handoff points are explicit and guardrails are in place, agents move smoothly through the work. When they aren't, the system gets stuck, repeats old mistakes, or, worse yet, generates new ones.


Prepare people for shifting roles. Humans will focus more on judgment, edge cases (rare, extreme, or boundary situations in software or systems that fall outside normal usage but can cause unexpected behavior, errors, or crashes), and empathy, while agents will handle setup, coordination, and monitoring. New responsibilities will emerge, such as overseeing agents, interpreting exceptions, and improving the workflow over time. For example, when Tampa General Hospital's call center deployed its digital agent, Aimee, "she" was instructed to handle routine inquiries and escalate complex cases to human staff. The call center's operations team maintained full control over Aimee's knowledge base and the specific logic governing her responses. When teams adapt alongside the technology, both keep getting better. Expect the definitions of teams that include humans and agents to evolve beyond traditional roles and responsibilities.

AGENTIC AI IS unleashing a fundamental restructuring of competitive dynamics across industries. For entrepreneurs, the barriers to entry have never been lower. For incumbents' executives, the choice isn't whether to embrace agentic AI; it's how quickly and comprehensively to do so safely. For two decades, incumbents competed with startups from a position of resource superiority. Agentic AI-based ventures have dramatically eroded that advantage. History suggests that waiting until the threat becomes undeniable is waiting too long. ■

HBR Reprint R2604G



VIVIAN S. LEE is an executive fellow at Harvard Business School and the former president of Verily Health Platforms at Alphabet. She is the author of *The Long Fix: Solving America's Health Care Crisis with Strategies that Work for Everyone* (W.W. Norton, 2020). **LINDA MANTIA** is a director on the boards of both AI-native and incumbent companies, including Alan, Dayforce, Liberty Mutual, and Maple Leaf Foods. She is the former COO of Manulife. **JON MCNEILL** is a cofounder and the CEO of DVx Ventures and was previously the president of Tesla. He is the author of *The Algorithm* (Penguin Random House, 2026).





AUTHORS

**Michael
Mankins**
*Partner, Bain
& Company*

**Matthew
Crupi**
*Partner, Bain
& Company*



BUSINESS &
SOCIETY

Bring Back Managing for Value

As capital costs rise, companies need a disciplined way to allocate resources without abandoning commitments to employees, customers, and the environment. Here's how to find that balance.



ILLUSTRATOR CHAD HAGAN



Over the past two decades managing for shareholder value has fallen out of favor. It has been portrayed as out of step with societal priorities—particularly the growing emphasis on environmental, social, and governance (ESG) concerns—or incompatible with the interests of other stakeholders.

In its place, many firms have embraced stakeholder-oriented or purpose-driven models, seeking to balance the interests of employees, customers, communities, the environment, and shareholders simultaneously.

That shift coincided with an unusually forgiving capital-market environment. In the decade following the global financial crisis, capital costs fell to historically low levels, muting the consequences of weak capital allocation and obscuring the trade-offs inherent in strategic choices. Growth masked weak returns; subpar performance was tolerated; and the discipline of rigorously testing potential moves against their ability to create value eroded.

The results were uneven at best. Starbucks expanded rapidly, boosting revenues but diluting returns on invested capital before closing underperforming stores and refocusing on profitable growth. The French food products giant Danone leaned heavily into a purpose-driven governance model, but after it suffered sustained underperformance relative to its peers, shareholder pressure forced a leadership change.

In both cases the central issue was not ambition or stakeholder concern—it was the absence of clear return thresholds guiding capital allocation. Now, as capital costs move back toward historical norms (see the exhibit “The Return of Normal Capital Costs”), companies are less willing to look at options that do not create value. Trade-offs that once seemed manageable are becoming untenable, making the lack of a coherent governing objective increasingly costly.

This article argues that it’s time to reimagine value-based management for a more constrained world. Properly applied, value maximization remains the only governing objective capable of guiding consistent, accountable decision-making in large, complex organizations. Managers must relearn what earlier generations understood about managing for value—and rebuild the organizational capabilities that atrophied during an extended period of exceptionally low capital costs.

A Return to Value

If value maximization is to guide decision-making once again, it must be applied differently. Leaders must incorporate what recent experience has made clear about the real and enduring impact corporations have on multiple stakeholders. We propose a simple but disciplined model, one rooted in linear programming: Retain value maximization as the governing objective—the maximand—but explicitly recognize stakeholder commitments as specific constraints that limit the set of feasible choices, not as objectives in and of themselves.

The intellectual foundations for this approach have existed for decades. Michael Jensen and other scholars argued, correctly, that value maximization was the only coherent governing objective for large organizations, with legal, ethical, and contractual obligations appropriately constraining managerial action. But those arguments largely operated at the level of principle. They explained what firms should aim to do in theory, not how leaders can address constraints in concrete strategic and capital-allocation decisions.

It’s critical that leaders make those constraints explicit, distinguishing between hard and soft constraints.

Hard constraints. Hard constraints are inviolable. Any alternative that fails to satisfy them is eliminated from



Managers must relearn what earlier generations understood about managing for value—and rebuild capabilities that atrophied during a period of exceptionally low capital costs.

consideration. Compliance with the law is the most obvious example: Every strategic option must be legal. Leadership may impose additional hard constraints to reflect binding commitments or nonnegotiable priorities. Examples include “no increase in the firm’s carbon footprint,” “no reduction in employee satisfaction,” or “no decline in customer Net Promoter Scores.” To be effective, hard constraints must be:

→ **Measurable and monitorable.** Hard constraints must be defined in objective, observable terms. Scholars and practitioners have developed rigorous tools to support this discipline—such as the Net Promoter Score framework created by Fred Reichheld and colleagues at Bain, Robert Kaplan and Karthik Ramanna’s proposed e-ledger accounting measures for carbon, and standardized approaches to measuring employee engagement from organizations such as Gallup and Glassdoor. Without credible measurement, a hard constraint is aspirational rather than binding.

→ **Specified as absolutes.** Hard constraints screen out alternatives. They must therefore be framed in a way that allows leaders to determine clearly whether a proposed action violates the constraint. If the constraint is ambiguous, it invites reinterpretation—and ultimately discretion—precisely what the framework is designed to avoid.

Climate commitments illustrate both the challenge and the promise of defining hard constraints. Many companies—including Amazon, Procter & Gamble, and HP—have adopted net-zero pledges for 2040. While laudable, long-term aspirations can be difficult to treat as binding constraints in near-term decision-making. It is often unclear whether a specific investment meaningfully advances—or undermines—a goal set 15 years in the future, leaving room for competing narratives.

By contrast, some companies have translated environmental commitments into measurable, time-bound targets.

Unilever, for example, has established science-based emissions-reduction targets for Scope 1, 2, and 3 emissions validated under the Science Based Targets Initiative. Unilever’s targets are tied to annual reporting and embedded in its climate transition action plan, enabling year-over-year accountability. While the company’s ultimate goals extend to 2030 and beyond, progress is tracked and reported annually—making the constraint more actionable.

Hard constraints need not be limited to environmental commitments. At Alcoa the former CEO Paul O’Neill made worker safety the company’s top priority and tied managers’ evaluations to measurable injury metrics, including lost-workday incident rates. Alternatives that compromised safety performance were unacceptable—even if they promised near-term productivity gains. Similarly, Toyota’s production system empowers frontline employees to stop the line when quality defects are detected, treating customer safety and product integrity as inviolable constraints on throughput and cost. At Lincoln Electric, a long-standing guaranteed employment policy has functioned as a hard constraint during downturns, forcing leaders to adjust compensation structures and operating practices rather than resort to layoffs. In each case the constraint was explicit, measurable, and enforced in practice—screening out a number of alternative solutions to a given problem or challenge.

Soft constraints. Soft constraints apply only after all hard constraints have been satisfied. They are most useful for differentiating among alternatives with similar—or roughly equivalent—intrinsic value. Examples include “minimizing disruption to local communities,” “increasing workforce diversity,” and “limiting near-term capital expenditures.” To serve their purpose, soft constraints should be:

→ **Measurable but not absolute.** Like hard constraints, soft constraints should be defined in objective terms. Unlike



IDEA IN BRIEF

THE PROBLEM

As capital becomes more expensive, many companies are finding that they lack a disciplined way to allocate resources.

THE REASON

Years of cheap capital allowed firms to pursue growth and stakeholder goals without rigorously testing whether investments created value.

THE SOLUTION

Reestablish value maximization as the governing objective but explicitly recognize stakeholder commitments as constraints that limit the set of feasible choices. Leaders can use this model to evaluate alternatives consistently, allocate capital to its highest-value use, and address legitimate stakeholder claims without sacrificing accountability.



hard constraints, however, they do not eliminate alternatives; they help differentiate among them. When two strategies create comparable value, soft constraints guide leaders toward the alternative that best aligns with secondary priorities. Consider Costco. The company explicitly affirms the primacy of shareholder value while emphasizing the importance of treating employees, customers, suppliers, and communities well. In practice, Costco has repeatedly chosen to pay higher wages and provide greater benefits even when competitors reduced costs—arguing that these decisions support long-term shareholder value. Critically, Costco does not subordinate shareholders; stakeholder preferences influence decisions only when financial outcomes are not materially worse.

→ **Prioritized and weighted.** Soft constraints cannot all carry equal importance, or leaders will lack a principled way to sort through competing claims. When two alternatives create similar value but differ in their impact on multiple soft constraints, leaders need a clear method for determining relative importance and arriving at an overall assessment. This situation commonly arises when decisions affect employees, suppliers, and customers differently. One option may impose short-term costs on employees while substantially benefiting customers or suppliers. Without clarity on the relative weight assigned to each group, leadership inevitably reverts to discretion. Prioritization and weighting reduce ambiguity and preserve discipline.

Both types of constraints must be agreed upon in advance. Waiting until after alternatives have been evaluated invites two failures: allowing hard constraints to quietly replace the objective function, and retroactively elevating soft constraints to veto power. Either outcome substitutes discretion for discipline—the very problem this model is designed to prevent.

Value-Based Decision-Making

Managing for value requires more than a stated commitment to shareholder returns. Leaders must be able to estimate—and compare—the intrinsic value of competing strategic alternatives. Twenty years ago many large companies had the financial information and processes necessary to support value-based decision-making. Today far fewer do. To

manage for shareholder value once again, companies must build—or rebuild—four foundational capabilities.


Alternatives-based decision-making. Most companies do not routinely formulate and evaluate alternatives when making strategic choices. In a Bain survey of executives from 760 large organizations, only one in five reported explicitly considering alternatives in strategy development. A common pattern is incremental substitution: “We are doing this; that looks better; let’s do that.” While that approach can improve performance, it does not ensure that capital and scarce resources are put to their highest-value use. Value maximization demands that leaders evaluate explicit alternatives against a common objective—subject to prespecified constraints.

Balance-sheet reporting at the business-unit level (and lower). A strategy’s intrinsic value reflects the discounted value of future equity cash flows—or, equivalently, the present value of future economic profits plus initial book equity. Both formulations require an explicit understanding of the balance sheet. Economic profits, by definition, equal earnings minus a charge for the equity capital employed; without accurate measures of invested capital at the business-unit or product level, leaders cannot determine whether returns exceed the cost of equity—or by how much. However, information on the capital requirements of businesses, products, and services is unavailable in many organizations today. We estimate that in 2005, prior to the global financial crisis, nearly one-third of the S&P 500 had the capability to gather it. But today practitioner and academic research suggests that relatively few large companies systematically allocate balance-sheet capital to operating units for ROIC or economic profit measurement. Instead, many rely on top-level metrics like EBITDA or the contribution margin and managerial judgment. But distributed capital measures are essential for rigorous value-based decision-making. A return to value-based decision-making will require most organizations to improve their management reporting.

Reliable estimates for capital costs. Most companies maintain estimates of their cost of equity, usually developed by treasury or finance teams. These estimates are typically sufficient for comparing strategic alternatives—provided they reflect current risk-free rates and equity risk premiums. Bain analysis suggests that the market rate of return—the average



Managing for value requires more than a stated commitment to shareholder returns. Leaders must be able to compare the intrinsic value of competing strategic alternatives.





cost of equity across large public companies—is slightly above 9% today, implying a market risk premium of roughly 4% to 4.5% over long-term government bond yields. For most large public companies, the cost of equity falls within a band of plus or minus 1.5 percentage points of that market rate. For smaller private companies, the appropriate cost of equity depends more directly on the risk tolerance and diversification of the owners. Because owners often have a larger share of their personal wealth tied to the business, they may require a higher return to compensate for that concentration of risk—implying a higher effective equity risk premium.

Disciplined strategic decision-making does not require calculating a precise absolute value for all alternatives. The goal is to understand the relative value of each. As long as the cost of equity is specified within a reasonable range, modest differences in the discount rate are unlikely to change the ranking of alternatives. Leaders should therefore resist assigning different discount rates to different strategies; it is generally more reliable to reflect differences in risk by estimating the variability in projected cash flows rather than by adjusting the discount rate itself.

A disciplined approach to financial forecasting.

Companies need their projections of future cash flows or economic profits to be realistic and grounded. While forecasting is inherently uncertain, leaders can impose discipline by using consistent assumptions, considering explicit and well-defined scenarios, and applying projections that reflect both industry attractiveness and the firm's competitive position under each alternative. The goal is comparability: using a common analytical structure that ties economic profit forecasts directly to the market economics and competitive position implied by each strategic option.

Together, these capabilities allow leaders to compare discounted cash flows across competing strategies. Some firms have retained these disciplines. But many will need to rebuild them.


The New Model in Practice

Strategy is ultimately a sequence of choices. Making good choices requires generating strong alternatives and then applying a rigorous selection process—one that combines a clear objective with explicit constraints.

The Return of Normal Capital Costs

For more than a decade following the global financial crisis, interest rates hovered near zero, risk premiums were low, the use of financial leverage increased, and the weighted average cost of capital (WACC) for most large U.S. firms fell to levels rarely seen in modern financial history. This period was an anomaly. Long-term data suggests that the cost of equity for U.S. corporations has typically ranged from 9% to 12%, while nominal WACC has been in mid-to-high single digits. Since 2022 capital markets have been reverting toward historical patterns.

U.S. market average capital costs across all industries



Note: Data is for companies in the Value Line Investment Survey.

Source: Bain & Company

A large aerospace manufacturer confronted that challenge when demand for commercial aircraft rebounded sharply once the pandemic-induced collapse in air travel ended. The company needed to ramp up production quickly, but many suppliers were unprepared to meet the sudden increase in demand. Senior leaders had to determine the best system-level response, accounting for both supplier readiness and internal production capacity. The leadership team began by specifying the governing objective. The team agreed that it should be value maximization. No alternative objective better balanced profitability and growth or near-term performance and long-term results.

Leaders then specified several hard constraints. First, the company had to meet the required production increase without introducing unacceptable schedule risk. Second, it had entered into binding agreements with state authorities to maintain employment levels in two localities; no alternative could violate those commitments. Third, to improve



The challenge is not merely to return to managing for value but to do so with a disciplined framework that preserves explicit stakeholder constraints.

supply-chain resilience, leadership imposed a constraint to eliminate sole-sourcing of critical components. Finally, leaders specified that the systemwide carbon footprint of manufacturing and assembly could not increase.

After eliminating alternatives that violated hard constraints, leaders turned to soft ones. When two or more options produced similar intrinsic value, they favored those that minimized disruption to existing operations and limited near-term cash outlays—reflecting concerns about execution risk and the company's still-fragile liquidity position.

Applying this framework ruled out several options that initially appeared attractive. One seemingly promising alternative involved relocating critical activities to lower-cost locations. Despite its cost advantages, the option failed on multiple dimensions: It introduced unacceptable schedule risk, required lengthy transitions for critical work steps, and violated the company's employment commitments. Another option would have marginally reduced ongoing labor costs by eliminating overtime, but it required substantial new capital investment—destroying value once capital costs were taken into account.

Ultimately, leadership chose to adopt new production technologies—already proven at one of the company's international sites—and to adjust its make-versus-buy mix. The company outsourced select components where internal bottlenecks would have constrained the production increases, but it retained core activities that were critical to quality and coordination. This alternative enabled the firm to meet its production targets over the subsequent 24 months without violating any of its hard constraints.

The Pendulum Is Already Swinging Back

Many companies are already recalibrating their strategies toward value creation. According to a 2025 survey by the Conference Board, 80% of large companies reported adjusting their ESG approaches in response to policy and market shifts, with many moving away from broad ESG positioning and toward initiatives that demonstrate clear business ROI and shareholder value impact.

BP offers an example of that shift. After pursuing an ambitious renewable-energy transition earlier in the decade, the oil giant announced a “fundamental reset” of strategy

in 2025, reallocating capital toward higher-return oil and gas investments and scaling back planned spending on certain low-carbon initiatives. BP's leadership framed the pivot explicitly in terms of improving cash flow, strengthening returns, and focusing capital on projects with stronger return profiles. The company's recalibration reflects a broader reality: As capital costs rise and trade-offs reemerge, firms are reasserting return thresholds in capital-allocation decisions. That shift is neither surprising nor inappropriate. The danger lies elsewhere.

The risk is that, in the process of restoring value maximization, companies swing the pendulum too far—treating shareholder returns as the sole objective rather than a governing one that is subject to explicit constraints. Over the past two decades executives have internalized the fact that corporate actions have a real and lasting impact on employees, communities, and the environment. Abandoning those lessons in a rush back to financial rigor would be as misguided as it was to abandon value-based management during the era of cheap capital. The challenge, therefore, is not merely to return to managing for value—but to do so within a disciplined framework that preserves explicit stakeholder constraints.

AS THE ERA of ultralow capital costs draws to a close, business leaders can no longer afford to rely on ad hoc stakeholder balancing to justify the allocation of scarce resources. Managing for shareholder value—properly understood and rigorously applied—provides clear rules. By pairing a single governing objective with explicit hard and soft constraints, leaders can evaluate alternatives consistently, allocate capital to its highest-value use, and address legitimate stakeholder claims without sacrificing accountability. In an environment where mistakes compound quickly and reversals are costly, this discipline is not ideological—it is essential. ■

HBR Reprint R2604H



MICHAEL MANKINS is a leader in Bain & Company's Organization and Strategy & Transformation practices and is a partner in Austin. He is a coauthor of *Time, Talent, Energy: Overcome Organizational Drag & Unleash Your Team's Productive Power* (Harvard Business Review Press, 2017). **MATTHEW CRUPI** leads Bain & Company's Corporate Strategy & Finance practice and is a partner in Dallas.



“ The clarity and relevance of the content to real executive decisions made the frameworks and insights useful. They helped sharpen thinking, structure conversations, and drove more-focused outcomes. ”

HBR Executive Subscriber



Lead with
insight.
Act with
impact.



HBR Executive empowers teams of any size to navigate uncertainty, stay ahead of emerging trends, and plan for a stable future.

Leadership teams succeed together.

Use HBR Executive for:

- AI activation
- Upskilling
- Peer networking
- Strategic alignment


HBR Executive is an annual subscription that includes content and conversation from the experts at Harvard Business Review.

Subscribe to HBR Executive today.
hbr.org/executive-teams



Experience

Advice & Inspiration



MANAGING YOURSELF

How Elite Sports Coaches Make High-Pressure Decisions

What they do before, during, and after critical choices.

by Alan McCall, Adrian Wolfberg, Johann Bilsborough, and Ricard Pruna

BUSINESS LEADERS ROUTINELY make important decisions under pressure, often with incomplete or conflicting information, in ways that significantly impact team and organizational performance—as well as their own careers. Elite sports coaches do the same, with two critical factors increasing the stakes: a compressed time frame (they often must call plays in just seconds) and constant public exposure (including live TV coverage and 24/7 criticism from fans and the media).



experience

Over the past several years we have studied 11 successful coaches working in the National Football League (NFL), National Basketball Association (NBA), and Major League Baseball (MLB) in the United States; Premier League, LaLiga, and UEFA Champions League (football) and Rugby Union in Europe; and the National Rugby League in Australia and New Zealand to better understand how they handle everything from in-game play calls and substitutions to recruitment and return-from-injury choices. We wanted to understand what they did cognitively, emotionally, and socially as they made decisions.


What emerged was not a checklist of traits or a new decision formula. Instead, we were able to carefully document how high-stakes decisions take shape before, during, and after moments of consequence. In this article we focus on specific practices the coaches employ during each of those three phases—and explain how business leaders can use them to increase the quality of their own decision-making.



BEFORE

Before a decisive moment arrives, coaches focus on creating the conditions that make clarity possible when pressure hits.

They anticipate future decision needs. The coaches we interviewed look ahead to the key in-season or in-game calls they'll need to make and rehearse various scenarios. One Rugby World Cup-winning head coach we interviewed explained, "You have to anticipate situations that require a decision [and] find the solution in advance."



Another elite rugby coach said, "It's all about planning so you can be decisive in the moment [and] not have a panic conversation on the spot."

This planning is practical rather than theoretical. Coaches mull over decision paths themselves but also test their thoughts with others by talking things through, inviting challenges, and sharpening calls *before* the moment arrives: If this happens, what do we do? Who decides? What evidence will matter? The goal is not to script the future but to remove ambiguity about roles, criteria, and authority. An MLB manager described how that played out in a recent World Series game. The night before the first game, he and his staff reviewed their batting lineup against the opponent's All-Star left-handed starting pitcher. They decided to start a promising left-handed rookie instead of a right-handed veteran, who would typically be favored in that scenario, because they thought the rookie would better match up against this particular ace. They also decided that if the other team brought in a left-handed reliever, they would replace the rookie with the veteran. When the

opposing team swapped in one lefty for another in the seventh inning, it might have been tempting to change the plan because the rookie had been hitting so well. But the manager didn't deviate. He subbed in the veteran, who later hit a home run and sealed the team's victory. The coach later reflected, "We had made the decision the night before."

They control how information reaches them. Anticipating calls is not enough. Coaches need data, and the best ones are deliberate about how data is given to them. "It needs to be purposeful and presented in a way that gives meaning to what you want to do," a world-championship-winning rugby coach explained. "Never more than three pieces of information." An NBA coach said, "Too much information from too many people makes [for] a cloudy decision. Just give me focused and valuable data."

Today analytics guide many decisions in sports. Baseball's use of data to drive talent decisions was made famous by the book *Moneyball*, and now teams use detailed models to simulate thousands of game scenarios. The



Coaches test their thoughts with others by talking things through, inviting challenges, and sharpening calls *before* the moment arrives.

NFL, NBA, and European football and rugby leagues also use GPS tracking to optimize player conditioning, video-tagging software to help with opposition scouting, and AI-driven predictive modeling to inform in-game decisions. Business leaders similarly have no shortage of data, as data platforms and AI tools increasingly offer decision recommendations.

Yet the coaches we studied did not think that analytics would replace their judgment. A highly successful rugby coach described how he gets the most valuable data. “I get my assistant coaches to identify the key data—not just analysts, because they don’t always see the context,” he explained. “Then I look at the people behind the data. Does the information match?”

They understand their people.

The third element of preparation is subtler but no less important. Coaches use moments when nothing is at stake to learn how their staff members and players behave. They constantly observe their team members in everyday settings, watching how they communicate and carry themselves, when they hesitate or speak up, and how they respond to challenges and uncertainty. An NBA coach we interviewed made a point of walking through offices, locker rooms, and shared spaces, noticing individual differences and changes. “If someone’s behavior shifts, I know something is off,” he explained. “That helps me make decisions later under stressful conditions.” In a recent game, when drawing up a last-second play, he factored in his earlier observations when deciding to change which star player should take the final shot.

Great business leaders rely on the same principles to set themselves up for better decision-making. While it’s impossible to predict the exact particulars of every scenario your team or organization will face, you can thoughtfully plan for different ones, from the most probable to the least. You can ensure that your team gives you relevant data and opinions rather than overloading you with unnecessary detail. And you can learn how to read your team members before a crunch time hits so that you can more effectively draw on them for support when it does.

Consider how the chief risk officer of a global investment company prepared for a looming decision on central-bank interest rates, knowing that even a modest policy surprise could force rapid portfolio adjustments. To control information flow, she insisted that every update come to her in a strict format, which included exposure level, downside scenario, and immediate action trigger. She then asked her team members to weigh two paths: to position the firm defensively or to hold current positions and be ready for rapid adjustment. She asked portfolio managers to recommend moves, paying close attention to their level of conviction, but communicated that the final call would sit with her. When the central bank moved early and markets began to swing, that preparation paid off. With disciplined inputs and advice from her managers, she could make decisions calmly and quickly.



DURING

When facing high-pressure choices, top coaches steady

their own emotional state, read people in real time, and draw on their preparation as momentum, confidence, and pressure shift around them.

They manage emotions. Their goal is not to eliminate emotion but to regulate it so that it does not distort their judgment. Positive feelings like excitement can have a biasing effect. Negative ones—such as anger and fear of undesirable outcomes—are more pernicious. One European football coach highlighted the danger of imagined reactions: “The biggest thing that can blur the picture is thinking about how the media will react. If you start adding ‘ifs,’ fear takes over. Fear doesn’t align with good decision-making.”

The coaches we talked to don’t suppress their feelings; they shrink them and bring their focus back to the moment: What matters right now and only right now?

They apply the same discipline to others’ emotions. “The feeling [that] coaches have in the box is not always matching what’s happening on the pitch,” one elite rugby coach told us. “In the moment, I need to influence what’s required”—that is, to not be affected by the players’ feelings. This isn’t cold detachment. It is an effort to prevent the wrong signals from driving the decision.

They read the room. Coaches pay close attention to people across multiple channels in real time—what they see, hear, and sense in the moment. That includes visual cues such as body language and movement, verbal signals such as tone and choice of language, and more-subtle shifts in energy. When people are not quite acting like



When there is debate, conflict, or resentment in the aftermath of a decision, the best managers circle back with all stakeholders and identify how to do better next time.

themselves, they recognize that something is off. This focus on deviation rather than absolute behavior allows coaches to identify both risk and readiness more accurately. A player who appears calm might not be energized enough for the moment; a player who is usually intense but suddenly withdrawn may not be reliable.

These impressions are not taken at face value. Coaches repeatedly described checking and rechecking what they see, often validating their impressions with others before committing. “I read body language. Is he nervous, flustered? If I start to see something, I ask the other coaches if they’re seeing the same thing. I’ll read it three or four times, and then I know I need to make a decision.” The final choice reflects not just tactics or data but a real understanding of people in the moment.

They turn preparation into instinct.

Decisions that appear to be gut reactions from the outside are rarely spontaneous on the inside. Instead, they reflect rapid pattern recognition built through experience and repeated exposure. Many likened this process to clinical reasoning: scanning the situation, matching it to known patterns, and narrowing options almost instantaneously without consciously working through every alternative.

“The decision is based on the circumstances and the resources you have, which is really your experience,” one rugby coach told us. “It becomes freeing. You can make logical choices on the fly.” Several coaches described this as knowing before being able to fully articulate why. The recognition comes

first; the explanation follows. In this way, preparation does more than inform decisions—it reshapes how situations are perceived. The right option becomes more obvious, not because the decision is simple but because the complexity has already been absorbed.

How do these principles work in more-traditional organizational settings? Successful leaders don’t ignore their own feelings or those of their team members, but they find ways to contain them so that decision quality isn’t adversely affected. They scan for signals from trusted team members and outside stakeholders. And then they trust their preparation and pattern recognition enough to follow their instincts rather than their emotions.

For example, the CEO of a fast-growing tech company was pulled into a crisis call early one morning. An enterprise client had just discovered a security flaw in a major software release. The firm’s reputation hung in the balance. As the client’s tone grew more critical, the CEO knew that everyone was watching her demeanor rather than listening to her words, so she deliberately slowed her cadence, kept her voice level, and asked one clarifying question. Other participants on the call settled down because she remained calm. Her CTO noted that the issue was probably a false positive and offered two rapid-response options. Neither was perfect. She looked at her colleagues. The security lead appeared confident, the product manager looked uneasy, and the risk officer hesitated. She ran a quick internal audit of her own thinking: *Am I reacting to a fear of reputational damage or responding to the facts?* Only then did she tell

the client to proceed with a controlled rollback while the team validated the alert. Later, the engineers confirmed the flaw was minor. The CEO’s measured approach turned out to be the right call.



AFTER

In business, the aftermaths of decisions are often reviewed. In pro sports, that step is a distinct phase of decision-making—one deeply embedded in the rhythms of competition. Coaches routinely watch game footage with their staff members and teams—analyzing key decisions, tactics, and moments under pressure—as part of their preparation for the next opponent.

They normalize being wrong. Top coaches are remarkably comfortable acknowledging when choices lead to failure. As one international team rugby coach put it, in “spur-of-the-moment decisions...the chances of making an error are high.” It’s therefore key to not only be aware of the risks but also accept one’s own errors and communicate about them with others to model a growth mindset. An experienced NFL coach has a standard response after he makes a mistake: “I may have got that wrong, but I always learn for next time.”

In one case, the star receiver suffered a hip injury that medical staff estimated would take three to six weeks to recover from. The coach and his team had to decide within 72 hours whether to put the player on injury reserve, ruling him out for a guaranteed four weeks but allowing them to bring in a replacement, or to keep him

on the roster in the hope of an earlier return. With imperfect information, they opted for the first course of action. When the star was ultimately fit to play earlier than expected, the coach acknowledged to his team and staff that keeping the star active would have yielded a better outcome but explained that he was still content to have erred on the side of caution. His aim, he told us, was to communicate that “there can be multiple good decisions.” The goal is not to be perfect but to choose well among viable options and own what happens as a result.

They repair trust. Across sports, people’s trust in one another directly shapes decision quality under pressure. That’s why after decisions—especially difficult or contested ones—coaches are deliberate about ensuring that any damaged relationships are fully restored to health.

One NBA coach described a heated in-game conflict with a staff member after a breakdown in communication around a player’s minutes restriction in a big game. After a cooling-off period, they both addressed it directly: “We [each] said what we could have done differently.” He reflected that the issue was not just the decision itself but how it had been handled—in both communication and alignment.

Trust is not automatic. It is constructed through daily interactions. What follows a decision—how it is explained, owned, and discussed—shapes whether trust is strengthened or weakened before the next one.

They upgrade the system. The most successful coaches do not stop at simply reviewing decisions or repairing relationships. They go further, asking not just what happened but why it happened and what needs to change. Their focus shifts from the decision itself to the conditions that shaped it.

One European football coach described how a misjudged contract signing—one that, in retrospect, he

wished he’d pushed back on—led him to fundamentally overhaul his decision-making process. He became more deliberate in how he gathered information, clearer on the role the player would fill, and more explicit in how he evaluated fit. “If I have the evidence, I can convince others why I want this player,” he said. The structure gave him the confidence he needed to challenge others, to hold his position, and when merited, to say no.

The best managers in business settings employ those same principles after a high-stakes decision. They acknowledge when they’ve made the wrong call. When there is debate, conflict, or resentment in the aftermath of a decision, they circle back with all stakeholders and identify how to do better next time. And then they implement changes that will improve the process.

After a leader of a pharmaceutical company made a tough call in a live meeting with representatives from the FDA—agreeing to delay a drug submission by six weeks so that her team could gather new safety data—she knew that her clinical lead felt blindsided, the regulatory team felt exposed, and commercial executives were furious. Recognizing that a hot debrief would only amplify the negativity, she asked the team members to regroup in 24 hours, not to relitigate the outcome but to walk through the process. When they did, she asked, Was the safety signal communicated consistently? Where did assumptions replace clarity? Did we have the right people in the room?

The discussion revealed important subtext: The clinical lead had been hesitant about the data but never raised his doubts with her, and the commercial

team had interpreted the risk of delay as lower because of mixed messaging. She’d made the right decision, but the process had been flawed. She made a change: All safety concerns would now be reported to her, and she would decide on and communicate any resulting timeline changes as soon as possible.

THE COMPARISON BETWEEN decision-making in sports and business has its limits. Sports offers a visible scoreboard and clear wins and losses, whereas business outcomes sometimes unfold more slowly and with greater ambiguity. Yet both environments present leaders with the same fundamental challenge: making high-stakes decisions under constraint, with incomplete and sometimes conflicting information.

Our study of elite sports leaders makes that process visible. What appears to be instinct is usually the product of preparation, emotional control, pattern recognition, social awareness in the moment, and accountability in the aftermath. Using the playbook of sports for inspiration, business leaders can develop capabilities and build systems that help them make better decisions when it matters most. ■

HBR Reprint R2604J

 **ALAN MCCALL** is a sports performance consultant and founder of Research & Innovation for Sports working with elite teams worldwide. **ADRIAN WOLFBURG** is an organizational scientist and the founder and CEO of Organizational Insight Consulting. **JOHANN BILSBOROUGH** is a sports scientist and the director of medical, performance, and innovation in the NBA. **RICARD PRUNA** is the head of medical services at FC Barcelona.



How We Did It

The Former CEO of Procafecol on Building a Global Growth Strategy for Juan Valdez Coffee

by Camila Escobar

WHEN I JOINED Procafecol seven years ago it was ready for a new stage of global growth. Our company, created by Colombia's National Federation of Coffee Growers (FNC) in the early 2000s to bring the Juan Valdez brand to retail consumers, had a strong foundation in our home market. But over the years we'd expanded into almost 40 other markets, from as nearby as Aruba to as far away as Australia.

These varied international experiences taught us some valuable lessons, but we recognized that to build a sustainable, scalable, and successful long-term global business, we would need to focus our efforts. The 2020 pandemic pushed us to be even more deliberate in our approach.

By 2022 we had shifted our expansion strategy. With Colombia remaining our home base, we selected five priority markets in which to search for exponential growth: Brazil, Mexico, Spain, the United Arab Emirates, and the United States. At the same time, we continued to look for ways to bring Juan Valdez to new areas and to seize attractive opportunities.

This evolution required important changes: a new mindset about our mission, a stronger leadership structure, and more-thoughtful decisions on partnerships, investments, and talent.

Four years later, we are already seeing positive results. We're outperforming our targets in each key market and logging double-digit growth. Last year, we achieved our highest EBITDA in history.

BUILDING THE BRAND

Procafecol's story began nearly a century ago, in 1927, when producers across Colombia came together to create the FNC. In 1959 the FNC created a brand to represent the origin, quality, and identity of Colombian coffee to the world: Juan Valdez, a fictional farmer with a signature hat, poncho, and mule standing in front of the mountain Cerro Tusa. The idea was powerful; Juan Valdez was not just a character but a symbol of our country's exceptional coffee growers.

Over time the brand became a solid representation of Colombian origin to coffee consumers, both at home and internationally. But by the early 2000s coffee consumption was changing. People weren't just drinking it at home; they wanted café experiences. Large coffee chains began to gain more of the market, while growers were receiving only a small percentage of sales—in some cases as little as 6% of the price of a pound of coffee.

That shift forced the FNC to ask an important question: How can we get closer to the consumer and capture the end value of a cup of coffee in favor of coffee growers? The answer was clear. We needed to move beyond a brand that represented an origin and start building a business around it. To make that happen, the FNC created Procafecol in 2002—a company that would use the Juan Valdez trademark, source its coffee

from growers, and pay royalties to the National Coffee Fund, which maintained ownership of the brand. The same year, the first Juan Valdez Café opened at Bogotá's El Dorado International Airport—a location chosen to attract both local and international customers. As we expanded across the country, Colombians embraced the cafés, not only for the coffee but also for what they represented—the opportunity to support 550,000 small coffee-growing families.

We began to look beyond Colombia's borders, and Latin America and the Caribbean were natural first steps. Then we expanded into more diverse markets, including the United States, Spain, Kuwait, and Malaysia. By 2020 we had opened 535 cafés; about a third of them were outside Colombia and operated as franchises and joint ventures. We were selling coffee in nearly 1,500 outlets across Colombia and close to 10,000 internationally, supported by institutional and e-commerce channels.

THE STRATEGIC SHIFT

My journey to Procafecol began during college, when I took an internship at the FNC. After graduating, I started my career at McKinsey and later moved into corporate strategy, where I developed a strong understanding of how to drive growth and make disciplined decisions.

I was fulfilled in those roles. But when former colleagues from the FNC approached me with the opportunity to lead Procafecol, I couldn't pass it up. It meant coming back to one of the most beloved brands in my home country—and being part of something more meaningful. I'll admit I wasn't a big

coffee drinker at the time, but that has certainly changed.

I stepped in as CEO in 2018, and shortly after, we began shaping the new global growth strategy, which required decisions about what to prioritize and what to deprioritize. Then in 2020 the pandemic hit and forced us to rethink. With stores closed and revenues under pressure, our growth strategy had to get even more focused.

Two things were clear: First, we needed to reduce our dependence on Colombia, which represented more than 80% of our sales. And second, our global expansion choices had to be about more than proximity and chosen much more deliberately.

We started looking at a combination of factors—market size, income levels, coffee consumption, pricing, and ease of doing business. We had to make tough calls; we walked away from markets that looked attractive on paper but where it would be too complex for us to execute well. What started as a plan to prioritize 10 markets quickly narrowed to targeting five. We also set an ambitious goal: to increase international sales from 20% to 40% of total revenue.

Some of the choices we made were not obvious. For instance, most of the regions we prioritized were already competitive coffee markets. But we reasoned that if major players were not there, there was probably a reason. At the same time, we don't see this as a zero-sum game—when coffee consumption grows, everyone benefits. Our challenge is to make Juan Valdez stand out. We want to be the most relevant premium Colombian coffee brand and to create not just economic but also social and

Top to bottom: a Colombian coffee grower working among the coffee trees; bags of coffee bearing the Juan Valdez logo

environmental value. (Indeed, in 2021 we became a Certified B Corporation, signifying our commitment to and progress on the latter goal, and we have committed to net-zero emissions by 2050.)

In each priority market, our preference is to build partnerships through joint ventures while keeping control over the brand and the strategic direction. In many cases this is a departure from our previous approach. In the United States, for example, we initially opened flagship cafés in large cities like Seattle and New York. They helped build awareness, but the business wasn't scalable. Today we are choosing U.S. locations in a much more targeted way. We started in Florida and expanded only after proving our model there works. At the same time, we pursue retail partnerships to help us reach the broader grocery market more efficiently.

Spain plays a similar role for us in Europe. Working with a partner that has a strong local presence, we started with three stores in 2021; we'll have more than 20 by the end of this year. Collectively they're already turning a profit.

In Latin America, Brazil and Mexico are essential for regional relevance. Brazil might seem like a risky bet given that it grows a lot of its own excellent coffee—enough to make it the biggest producer in the world. But we realized that there might still be interest in Colombian offerings after we were invited to a coffee expo there in 2022. Our stand was surprisingly popular, and as we were closing up many people asked if they could buy the beans we had left. They saw our selection of coffees as premium, milder alternatives to their homegrown ones. Studies also show that



Brazil is embracing international coffee chains and global brands in general. So we searched for the right partner and opened our first café last year. Within three months it achieved the returns we had expected to see in 12. Our grocery sales in Brazil are also picking up.

Mexico remains a work in progress. Although consumers there love our brand, it has been a challenge to find partners. Before I arrived at Procafécol, we had 15 cafés in Mexico with one franchisee, but the business wasn't successful enough for him to continue. Now we have two franchise operations there, and we hope their success will attract

interest from large investors so that we can sign a national joint venture.


Finally, in the Emirates, in partnership with a local company, we have three cafés.

While we stay disciplined in our priorities, we are also committed to consolidating our presence in mature markets for our brand where existing joint ventures are working well, such as Chile and Ecuador. And we remain open to opportunities. Markets that were not part of our initial focus but have performed strongly (such as Turkey, where we have more than 30 cafés, all overshooting their targets) remind us of the

Courtesy of Juan Valdez

Procafecol

Founded: 2002
Headquarters: Bogotá, Colombia
No. of employees: 2,800



Source: Procafecol

importance of balancing focus with flexibility. By the end of 2025 we had more than 675 stores and 17,500 points of sale in supermarkets around the world.

SAME VALUES, NEW MINDSET

How do you transform a company that has been primarily focused on its home market and dabbling everywhere else into a more strategic global operation?

It starts with culture. Our core values haven't changed: service (collaborating with colleagues and partners to deliver great products and experiences for customers); integrity (always doing the right thing, even when nobody's looking, and not doing bad things, even if others in the market might be); diversity (bringing different perspectives to the table); and *colombianidad* (embracing the Colombian spirit of being welcoming and kind, innovative, and resilient).

However, now our people know that our mission is to take those ideals around the world. For some team members,

the shift stung at first. It meant giving our domestic operations proportionally less investment and talent. But because we've carefully outlined our new plans—and explained why they're key to our long-term growth—everyone understands the trade-offs, and many have volunteered to join our operations around the world as local hires, ensuring that we bring our Procafecol and Juan Valdez knowledge and best practices to the new ventures. Some have even moved from one foreign location to another, so talent has begun to flow naturally between our operations.

Structurally, we have reorganized the executive team. Whereas before we had one chief commercial officer for Colombia and one for the rest of the world, we now have three chiefs for three clusters: North America, Latin America (including Colombia), and Eurasia. We appointed a global chief marketing officer because it's important that we have one distinct global brand. And last year we unveiled a new global marketing campaign.

Perhaps the most critical piece of our new strategy is the selection of joint venture partners. Our values lead the way. We look for groups that share our service mindset, that operate with the highest ethical standards, that are excited by different ways of thinking, and that embrace the *colombianidad* spirit. When we meet with potential investors, it's a simple determination: They are either moved by the same things we are and we continue to engage, or they're not and we don't.

Every café we operate is designed to offer an upscale experience and emphasize Colombian culture; some feature work by Colombian artists. They

all employ baristas rather than having automated machines. And our procurement strategy is the same everywhere: We source from FNC farmers only, and we send their coffee, in many different varieties, to every location in which we operate. At the same time, we do collaborate with our local partners on some customization, for example by incorporating country-specific preferences. While a cappuccino, a latte, and a filtered coffee are standard everywhere, a cold brew in the Emirates might be flavored with dates and so we do a Turkish coffee preparation for that market. Our food offerings are both seasonal and super-local—overseen by our central marketing team but deeply informed by our partners. We invest in some local trade marketing and, of course, adapt prices, most recently adjusting for U.S. tariffs, although we are absorbing some of that extra cost to stay affordable.

SIXTY-SEVEN YEARS AFTER Juan Valdez made his debut, Procafecol is excited to be building a robust global business around him. It is still early days. But by committing to a more outward-looking mindset, targeting and reorganizing around growth markets, choosing joint venture partners wisely, and working with them to best serve our consumers, the company and brand are now on an upward trajectory. Although I recently stepped down from the CEO role to lead a nonprofit foundation in the education sector in Colombia, I remain excited about Procafecol's future. We hope other companies with international ambitions can learn from our process. ■



experience

HBR Reprint R2604K

MassMutual

175
YEARS



A legacy of purpose. A future of possibilities.

For 175 years, MassMutual has helped people secure their future and protect the ones they love. Our purpose is more than our history. It's who we are today – and tomorrow.

HBR's fictionalized case studies present problems faced by leaders in real companies and offer solutions from experts.



A Key Supplier Failed a Labor Audit. Now What?

by Sandeep Puri

PAULA SANTOS, the chief sustainability officer of Vestenborg Home, listened as the third-party audit team explained its findings. It was the final review session of her annual supplier audit cycle, part of a commitment the Danish furnishings company had made several years earlier. The goal was to ensure that 100% of Vestenborg's vendors complied with strict ethical standards covering labor practices, health and safety, and environmental impact. The auditors had spent weeks in India, Vietnam, and Thailand to inspect facilities, review

records, and interview workers. Most of what Paula had seen so far was familiar: a few documentation gaps and delays in conducting safety drills.

But then came the report on Greenspan. For 14 years the Vietnamese foam producer had been a cornerstone of Vestenborg's supply chain. Currently it was tasked with providing the foam for the company's new EcoLuxe mattress, its most ambitious product launch in a decade and the centerpiece of its global sustainability campaign. Vestenborg, meanwhile, represented roughly



a third of Greenspan's total production volume.


Previous audits of the supplier had identified minor issues, such as administrative lapses and incomplete logs. But the reviews were otherwise clean, and corrective action had been swift. Greenspan had always been exceptionally reliable and responsible—until now.

"Our three findings are all serious issues," said Tom Barney, the audit lead. The first problem was excessive overtime. Greenspan employees had routinely logged between 70 and 84 hours per week over the past few months. Vestenborg's cap was 60.

The second finding had to do with safety standards: An emergency exit in the factory had been partially blocked by stored raw materials.

"The final problem is the most serious," said Tom. "Two workers at Greenspan were only 14 years old. Both were hired through a subcontractor, their identification documents apparently falsified. The discrepancy surfaced during employee record cross-checks and confidential worker interviews."

Paula's head was spinning. The timing could not be worse, because Vestenborg's mattress launch depended on Greenspan's foam core. Internal forecasts projected strong first-year revenue and margins for the product. Marketing calendars were fixed. Retailers were aligned. Greenspan was the only supplier capable of producing the foam within the launch window; finding



another would take months. If Paula had to tell Greenspan to halt production, EcoLuxe would be delayed.

As the auditors packed up their materials, Paula made a decision. Greenspan had been one of Vestenborg's most trustworthy suppliers for over a decade. So, before she shared these findings with the rest of the executive team, she wanted to understand what had happened—and whether it was contained. She asked her assistant to put her on the next flight to Vietnam.

RULES COLLIDE WITH REALITY

Greenspan's factory in Binh Duong was humming when Paula arrived. Forklifts moved steadily between curing lines. Workers in uniforms trimmed foam cores for what she knew were

EcoLuxe mattresses; a finished pile was stacked in labeled batches along the wall. The whole factory seemed clean and well-organized. An executive assistant led Paula upstairs to a modest office, where a young man rose to greet her.

"You must be Nguyen Van Minh," Paula said. From past interactions with Minh's proud father, the longtime head of Greenspan, she knew that Minh had joined the family business after spending years in Australia for university and business school. But she was surprised that his father wasn't present.

"Very nice to meet you," Minh said. "My father has stepped back due to health issues. I lead the company now."

"Oh. When did you take over?"

"Almost a year ago. Dad isn't going to come back to work, but we only realized



When faced with a conflict between compliance and what he saw as a morally better outcome, Minh had followed his own judgment rather than an auditor's checklist.

that recently. I wanted to tell your leadership team in person.”

While the assistant poured two cups of tea, Minh jumped into a speech he had obviously rehearsed. “This is a new era,” he began. “We are upgrading facilities, investing in automation, improving conditions for workers. We want to grow with Vestenborg for the long term.”

“I understand the partnership is important to you,” Paula said. “That’s why I’m here.” She pulled out the audit report and opened it to the section on Greenspan. “We need to talk about the underage workers—that’s a zero-tolerance policy for us,” she said.

“Those workers came through a subcontractor. We terminated them.”

“Yes, but they were on your floor when the auditor visited.”

There was a pause. “This is not the full story,” Minh said carefully. “The boys were orphans working at a roadside food stand. A supervisor brought the situation to my attention. We believed giving them employment was better than leaving them on the street.”

“But the age discrepancy?”

“We knew their documents were... unclear,” he admitted. “We believed their jobs here would be temporary while we helped them find a home. We intended to resolve it quietly. Unfortunately, it took several months.”

“Was their hiring about labor costs?”

“No. We pay all workers the same. It was not cost-driven.”

“I would be more inclined to believe you if there weren’t other violations.”

Minh pushed his chair back. “Let me show you something,” he said, and led her downstairs to a section of the factory that had been reconfigured. Against

one wall was a rest area with bunk beds, neatly arranged with clean sheets and a wall of lockers. “We installed this during the EcoLuxe ramp-up,” he said. “Workers were doing long shifts. It was difficult for some to commute home and back between rotations. We wanted them to rest properly. We have a goal of eliminating all industrial accidents, and a proper rest area was part of our plan.”

“But why were they working long shifts in the first place?” Paula asked.

“Peak months were heavy,” he said. “Many employees worked more than 60 hours per week. We couldn’t find and train more people to avoid the overrun. I regret this; it was poor planning. But we compensated generously. Overtime was paid at premium rates. Many employees preferred the extra income.”

Paula glanced toward the emergency exit that had been obstructed. “The materials blocking the exit?” she asked.

“We moved inventory to create this resting space. It was not permanent.”

“When you made these decisions, did you consider informing us?”

Minh hesitated. “As the managing director, I must make hard decisions. I believed we were acting responsibly. We did not want to cause alarm.”

Paula sighed. On paper the violations were clear: underage labor, excessive overtime, a blocked exit. In practice they were wrapped in motives that were harder to dismiss. Having taken over for his father, Minh was trying to modernize the family factory, improve worker welfare, and meet the demands of his client’s most important launch.

But he had also demonstrated something else: When faced with a conflict between compliance and what he saw

as a morally better outcome, he had followed his own judgment rather than an auditor’s checklist. Paula wondered whether he would do so again.

As she returned to the hotel, a question crystallized in her mind: Could a sustainability-driven company like Vestenborg Home tolerate a partner that believed rules were flexible when the cause seemed noble? She decided to explain this complexity to the rest of the management team and ask for advice.

WHAT KIND OF COMPANY ARE WE?

At Vestenborg headquarters two days later, the executive conference room was quiet as Paula described the audit and her visit to Vietnam.

A few seats down, Aylin Demir, the head of corporate affairs, broke the silence. “If this surfaces publicly—particularly the underage labor—the nuance that Paula discovered disappears,” she said. “The news stories won’t lead with ‘compassionate intent.’ They’ll lead with ‘Vestenborg supplier knowingly employs children.’ The sustainability positioning of our green product will seem hypocritical.”

Erik Lund, the chief operating officer, spoke next. “Let’s be clear about the operational side,” he said. “If we suspend Greenspan, the EcoLuxe launch slips. That’s a material revenue impact.”

“A delayed launch is painful,” said Freja Kristensen, the chief marketing officer, folding her arms. “But reputational damage tied to hypocrisy? That’s far worse. If we’re seen as compromising on child labor, the ‘eco’ mattress campaign collapses anyway, and it would be the least of our problems.”



“What’s our legal requirement?” cut in Daniel Hartmann, the CEO.

“There’s no automatic trigger,” replied Niels Thomsen, the general counsel. “European due diligence frameworks require remediation and documentation. They don’t mandate

terminating the supplier relationship. North American and Japanese stances are similar—we’re expected to demonstrate responsible action.”

Paula took a breath. “Since I returned from Vietnam, I’ve spoken to several other chief sustainability officers off the record,” she said. “Some suggested suspending the supplier immediately. Others would remediate intensively and keep production going. Many companies publicly claim

to have zero tolerance in these situations, but often they make contextual judgments.”

“So there’s no industry standard,” Daniel said, nodding slowly. “Then let’s define our options.”

“Option one,” Paula began, “is immediate suspension until Greenspan can demonstrate full compliance through independent verification. That would ensure our standards are met but would almost certainly delay EcoLuxe.

How should Vestenborg respond to Greenspan’s violations?

THE EXPERTS RESPOND.



NOEL KINDER is the SVP of sustainability at Lululemon.

First, Paula should have a sit-down with Minh.

She must reaffirm Vestenborg’s values, policies, and code of conduct, making clear that compliance isn’t a bureaucratic exercise: The company’s employees, customers, and investors expect certain non-negotiable standards.

From there she should work with her operations and sourcing teams to develop a remediation plan for the audit violations. That often means giving the supplier a defined period—say, 60 days—to create a credible plan, followed by

another audit to assess whether it’s translating into real change. The goal is to fix problems and to answer a deeper question: Does Greenspan still have the willingness and the capacity to meet Vestenborg’s standards?

How the supplier responds will be telling. If Minh continues to explain away violations, that signals misaligned values. But if he is collaborative and receptive, there is hope for a path forward. One meeting is not enough to make that judgment; Paula needs to watch how he acts over time.

Several of the violations point to capacity constraints rather than malicious intent. Take the bunk beds that Minh proudly shows Paula. While they’re presented as a benevolent gesture toward workers, they may indicate that the factory can’t absorb Vestenborg’s production demands. That should trigger a conversation about why breaking the rules felt necessary and what needs to change so that it doesn’t happen again. In many cases the answer involves better production planning, more-realistic lead times, or investments in workforce systems—areas

where brands often bear some responsibility.

Paula then needs to establish a rhythm of engagement, including specific milestones and check-ins, to repeatedly reinforce expectations. If remediation fails, she has an obligation to disengage responsibly. That means giving the factory sufficient notice—often months, sometimes longer—to replace the lost volume and avoid mass layoffs while Vestenborg moves production elsewhere.

The need for an extended assessment means it’s safe to assume that Vestenborg is committed to Greenspan for the mattress launch, so Paula has to manage internal dynamics. Finance, sourcing, marketing, and communications should align around a shared narrative. If questions arise from nongovernmental organizations, journalists, or investors, the company should have a clear, honest response: This is a long-term partner, violations were identified, and Vestenborg is working with the supplier to investigate and remediate them.

Ultimately, Paula’s challenge is to run a disciplined, collaborative



“If we’re seen as compromising on child labor, the ‘eco’ mattress campaign collapses anyway, and it would be the least of our problems.”

“Option two: Continue production under strict oversight. We’d need on-site compliance monitoring, weekly reporting, and formal corrective milestones. This option would preserve the launch timing—but if the violations are discovered, it could appear as though we continued business with a supplier that tolerates child labor.

“Option three: a hybrid. Suspend Greenspan, similar to option one, but court other suppliers in case we need a

new foam producer. This would likely delay EcoLuxe’s launch, but it could give us more options later on.”

“It’s a tough call,” Daniel said, his candor loosening the tension in the room. He looked at Paula. “I’ll support your recommendation. We know what’s at stake. The wrong decision could derail our flagship launch—or deal a reputational blow to everything we claim to stand for.”

Paula glanced at the EcoLuxe revenue projections and then at the

audit report on the table in front of her, the team’s comments pushing her thoughts one way and then another. She knew that this decision was not just about a supplier, a product, or even a launch date—it was about what Vestenborg’s sustainability promise truly meant. ■



SANDEEP PURI is a professor of marketing at the Asian Institute of Management.

process that reveals what kind of partner Greenspan really is. If she does that methodically and transparently, she’ll be able to make a decision she can stand behind, whatever the outcome.



AURET VAN HEERDEN is the founder and CEO of Equiception Business and Human Rights.

Paula should audit Greenspan again and observe the factory over several days.

This case may seem benign relative to other supply chain scandals, but that appearance rests heavily on Minh’s explanations, which may be accurate—or may not. Although Vestenborg has a long history with Greenspan, Minh is new to the relationship; Paula didn’t even know about his leadership transition. She hasn’t built trust with him, so the situation requires deeper investigation.

Some of Minh’s explanations are plausible. Employing orphans or vulnerable youth does happen in parts of Asia, often with local government knowledge or approval. But given current labor shortages, Vietnamese factories sometimes turn to questionable labor brokers or informal arrangements that are far more nefarious than what Minh describes. Best case, Paula confirms that the violations are limited and explainable. Worst case, she uncovers risks that make the Greenspan relationship untenable.

If the next audit finds true malfeasance, Paula should terminate the contract, delay the EcoLuxe launch, and find a new supplier. But otherwise she should focus on remediation. International

guidelines from the United Nations and the Organisation for Economic Co-operation and Development encourage doing so unless a supplier is a serial offender or acting in bad faith, because abruptly cutting ties leads to job losses and worse outcomes for workers. Importantly, Paula is on solid ground legally. Most due diligence regimes require good faith efforts to address problems, not guaranteed outcomes. Remediation provides that safe harbor while preserving leverage over the supplier.

Paula’s efforts should begin with a root cause analysis, and it might point to Vestenborg. Overtime violations in large brands’ supply chains are often driven by rigid deadlines, short lead times, and unrealistic forecasts. If EcoLuxe’s ramp-up put too much pressure on Greenspan, demanding compliance without adjusting expectations will guarantee repeat violations.

Greenspan will need hands-on support during remediation. Vestenborg could provide experts to help improve scheduling, safety systems, and workforce planning—changes that reduce overtime and

accidents while lowering costs. In more difficult cases companies have had experienced executives shadow local management.

Paula must also prepare for outside scrutiny. When companies with environmentally responsible products are accused of unethical supply chain management, they become easy targets for watchdog groups. Paula should proactively engage relevant NGOs, confidentially explaining the situation and outlining the remediation plan. Brands that build trust with such stakeholders are often granted time to fix problems quietly, which helps protect workers.

A commitment to improvement, not perfection, is how organizations develop fully ethical supply chains. If Paula investigates rigorously, remediates seriously, and manages risk transparently, she can protect both the EcoLuxe launch and Vestenborg’s credibility. In the real world of global manufacturing, that is what responsible leadership looks like. ■

HBR Reprint R2604L

Reprint Case only R2604X

Reprint Commentary only R2604Z

Executive Summaries

July–August 2026

CHANGE MANAGEMENT

ORGANIZATIONAL DECISION-MAKING

STRATEGY

LEADERSHIP

CUSTOMER EXPERIENCE



The False Alignment Trap

Julia Dhar, Kristy R. Ellmer, and Philip Jameson | page 46

Most organizational change efforts fail not because of poor execution but because senior leaders fall into a *false alignment trap*—believing they agree on why, what, and how to change when they actually do not. Drawing on research, behavioral science, and a case study of the jewelry group Pandora, this article explains how false alignment arises from vague discussions, avoidance of disagreement, and pressure to move too quickly and shows how it leads to paralysis, wasted activity, misdirected progress, or watered-down compromises. The authors propose a disciplined process for reaching true agreement, emphasizing specificity, early and safe dissent, rigorous debate, clear decision rights, formal commitment, and unified communication. They argue that investing time up front to resolve disagreements ultimately accelerates execution and improves outcomes.

HBR Reprint R2604A



What Companies Get Wrong About Decision Rights

Lindy Greer, Jennifer Jordan, and Maxim Sytch | page 58

Many organizations use tools like RACI to clarify who should provide input on decisions, who should make them, and who should carry them out. But in practice these frameworks often fail because teams don't apply them properly. Drawing on their research and their work with more than 100 companies, the authors identify four common mistakes: setting decision roles before clarifying goals, treating decision rights as a static list created by a single senior leader, misunderstanding what each role actually means, and letting hierarchy override the assigned roles. To avoid these issues, leaders should cocreate roles with the people involved in or affected by the decisions, define clear behaviors for each role, revisit role assignments regularly, and have ongoing conversations about them. Used well, tools like RACI will reduce conflict, improve collaboration, and increase the quality of decisions and the speed at which they're made.

HBR Reprint R2604B



The Power of Strategic Centering

Rita McGrath | page 68

In a dematerializing economy—where value has shifted from physical assets to intangibles such as data, software, and capabilities—traditional strategy frameworks no longer provide enough guidance. Organizations need a clear *center*: a single, coherent organizing principle that defines what they are really about. A strategic center guides resource allocation, opportunity selection, and organizational identity. Columbia Business School professor Rita McGrath draws on real-world examples to show how companies that choose and commit to a center—whether it's around a mission, a customer, a technology, a regional or national ecosystem, or friction erasure—gain clarity, speed, and coherence while reducing internal politics and paralysis. Effective centering is not about rigid positioning, she explains, but about choosing the dimension along which a company will pursue coherent opportunity sets as industries and old strategy anchors dissolve.

HBR Reprint R2604C



Making the Leap from Corporate Leader to PE-Backed CEO

Samantha Hellauer et al. | page 80

Private-equity firms have long favored CEOs with prior experience at portfolio companies, but talent shortages and the rapid expansion of PE ownership are pushing investors to recruit more leaders from the corporate world. Not all make the transition successfully. Leading a PE-backed firm requires five capabilities: a practical commercial orientation, the ability to tackle strategy under pressure, the ability to wield influence widely to achieve impact, a willingness to take risk (especially on talent), and a broad interpersonal range. Unlike most corporations, PE firms demand immediate value creation, lean execution, and direct, high-candor engagement with investors and boards. Rigorous assessment of candidates' ability to deliver those things—and honest evaluation of their motivation and readiness—can help portfolio companies increase the odds of finding CEOs who are a good fit.

HBR Reprint R2604D



Why the Best Immersive Experiences Succeed

Joseph C. Nunes and Wendy Heimann | page 88

Immersion is increasingly defining customer experiences in a variety of industries, yet too often creators of immersive environments rely on spectacle, advanced technology, or scale, which fails to create a lasting impact. The most effective experiences are based on the psychological state of absorption and guide participants through a progression in which they seek answers to six questions: Where am I? Who am I with? What can I do? What is happening? Am I making progress? Why does this matter? Many experiences address one or several of these dimensions, but those that help people answer all six questions are the ones that draw them in, keep them engaged, and culminate with a coherent sense of what it all meant. For companies, the payoffs of building effective immersive environments can be substantial, from increased revenues to higher customer satisfaction.

HBR Reprint R2604E



Six Ways Leaders Harness Stress

Jon Miller and Drew Keller | page 100

Today's leaders face increasing pressure on all sides, and their stress levels are higher now than they were even at the peak of the pandemic. Though stress can sharpen performance briefly, over time it erodes judgment, narrows perspective, and increases the risk of costly missteps. Most leaders have distinct default responses to it. This article outlines the six most common patterns: the calm lighthouse, the reinvention-oriented alchemist, the action-driven firefighter, the disciplined stoic, the relationship-focused diplomat, and the control-driven container. Each style has both strengths and blind spots that pressure can amplify. Leaders can increase their ability to perform under duress by identifying and understanding their default responses and then deliberately expanding their range of reactions—using simple tactics to regulate themselves, share the cognitive load, and alter their style in real time as conditions change.

HBR Reprint R2604F



How Agentic AI Supercharges Startups and Threatens Incumbents

Vivian S. Lee, Linda Mantia, and Jon McNeill | page 112

Agentic AI is reshaping entrepreneurship and raising the stakes for incumbents. Startups are deploying coordinated systems of AI agents that can plan, act, and adapt autonomously. This approach dramatically compresses the time, capital, and head count needed to launch and scale up a company. Products that once took large teams more than a year to develop can now be built and refined in weeks by a handful of people. AI-native startups build proprietary workflow knowledge by using AI agents, which creates compounding advantages. Established companies face a structural challenge: Siloed data, legacy workflows, and rigid roles limit the benefits of agentic AI. Leaders must redesign processes before automating, strengthen data quality, clarify human-AI handoffs, and prepare employees to focus on work that requires human judgment or is outside the norm.

HBR Reprint R2604G



Bring Back Managing for Value

Michael Mankins and Matthew Cupri | page 122

After years of cheap capital, many companies drifted away from managing for shareholder value in favor of loosely defined stakeholder models. Now capital costs are moving back toward historical norms, and value maximization should once again guide decision-making. However, it must be applied differently than it has in the past. This article proposes a simple but disciplined model: Retain value maximization as the governing objective but explicitly recognize stakeholder commitments as constraints that limit the set of feasible choices. The authors distinguish between hard constraints (nonnegotiable commitments) and soft constraints (which guide choices among value-creating alternatives) and describe four foundational capabilities that companies must rebuild. Leaders can use this model to evaluate alternatives consistently, allocate capital to its highest-value use, and address legitimate stakeholder claims without sacrificing accountability.

HBR Reprint R2604H



How Elite Sports Coaches Make High-Pressure Decisions

Alan McCall et al. | page 131

Drawing on interviews with 11 elite coaches across major U.S., European, Australian, and New Zealand professional sports leagues, the authors show that high-pressure decisions aren't a single "formula" but a set of practices that unfold before, during, and after pivotal moments—and translate directly to business leadership. Top coaches combine disciplined preparation with emotional control and social awareness during crunch time and accountability and continual system improvement after the fact.

HBR Reprint R2604J



The Former CEO of Procafecol on Building a Global Growth Strategy for Juan Valdez Coffee

Camila Escobar | page 136

Procafecol was created in the early 2000s to bring the Juan Valdez brand to retail consumers. The company has recently entered a new stage of global growth, maintaining its strength in Colombia but paring down its foreign priorities to five markets, with the goal of significantly increasing international sales. The strategy is built on disciplined selection, joint ventures that keep brand control, and a plan to scale at the right pace in each place. To execute this shift, Procafecol reemphasized its values, reorganized its executive team around regional clusters, and chose partners with aligned mindsets.

HBR Reprint R2604K


POSTMASTER

Send domestic address changes, orders, and inquiries to: *Harvard Business Review*, Subscription Service, P.O. Box 37457, Boone, IA 50037. GST Registration No. 1247384345. Periodical postage paid at Boston, MA, and additional mailing offices.

Printed in the U.S.A. *Harvard Business Review* (ISSN 0017-8012; USPS 0236-520), published every other month for professional managers, is an education program of Harvard Business School, Harvard University; Srikanth Datar, dean. Published by Harvard Business School Publishing Corporation, 60 Harvard Way, Boston, MA 02163.



“You can teach skills, but you can’t teach passion. If you never stop asking, ‘What can I learn today?’ you are my kind of person.”



José Andrés

*Having learned to cook in his parents' kitchen, Andrés trained under several of Spain's top chefs before immigrating to the United States and leading his first restaurant as head chef in Washington, DC. In the three decades since, he has not only built a culinary empire that stretches from Los Angeles to the Bahamas but also founded a high-profile nonprofit, World Central Kitchen, that provides fresh meals on the front lines of crises. His latest cookbook, *Spain My Way*, came out this spring.*

Interview by Alison Beard

HBR: What lessons did you learn from your famous mentors?

ANDRÉS: Ferran Adrià taught me that you can't be a good chef if you don't love to eat. And you become a great chef by wanting to pass that love to others through your cooking. He was also always reading books about traditions, dishes, great chefs, and restaurants of the past, but at the same time, was the guy who said, “When we master this, we need to forget about it. We should not be settling for what we are taught. Nothing is so sacred that we cannot change or improve it.”

What kind of leader are you?

I think I'm a leader in the making. The more I know, the more I know I know nothing. There is no perfect recipe that works every time in every circumstance for every person. It's also important to understand that our business is one of the most difficult to be in. A lot of restaurants that seem to be doing well can't pay the bills. So I am always clear that we must have not only great food and great service but also an unbelievably good sense of the economics, so we can keep doing what we love.

How do you create new, innovative restaurants under one consistent José Andrés brand?

The early days were easier because I had all my restaurants around where I lived and worked, so I could walk from one to the other. That was very fulfilling because in two hours, I could spend 15 minutes in each and then decide which one to give more time to that day. And I really watched the restaurants create a neighborhood. For me, new

restaurants are a way to learn other stories. Spain was easy, but I like Lebanese, Greek, Turkish, and Mexican cooking, too. So why not put my creativity and curiosity into exploring these other cuisines? Some say this is appropriation, but give me a break. You're expanding, which is what immigrants do. The restaurants made me work harder to learn.

What did you bring from your culinary career to your humanitarian work?

When a big earthquake hit Haiti in 2010, I was nearby in the Cayman Islands and said, “I'm going.” I began cooking in a few camps and realized that while it's not easy to feed people in a disaster, it's also not so difficult because there is always food somewhere, and there are always volunteers, chefs, cooks, caterers, warehouses, and distributors. So I brought this mindset that the restaurant community is the world's biggest and most powerful. We are all one phone call away from each other, and with this network we can very much make sure that every problem is solved. The other thing about restaurants is that nothing ever goes as planned. The table of eight is now 14. You are a seafood restaurant, and one diner eats only vegetables. A pipe breaks, or the computer system goes down. This is what the chef in me brought to the way we manage at World Central Kitchen: You have to adapt to whatever situation you're in and ask, “How do we produce food and keep doubling, tripling, increasing exponentially until everybody is fed?”

HBR Reprint R2604M

Lexey Small

A woman with dark hair and glasses is speaking into a microphone. In the background, several men are listening attentively. The scene appears to be a professional conference or seminar.

BUILD A WORLD-CLASS FAMILY OFFICE WITH CLARITY AND CONFIDENCE.

The Future of Your Family Office

This new, five-day executive education program equips senior family office leaders—both active family principals and family and non-family c-suite executives—with the frameworks and strategic clarity to design a purpose-driven, multigenerational roadmap.

2026 DATES

November 2–6

Location: Gleacher Center in downtown Chicago.

Learn more and apply at chicagobooth.edu/ffo.

+1 312.464.8732 | exec.ed@chicagobooth.edu

CHICAGO BOOTH
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS



Executive Education

CHICAGO | LONDON | HONG KONG

anchor and steel
Hermès, crossing horizons



HERMÈS
PARIS