

BEYOND COMPLIANCE

Rethinking Safety Through
Leadership and Culture

DAVID G. BROADBENT
— Safety Psychologist —

BEYOND COMPLIANCE

Rethinking Safety Through Leadership and Culture

David G. Broadbent

Safety Psychologist

Copyright © 2026 by David G Broadbent

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, scanning, or otherwise — without the prior written permission of the publisher, except in the case of brief quotations embodied in critical articles or reviews.

This book is a work of nonfiction. While the author has made every effort to ensure accuracy and reliability at the time of publication, this book is provided for informational and educational purposes only.

Disclaimer — Legal, Technical, and Safety

The information contained in this publication is intended to support and promote discussion and reflection on workplace safety culture, leadership, and systems thinking. It does not constitute legal advice, professional certification, or a substitute for compliance with applicable regulatory frameworks, legislation, or workplace obligations.

All safety-related procedures and systems referenced should be reviewed and validated in accordance with your organisation's policies, relevant jurisdictional laws, and guidance from qualified safety professionals. Readers are advised to consult appropriate professionals or regulatory bodies before acting on any information contained in this book.

The author and publisher disclaim any liability, loss, injury, or risk incurred directly or indirectly as a result of the use or application of any information presented herein. Case studies, scenarios, or illustrative content are used to highlight critical safety principles. Where individuals are referenced, identities may be anonymised or fictionalised to preserve privacy.

First Edition Published by Transformational Press

For permissions, inquiries, or speaking engagements, contact: David G. Broadbent
<https://www.transformationalsafety.com>

About the Author

David G. Broadbent is internationally recognised as one of the world's leading authorities on Safety Culture and Safety Leadership, bringing a unique blend of psychological expertise and lived experience to workplace safety transformation.

David holds a Bachelor of Arts (Honours) in Psychology from the University of Newcastle, New South Wales, along with specialised qualifications including a Diploma of Work Health and Safety and an honorary doctorate in Safety Science from Dunster Business School in Switzerland. Since 2004, he has been a sought-after international speaker, presenting at prestigious venues including the International Congress of Psychology in Beijing, China. This pivotal presentation led to the development of his Transformational Safety® Leadership System.

David has pioneered revolutionary approaches to workplace safety, including the SAFE-T-SOS Psycho-Behavioural Safety Systems, developed in response to the 70%+ failure rate of traditional BBS programmes. By integrating behavioural psychology with Transformational Safety® Leadership competencies, these systems deliver unprecedented and sustained organisational outcomes. He also created amongst the world's first safety development workshops, incorporating Resilience Engineering and High Reliability Organising in his "Building a Resilient Safety Culture" program.

In late 2008, David experienced a near-fatal cycling accident resulting in multiple punctured lungs and over 24 fractures. His surgeon described the injuries as though "a hand grenade had exploded on the inside." Requiring extensive titanium reconstruction of his upper left arm and shoulder, this experience provides David with unparalleled experiential authority on accident causation, risk tolerance, and trauma recovery — insights few safety professionals possess.

David's early experience as a reservist soldier with the Australian Army set him up well for a detailed understanding of risk and he has been working with the military and first responder populations for decades — assisting these populations better navigate the barrage of personal and organisational trauma exposures.

David is one of a handful of psychologists in the ASEAN region qualified as a Master Practitioner of Accelerated Resolution Therapy. One of the world's fastest growing treatment approaches displaying great promise in the area of trauma recovery.

David is also a powerful critic of traditional safety metrics, challenging the industry's focus on lagging indicators with his observation: "If we just keep measuring what we are getting wrong by way of accidents and disasters, we make invisible what we are getting right — so don't be surprised when you don't see the 'bad stuff' coming."

David is also a Fellow of the International Institute of Risk & Safety Management (FIIRSM).

Based in Maitland, Australia, David continues his global mission to transform workplace safety culture through consulting and speaking engagements across multiple continents.

PLEASE RESPECT COPYRIGHT

Dedication

To the leaders who have felt the sting of hindsight—who have learned that systems can be compliant yet unsafe, polished yet hollow, reassuring yet disconnected from the reality of work.

To those who have discovered, sometimes painfully, that safety is not a matter of documentation but of understanding, presence, and the willingness to confront truths that metrics cannot reveal.

To the professionals who work quietly in the background—coordinating responses, supporting teams, interpreting risk, and holding organisations together through diligence, compassion, and integrity—often without recognition, often without visibility,

And to the families whose lives have been shaped by organisational decisions they never saw, whose futures were altered by risks they never agreed to carry, and whose grief stands as a permanent reminder.

Their stories are the reason this book exists, and the reason leadership must never again mistake compliance for safety.



PLEASE RESPECT COPYRIGHT

Preface

This book begins, and must begin, with an acknowledgement that will be uncomfortable for many readers who spend their working lives in the safety profession.

Most organisations that have suffered catastrophic safety failures were not, at the time of their failure, organisations that ignored safety. They were organisations that believed, with considerable evidence to support the belief, that they were managing it well. Their safety management systems were documented and certified. Their training matrices were current. Their incident statistics, reviewed at regular intervals by people who cared about the numbers, told stories of steady improvement. The boards and executives who received those reports did not misread them. They read what the system produced, and the system produced reassurance, because it had been built, without anyone deciding this was the goal, to produce reassurance.

The question this book is organised around is not why that happens. The evidence is now clear enough, across enough industries and enough decades, that the mechanism is well understood. The question is what changes when leaders understand it: not as a theoretical proposition about organisations in general, but as a practical and personal accountability for the organisations they actually lead.

I have spent more than twenty-five years working with high-hazard industries across several continents, and the most reliable predictor I have found of an organisation's real safety performance, not its measured performance but the actual state of its risk controls and the actual margin its workers operate within, is the quality of engagement between its leaders and operational reality. Not the quality of its documentation or the sophistication of its management system. The engagement between leaders and the work as it is actually done, in the conditions that actually exist, by people who have developed ways of working the documentation may or may not describe.

None of this is an argument against safety management systems, which every operation needs. It is an argument about what they cannot do on their own, and about the gap between how work is planned and how it is done, where risk accumulates quietly and where no system, however well designed, can see without leadership looking for it.

Beyond Compliance is about that leadership task. It is about the practical, sometimes uncomfortable work of leading safety rather than managing it: of

understanding the difference between an organisation that has met its compliance obligations and one that has actually reduced the likelihood that a family will receive news they cannot recover from. That distinction is not captured by any metric most organisations use, and it tends to become visible only when something has already gone terribly wrong.

The chapters that follow do not begin with what to do. They begin with the understanding that has to come first: an honest examination of how organisations come to believe they are safe when they are not, and how safety systems become virtual, present in form and absent in substance. This is not comfortable territory. It implicates the safety function, the auditors and regulators who shape how safety is measured, and, most directly, leaders who have, with the best of intentions and in good faith, managed safety without ever quite leading it.

I do not write as a detached observer. I have been part of systems that did not perform as they should have. I have sat in management reviews at which real evidence of risk was present in the room and not recognised as such, including by me. I have seen safety cultures that looked, from every conventional vantage point, like real achievements reveal themselves under pressure as something considerably thinner. Those experiences are the origin of this book, and the source of its central conviction: that safety excellence, the kind that actually protects people rather than appearing to, is not only a systems problem. It is a leadership problem, which means it is a personal problem, which means it begins with individual leaders deciding to understand things they currently do not.

This book is not written for the safety professional seeking validation of existing practice. It is written for the leader, at any level and in any role, who suspects, perhaps without having put it into words, that the gap between the organisation's safety performance and its safety documentation is larger than they know, and who has the honesty and the courage to look at it clearly. The tools and the evidence are in the chapters that follow. The decision to use them belongs to you.

The question this book answers is: what does it mean to lead safety differently?

Table of Contents

About the Author	ii
Preface	vii
 PART I - THE FAILURE OF COMPLIANCE-DRIVEN SAFETY	
1. The System Performs	1
2. Virtual Safety Systems: Manuals on Shelves, not Behaviours in Practice.....	19
3. The Myth of ‘Safety as Priority’	31
 PART II - THE ARCHITECTURE OF SAFETY CULTURE	
4. What Culture Really Is: Artefacts, Values, Assumptions	41
5. Why Culture Fails: Anxiety, Identity, and Misalignment	53
6. The Condition Beneath the Culture: Psychological Safety	63
 PART III - LEADERSHIP AS THE ENGINE OF TRANSFORMATION	
7. The Anatomy of Weak and Strong Cultures.....	75
8. Why Safety Leadership Matters More Than Safety Management.....	85
9. The Full Range Leadership Model® and Its Safety Implications.....	95
10. The Transformational Safety® Leadership Constructs	105

PART IV - TRANSFORMING THE SAFETY WORLD

11. Leadership, Engagement, and Discretionary Effort.....	117
12. What Safe Organisations Actually Look Like: High Reliability Organising	125
13. The Subjective and Objective Worlds of Safety	143
14. The Communication Gap: Why Safety Messages Don't Land	157

PART V - THE PATH TO OPTIMAL SAFETY CULTURE

15. Belief, Zero Harm, and the Psychology of Credibility.....	165
16. The Transformational Safety® Integrated Model.....	175
17. The Maturity Journey: From Lethal to Transformational.....	193
18. Reading the Organisation: Assessing Safety Culture and Leadership.....	205
19. The Transformational Safety® System	217
20. Moving the Organisation: From Assessment to Transformation.....	229
21. Beyond Compliance: A New Philosophy of Safety	241
Annotated Bibliography	251

CHAPTER ONE

The System Performs

The street was quiet in the way that streets are only quiet in the hours before a suburb begins to wake itself. Dawn had not yet committed to the approaching day. The light sat grey and shimmering over rooftops still beaded with the condensation that only seems to appear in those chilly early hours. The only movement was a pair of magpies working their way across a nature strip two houses down. They had parked the car a block away, without quite knowing why? Now the two men walked side by side along the footpath in a subdued silence that felt appropriate to the hour and entirely inadequate to what they were about to do.

The house was unremarkable in the way that most houses are unremarkable: a rendered brick veneer set back from the street behind a garden that was maintained without being cared for, a clothes line visible above the fence of the side yard with clothes hanging limply, awaiting attention, and a light over the garage still on from the night before. A red bicycle lay on the front lawn at the angle of things dropped rather than placed, one pedal pressing into the grass, the front wheel turned as though caught mid-corner. There were stickers along the frame — cartoon figures, mostly peeled, the remains of last year's infatuation. A helmet sat beside it, chin-strap dangling.

The two men had not done this before. They knew enough to know that there was no version of this that went well, and no preparation that made any difference to what happened when the front door opened. The senior of the two, the operations manager, though his title felt like a different language now, had rehearsed nothing. He had been told, years before, that rehearsed words came out like rehearsed words, and that was worse than silence. The other man, younger, the safety manager, kept his hands in his jacket pockets and watched his own feet, step after painful step, on the path. He was there, and not there.

They reached the front step. The sensor light above the door came on and flooded the entrance in a flat white that seemed to announce their arrival to the whole street. On the mat beside the door sat a pair of school shoes — black, laces loosened but still tied, set neatly side by side with a precision that children produce only under instruction, often reluctantly. The sight of them stopped something in the operations manager's chest. Quietly, without drama, he pressed the doorbell before he could think about it further.

From inside came the sounds of a morning in motion: the particular acoustic geography of a household that has not yet received the news. Somewhere toward the back, a television murmured at low volume, the morning shows, background sound really. A kettle was on the boil, its whistle seeking somebody's attention. Why hadn't it turned itself off? A cupboard door. A muffled voice too indistinct to resolve into words. Normal sounds. Ordinary and irreplaceable, the last of their kind.

Footsteps approached the door. Light, quick, almost skipping. Not hers.

The door opened and a young boy looked up at them — ten years old, perhaps. Still in pyjamas with red bicycles on them, a piece of toast in one hand. He had his father's colouring and something in his expression that had not yet learned to conceal itself. He looked at the two men with the open assessment of a child who knows that something is out of place but has not yet assigned meaning to it. He turned his head, hesitantly, back toward the hallway.

"Mum?"

That one word struck the two men like thunder. There was a pause in the kitchen sounds, and then the footsteps that came were different: adult, unhurried, not yet alarmed.

The two men stood on the step in the white light of lament and waited.

* * *

This book is about what happened before that morning. I have spent more than twenty-five years, across mining operations, petrochemical facilities, construction environments and healthcare systems on six continents, trying to understand it with sufficient precision to make a difference.

So, let me be clear. I shall not be speaking about the incident itself. There are more than enough books trying to find their way through those weeds.

What I will be exploring is what happened in the months and years that preceded it: the decisions that were made and not made; the signals that were sent, the signals that were received, and the signals that went unanswered; the habits of thought that persuaded an organisation it was safe when it clearly was not, and the habits of leadership that made those habits of thought possible. Indeed, inevitable.

The walk up that footpath, toward that front door, is among the unspoken nightmares that so many leaders dread. What is so unbearably sad is that, for the vast majority, it takes someone to die before they might consider acting upon that nightmare.

It is important to appreciate that this family did not lose a father to an act of random industrial violence. They lost him to a system — a safety system, a management system, a leadership system — that had been functioning precisely as it had been designed to function, and precisely as it had been led. The most uncomfortable truth this book will ask you to sit with is this:

Every number that organisation produced told a story of safety, and the story was false. Not because anyone falsified the numbers, but because the organisation was measuring the wrong things, in the wrong way, for the wrong reasons, and no one in a position to know differently chose to ask.

..... there was nothing unusual about the organisation. It looked, by almost every measure available to it, like a well-managed operation. Its documents were current. Its training was delivered. Its statistics, reviewed at the monthly executive meeting, consistently told a story of steady improvement.

The Illusion of Safety

Let me begin with the observation that has shaped everything I have come to understand about safety culture. Almost every serious industrial accident in the modern era has been preceded by a period in which the organisation believed it was safe. This is not a coincidence. It is a structural feature of how safety is understood and managed in most organisations, and until that feature is confronted honestly, the understanding cannot change.

The belief in safety — the board-endorsed, operationally embedded belief that an organisation's risks are under control — is almost always the product of something other than evidence of safe conditions. It is the product of documentation, of training completion rates, of incident frequencies that have trended downward over the past three reporting periods, of audits passed, certifications maintained and corrective actions closed, and of a safety management system that is comprehensive in its structure, current in its revisions and largely disconnected from the physical reality of the work being done.

This is the illusion of safety: not deception, but displacement. The organisation has substituted the administration of safety for the reality of it, and having done so convincingly, to itself as much as to anyone else, it can no longer see the difference between the two.

The substitution is rarely deliberate. It accumulates over time. Organisations begin with good intent: a commitment to keeping people safe, translated into systems, procedures and oversight mechanisms. Those systems are necessary, and nobody who works in these industries should mistake this argument for a case against them. A well-designed safety management system is the essential architecture of safe operations. Nonetheless, it is architecture only. It is not the building as occupied, the rooms as used, the decisions made at two in the morning by a maintenance crew three weeks behind schedule. It is the plan, and it is in that gap between the plan and the reality that people die.

This gap is invisible from the top of an organisation because the information that reaches the top has passed through layer upon layer of a system designed, structurally rather than maliciously, to produce reassurance. Numbers are aggregated. Incidents are categorised. Near-misses are counted when they are reported, which means they are systematically undercounted by an amount that is never known but is always significant. The executive receives a dashboard that tells the story of how well the system is working, compiled by people who know, with or without being told, that the executive prefers evidence of progress. This is not conspiracy. It is culture, and it is almost universal.

James Reason¹, whose work on organisational accidents remains among the most intellectually honest accounts of how safe systems fail, described this in terms of what he called latent conditions: failures built into the system long before the moment of incident, invisible until the moment of contact with

sharp-end reality. His Swiss Cheese Model — multiple defensive layers, each perforated, aligning to allow a trajectory of harm — has become so widely used that its central insight is now frequently misunderstood. The insight is not about the holes. It is about who makes them, when, and why the organisation cannot see them accumulating. The holes exist because the system is maintained in an environment of competing pressures, inadequate information, and leadership that has confused the appearance of safety with its reality.

A board reviewing incident statistics that show a fifty per cent reduction in recordable injuries over five years concludes that the safety programme is working, and working fifty per cent better than it was. The conclusion may be reasonable given the data available to it. It may also be entirely wrong. A reduction in recorded injuries is a reduction in recorded injuries. It says nothing reliable about whether the conditions that produce serious harm have improved, stabilised or worsened. It says something about what is being reported, under what definitions, in a culture shaped by what people understand to be acceptable to report. That culture is shaped, in turn, by the behaviour of leaders.

I want to be clear that this is not an historical curiosity. It is the operating condition of a significant proportion of organisations right now. Their executives believe what the numbers tell them because they have been given no compelling reason to believe otherwise, and because believing otherwise would require them to ask questions that the system, as currently constituted, is not designed to answer.

..... The Deepwater Horizon disaster did not happen to an organisation that was completely indifferent to safety. It happened to one that had received safety awards, maintained an elaborate management system, and genuinely believed — at corporate level — that its operations were under control. They were not.

I often cite one detail when I work with executive teams, because nothing illustrates the disconnect more sharply. On the day the Deepwater Horizon exploded, killing eleven people, several BP executives had flown out to the rig to present the crew with a safety award². Within hours of that presentation,

the explosions occurred. It has been widely cited³, rightly, as emblematic of the divide between BP's safety performance on paper and what was happening operationally.

The investigations that followed found a pattern of decisions, pressures and normalised deviations stretching back years, each individually defensible, collectively catastrophic⁴. The organisation had come to manage risk information rather than risk itself.

What makes this more compelling still is that it was not the first time. This was not BP's first visit to the chaos of disaster. On 23 March 2005, only five years earlier, the Texas City refinery explosion had killed fifteen workers and injured one hundred and eighty others⁵.

The question must be asked: did BP learn anything from Texas City? My conclusion is: not much.

From a safety culture perspective, the two disasters share several disturbing structural failure patterns.

In both cases, warning signs had been present for extended periods and had been progressively accepted as normal. At Texas City, the raffinate splitter had experienced similar near-miss overfilling events before it catastrophically failed. At Deepwater Horizon, anomalous pressure readings during the negative pressure test on 20 April were misread and dismissed. It has also been noted that Dewey Revette quite forcefully expressed his concern about immediate and developing process safety failures; only to be ignored by senior leadership⁶. In the case of both Texas City and Deepwater Horizon, the organisation had lost the capacity to recognise what it was seeing, because deviation had become the routine standard.

The Baker Panel Report⁷, following Texas City, found that BP had systematically underfunded process safety while maintaining a strong focus on personal safety metrics such as lost time injuries, slips and trips. The organisation looked safe on paper while the underlying process risks deteriorated unseen. In the period before the explosion, the refinery's conventional safety statistics, its lost-time and recordable injury rates, had been acceptable by the measures the industry used to judge itself. The numbers were not wrong. They were simply measuring something other than the risk that killed fifteen people.

At Deepwater Horizon, the well was forty-three days behind schedule and significantly over budget. Decision after decision, including the choice to

remove drilling mud before setting the final cement plug, reflected the weight of schedule and cost pressure on operational judgement.

Neither disaster can be attributed solely to frontline error, although BP tried to sell it that way. Both the Baker Panel⁷ and the National Commission on the BP Deepwater Horizon Oil Spill⁸ found systemic leadership failures reaching well into the corporate hierarchy. Leaders were not asking the right questions, were not present in ways that mattered, and had not built environments where workers felt safe to raise concerns. Sound familiar?

Both disasters also exposed the inadequacy of the regulatory frameworks governing them. OSHA's process safety management standard had been in place since 1992 but was poorly enforced in petroleum refining. The Minerals Management Service, which oversaw offshore drilling, had developed an excessively cooperative relationship with the industry it regulated, and lacked both the technical capacity and the cultural independence to provide meaningful oversight.

What connects them most profoundly is the gap between declared safety values and enacted safety behaviour: the organisation saying one thing and doing another, at almost every level.

The Gap Between Design and Reality

Here is the thing about safety systems that took me many years of experience to fully appreciate. Every safety system is only a theory. It is a set of assumptions about how work is done — the sequences, the decisions, the conditions under which people operate — translated into procedures, controls and verification mechanisms. The theory is constructed by people who understand the work in principle and may or may not understand it in practice: at the pace, under the pressures and with the resources that actual workers have on an actual Friday when two people have called in sick and the client is on the phone about Monday morning's delivery.

Erik Hollnagel⁹, one of the most important thinkers in contemporary safety science, draws a distinction that is deceptively simple and profoundly consequential: the difference between Work-as-Imagined and Work-as-Done. Work-as-Imagined is fixed at the moment of design; Work-as-Done responds continuously to the variability of real conditions. Workers adapt — they have no choice but to adapt — and the gap between the two is not negligence. It is a normal feature of every organisation operating in conditions its system

cannot fully predict. The administrative structures built around that gap, the audit cycles, the corrective action registers and the management reviews, ask whether the right activities have been completed, not whether the gap itself is narrowing. The question is whether leadership makes the gap visible or leaves it to accumulate in silence.

This matters because the gap is where risk lives. The accidents that destroy organisations and end lives almost never happen in the parts of the system that are visible and controlled. They happen in the invisible space between procedure and practice, between what the system expects and what the work requires, between the version of operations that leadership sees and the version that frontline workers navigate, seemingly successfully, every day.

This is what makes the illusion of safety so dangerous. The system most in need of challenge is the system least likely to generate the information required to challenge it. Safety management systems, by their nature, pass upward the information that confirms their own effectiveness. They are designed to detect deviations and correct them. What they are not designed to do, and what requires leadership specifically to counteract, is detect the ways in which the system's own assumptions are simply wrong.

Adaptation Becomes the System

The concept I find myself returning to most often in my consulting work is one that Diane Vaughan¹⁰ introduced through her analysis of the NASA Challenger disaster that occurred on January 28th 1986. The normalisation of deviance. The engineers' concerns had been documented and raised through formal channels, and the organisation absorbed them. The gap between what the engineers knew and what the launch decision reflected was not ignorance. It was culture. The phrase, normalisation of deviance, is now sufficiently well known that it risks becoming cliché, stripped of its analytical force by overuse. What Vaughan was describing was a specific social process: the gradual acceptance, within an organisation's decision-making culture, of conditions and practices that would, examined in isolation, not be recognised as departures from safety standards. The normalisation is not a dramatic event. It is a slow accumulation of small acceptances, each individually justified, each quietly raising the floor for what comes next.

The process works because each step of adaptation is preceded by a step that was adapted and survived. Workers do not begin their careers making

reckless decisions. They begin by doing the job as they have been trained to do it, in conditions that are rarely exactly as training described, and they adapt intelligently to the difference. That adaptation works. The job gets done, nobody is hurt, the outcome is acceptable. The adaptation becomes part of how the job is now done. A second check is dropped because the first one has never found a problem. A step is shortened because the equipment has always behaved. New workers learn it not as a deviation but as practice. The procedure manual says one thing; the person who has been on site for seven years shows you how it is actually done. How it is actually done works, most of the time, well enough.

What has shifted, and what takes years of professional experience to recognise, is the distribution of that risk. Safe enough, most of the time, is not the same as safe. Each adaptation uses up a little of the margin between normal work and a serious event, and no one is keeping count of how much of that margin has been used. The gap between what the system describes and what the work requires does not remain constant as conditions change. It widens. It widens when production pressure increases, when maintenance is deferred, when staffing is reduced, when equipment ages, and when experience leaves the organisation and is replaced by people who inherit the adapted practice without inheriting the judgement that originally constrained it. The margin that experienced workers maintain — the intuitive, inarticulate sense of how far is too far — is a product of experience. It does not transfer automatically, and organisations that treat it as though it does are operating on borrowed time.

This is the mechanism by which cultural drift operates, and it is worth being clear about what it is. Cultural drift is not the deterioration of a culture that was once strong. It is the slow rotation of a culture's interpretive frame: the gradual shift in what is considered acceptable, normal and unremarkable, driven not by decisions but by the accumulation of outcomes. If conditions that should be alarming produce outcomes that are acceptable, the conditions stop being alarming. The culture does not decide to accept them. It simply stops seeing them as unacceptable, because the evidence available to it — continued operation, statistical safety, the absence of harm — does not support alarm. No single day looks different from the one before it. Only a comparison across years would show how far the culture has moved, and nobody inside it is making that comparison.

The leadership dimension of this process is the one most consistently misunderstood. Leaders are not, in most cases, actively promoting unsafe

conditions. They are, in almost all cases, responding to the same information environment as their organisations: an environment in which the absence of harm is read as the presence of safety, and in which the gap between Work-as-Imagined and Work-as-Done stays invisible because neither the systems nor the culture are structured to make it visible. Leadership passivity in the face of cultural drift is not neutral. It is a signal, and it is received as one by almost everyone in the workplace.

When a senior leader visits the worksite and does not ask the questions that would expose the gap, workers observe the absence of those questions. The questions are not complicated. What are you working around at the moment? What would you change if it were up to you? What does the procedure say that never actually happens?

When a safety concern is raised in a meeting and addressed administratively — a corrective action opened, a deadline set, a box ticked — without the kind of inquiry that would establish whether the concern reflects a systemic condition or an isolated event, the organisation learns something. It learns that concerns should be raised in formats that can be processed and closed, not in forms that require real engagement. The silence that follows is not alignment. It is the sound of a workforce that has learned what kind of information is welcome and what kind is not, and has organised its communication accordingly.

A Wednesday Afternoon

I want you to consider a worker I have encountered, in different forms, on almost every site I have worked on. Fifteen years in the industry, most of them on the same site. He is not reckless, and never has been. He has three children. He has watched colleagues get injured, attended site memorial services, and completed more safety training than he can accurately count. He is, by any reasonable assessment, a safety-conscious professional.

On a Wednesday afternoon in the third week of a scheduled maintenance shutdown, he is working on a pressure vessel isolation. The procedure calls for a specific sequence of valve operations, followed by a pressure check, followed by sign-off before work begins. The sequence is correct. The procedure was written by engineers who understood the process chemistry, and designed the isolation to protect against the specific hazards present. The procedure is right. No question about it.

The shutdown is running four days behind schedule. The client has been on the phone twice this week. The production manager, not directly involved in the maintenance work but visibly present on site and visibly watching the schedule, has not said anything to the technician or his team. He has not needed to. His presence is sufficient. The technician has worked in this industry long enough to read the ambient pressure of a site that is behind schedule, and to understand what that pressure means for the pace of work.

The pressure check will take forty minutes. There is a point in the isolation sequence at which an experienced technician, working in conditions that look right, on vessels that have been depressurised and locked out to a standard that has always worked before, makes a judgement about whether the check will tell him something he does not already know. Most of the time, in most shutdowns, it does not. Most of the time the experienced technician's judgement is correct, and nothing goes wrong. The procedure was designed for conditions that, on most days, the experienced technician can verify by other means: by reading the instruments, by knowing the system's history, by the particular combination of experience and pattern recognition that fifteen years of this work produces.

The technician does not decide to skip the pressure check. He decides that the conditions are consistent with what he has seen many times before, that the risk is manageable, and that the additional forty minutes will not serve any purpose that his experience has not already served. He has made this decision before. Nothing bad has happened. His supervisor, who has also seen this done, says nothing. The production manager, checking the schedule at the end of the shift, sees that the job is back on track. No one has told the technician that what he did was acceptable. No one has told him it was not. The system has received the outcome it required and has nothing further to say. The entire process was invisible.

This is not a portrait of negligence. It is a portrait of rational adaptation: intelligent, experience-based decision-making in an environment shaped entirely by the signals of the organisation around it. The technician is not ignoring the safety system. He is responding to the full system — the production system, the cultural system, the leadership system — of which the safety procedure is only one element, and not, in the moment of that decision, the most powerful one.

The question any serious organisational analysis must ask is not why the technician made that decision. The answer to that is obvious: he made it

because the system in which he operates gave him every reason to make it and no compelling reason not to. The question worth sitting with is why the system produced that environment. The answer is leadership. Not leadership as an individual failure, not the production manager who should have said something or the supervisor who looked away, but leadership as a systemic pattern.

It is the pattern of an organisation whose leaders have communicated, through their presence and their absence, their questions and their silences, their celebrations and their corrections, that production schedule and safety compliance are both valued but not equally weighted when the pressure rises.

It is the pattern of an organisation that measures what can be measured easily and has, over time, come to confuse ease of measurement with importance of outcome.

This brings me to the question I am asked more than almost any other. If the system is producing good numbers, why isn't it working?

That question takes us to the measurement problem that underpins everything else in this chapter. The metrics most organisations use to manage safety are built around events: injuries, incidents, near-misses, audit findings, training completions. These are lag indicators, which measure outcomes that have already occurred, or activity indicators, which measure what has been done rather than what conditions currently exist. Both share a structural limitation that is rarely acknowledged candidly in safety management literature: they can only tell you about the past. They tell you what has happened and what has been administered. They do not, and cannot, tell you about the state of the system right now, at the leading edge of operational risk. They cannot tell you whether the conditions for a serious event are accumulating, whether the gap between Work-as-Imagined and Work-as-Done is widening, or whether the adaptive practices that have developed at the worksite are approaching the edge of the margin that experience can still hold.

The distinction between lag and lead indicators has been widely discussed in safety management for well over two decades without substantially changing organisational practice¹¹. The reason is not that organisations refuse to understand it. It is that lead indicators — measures of the actual condition of risk controls, the actual state of safety-critical systems, the actual degree of alignment between work as designed and work as done — are hard to generate and harder still to interpret. They require the kind of close engagement with operational reality that audit systems cannot provide and that most executive

leadership does not attempt. So organisations continue to measure what is easy: injuries per million hours worked, the percentage of corrective actions closed on time, training compliance by department. They draw comfort from the trends those measurements produce, even as the conditions that produce incidents quietly evolve beneath them, in those invisible spaces.

There is a particular quality of organisational decision-making that becomes visible when you examine the pre-incident history of major accidents. It is not stupidity, and it is not recklessness. It is a kind of confident miscalibration: a state in which the leadership team's model of what is happening in the organisation is detailed enough and consistent enough to feel accurate, while being systematically wrong about the things that matter most. The model is consistent because it has been constructed from consistent data. The data is consistent because it has been generated by systems organised around the production of consistency. The system does not lie. It reports what it has been designed to report, and what it has been designed to report is largely a function of what the organisation has been willing to look at.

The Return to Stability

I have watched this cycle play out more times than I care to count. In the aftermath of serious industrial accidents, there is a characteristic organisational response.

The first phase is shock: the recognition, often at great personal cost to those involved, that something the organisation believed about itself was wrong. Plans are suspended. Reviews are commissioned. Leaders who have spent months explaining why their safety metrics demonstrate a positive trend are now compelled to explain why those metrics are insufficient guides to actual safety performance. New initiatives are announced. Resources are committed. The corrective action register is opened and rapidly populated, and already, regrettably, the cycle that brought them here is starting over.

The second phase is stabilisation. The immediate crisis recedes. The organisation returns to operation. The investigation report is received, and its recommendations are addressed with varying degrees of thoroughness and commitment. Some recommendations change things substantially. Many produce documentation — new procedures, updated training packages, revised risk assessments — that represents an administrative response to what

was found without substantially altering the conditions that produced the finding. The system processes the investigation as it processes everything: by converting it into actions that can be tracked, assigned and closed.

The third phase, which typically arrives between eighteen months and three years after a serious event, is the gradual return of the confidence that existed before it. Incident statistics have improved. They almost always improve in the immediate aftermath of a serious event, for reasons that are partly cultural and partly the result of improved conditions. The corrective actions are closed. The investigation's recommendations have been addressed, in the sense that something has been done in response to each of them. New leaders have arrived who did not live through the event and for whom its details are historical rather than visceral. The organisation begins to believe, again, that it is safe.



This cycle is not inevitable. It is not a law of organisational physics. It is the natural outcome of what happens when the response to an accident is administered by the same thinking that produced it. The same leaders review

the same kinds of findings, approve the same kinds of actions, and measure success by the same indicators that failed to warn them the first time. The report is tabled, the actions are assigned owners and due dates, the tracking spreadsheet turns from red to green, and the board is told the matter is closed. If the fundamental relationship between leadership behaviour and safety system effectiveness is not examined and changed — if the investigation identifies the immediate and contributing causes without interrogating the cultural and leadership conditions that sustained them — then the corrective actions, however thorough, are applied to a system that remains structurally oriented toward the same outcomes. A new procedure is written for a workplace that never followed the old one. A refresher course is delivered to people who already knew the rule and had their own reasons for working around it. So yes, an enormous amount of organisational and personal introspection has taken place in a sincere attempt to make things better. Sadly, all it has really done is recreate the mirage that existed in the first place, albeit with prettier pictures in the corporate publications.

The organisations that actually change after serious events are those in which leadership takes the investigation as the starting point for a reckoning with what their own behaviour has produced, rather than the end point of an obligation. That reckoning is not comfortable, and it should not be. It involves sitting with the understanding that the safety briefings that were delivered, the KPIs that were tracked, the recognition that was given for safe performance and the signals that were sent, with the best of intentions and in good faith, were insufficient to the task they were supposed to serve. Not because the people were inadequate, but because the model of safety leadership they were operating from was not equal to the complexity of what they were being asked to lead. It means asking questions that rarely appear in an investigation's terms of reference. What did we reward that we should not have? What were we told that we did not hear? Who stopped telling us things, and when? And why were we not looking where the danger actually was? These questions cannot be handed to the safety department. They have to be answered by the people whose behaviour shaped the conditions.

Leaders who manage safety treat it as a domain to be administered: the system runs, the obligations are met, the numbers move in the right direction, and they conclude it is working. As we have seen, that conclusion may be entirely wrong. The audit can be clean, the injury rate can be falling, and the next serious event can already be taking shape in the gap between how the

work is described and how it is actually done. The gap between administered performance and the lived safety reality of the workforce is the territory this book is concerned with. It is territory that management, however competent, cannot see, and that only

A Leadership Decision

I want to return, for a moment, to the street. The grey light, the magpies, the bicycle on the lawn, the school shoes set neatly beside the door. The two men on the step just about to knock on that door.

The outcome they were there to deliver was not industrial fate. It was the foreseeable conclusion of a system constructed with care, maintained with diligence and managed with good intent, and which had, for each of those reasons, remained entirely intact as it drifted, year by year, toward the one condition its defences were no longer adequate to meet. Inside, the family were going about the ordinary business of a morning: the kettle, the television, the toast. These were the final seconds of a version of their lives made possible, and then impossible, not by any single decision but by the accumulated decisions of an organisation and its leaders over time.

On the day this father, husband, brother and son died, the training had been delivered. The procedures were current. The permits had been issued and the risk assessments completed. The audit cycle had run on schedule. They even had an ISO 45001 stamp on the company letterhead, telling the world how good their safety system was. Every measure the organisation possessed told it that it was doing what was required. In a narrow, administrative sense, it was.

But the measure of a safety system is not whether it is administered correctly. It is whether the people who work within it go home, physically and psychologically intact. That man was not failed by a single leader or a single decision. He was failed by a model — one that treated compliance as the destination rather than the floor, measured activity rather than conditions, mistook the absence of recorded harm for the presence of safety, and allowed the gap between how work was designed and how was done to grow too wide for experience and judgement to bridge on a Wednesday afternoon in the third week of a shutdown.

That model is not unusual. It is the operating model of most organisations in most industries, and it gets them through most days without incident. It also

produces, with a regularity that should be confronting but has instead become normalised, the kind of morning that two men walk toward in the grey light.

Whether that model needs to change has been answered by a hundred accidents in a hundred industries. The real question is what replaces it. If safety management is necessary but insufficient, and it is, and if the gap between system and reality is a function of leadership rather than administration, then accountability for real safety sits not in the safety management system but in the behaviour of leaders. Every leader, at every level, in every interaction with the operational reality of their organisation.

Safety, then, is not a compliance obligation, though it encompasses compliance. It is not a system to be managed, though it requires robust systems. It is a decision, made daily by leaders, about how closely they are willing to look at the work being done under their authority.

The system performs exactly as it is led. Not as the leader intends it to perform. As it is led — in the daily, often invisible ways that leadership behaviour shapes the environment in which every other decision is made. When leaders accept that, not as a theoretical proposition but as a personal accountability, the model changes, and everything else changes with it.

What that leadership looks like — not as it appears in competency frameworks, but as it operates where real work is done under real pressure — is what the rest of this book sets out to illustrate.

PLEASE RESPECT COPYRIGHT

CHAPTER TWO

Virtual Safety Systems: Manuals on Shelves, not Behaviours in Practice

The email arrived at 8:47 on a Monday morning. The external safety audit was scheduled for Thursday. By Tuesday afternoon the site had produced three updated procedures, closed eleven corrective actions (several of which had been open for more than six months) and conducted an emergency review of the incident register to ensure that every entry was complete, categorised and reconciled with the management system's coding structure.

The operations manager had not requested any of this. The safety manager had not instructed it. It happened because it always happened in the week before an external audit, in this organisation and in a thousand organisations like it, because everyone on site understood, without being told and without discussion, what an audit was for and how to respond to its approach. The understanding was not cynical. It was not resistance. It was a fluent organisational response to a familiar stimulus: an external observer who would assess the organisation not by watching how work was done but by examining whether the right documents existed and the right activities had been recorded. In other words, whether the right boxes had been ticked.

On Thursday, the auditor spent seven hours on site. They reviewed the safety management system documentation, inspected the training matrix, examined the corrective action register, interviewed five supervisors, and walked the main production area with the safety manager for approximately forty-five minutes. They found the organisation substantially compliant. The report, received three weeks later, noted seven minor nonconformances and two opportunities for improvement. It did not note the pressure vessel that

had been tagged out of service six weeks prior and whose return to service had been deferred twice because the shutdown crew was occupied elsewhere. It did not note the three near-misses discussed at the morning toolbox meeting in the fortnight before the audit but never formally reported, because the crew had assessed them as minor and the reporting system required a computer login that most of them could not readily access from the field. It did not note the unspoken understanding that had developed over eighteen months between the day shift supervisor and his team about which parts of the energy control (lockout/tagout) procedure were actually necessary and which were, in the words of one of the team, “for the clipboard”.

The auditor’s report described an organisation with a functional safety management system. It was an accurate description. The system was functional. It occupied twelve folders on a shelf in the safety manager’s office, it was updated on a quarterly schedule, it had been certified to the relevant standard, and it reliably generated the documentation an external observer required to confirm its existence. What it did not do, and what it had not been designed, implemented or resourced to do, was change the decisions people made when the documents were not in front of them, which was most of the time.

Accident investigations describe this phenomenon with remarkable consistency: a safety management system that exists in full administrative detail, entirely separate from the work it is supposed to govern¹². Not a fraudulent system. Not an ignored one. A virtual one: present in form, absent in substance, maintained with real effort and disconnected from the conditions it claims to address.

The Architecture of Reassurance

Safety management systems were not designed to be virtual. They were designed to do something real: to identify hazards, assess risks, implement controls, and verify that those controls are working. The logic is sound. The architecture, translated carefully from engineering principles into organisational structures, is coherent. At its best, a safety management system is the materialisation of an organisation’s understanding of its own risk landscape: a living record of what the organisation knows about how harm can happen and what it has decided to do about it.

Beyond Compliance

The difficulty is not the design. The difficulty is what happens to any designed system when it encounters the complexity of organisational life. Organisations are not engineering problems. They are social systems: collections of people with different priorities, different understandings and different relationships to authority and to risk, operating within structures that evolved over time in response to pressures that were not always rational. Those structures have accumulated layers of informal understanding, tacit practice and unspoken rules that exist entirely independently of any formal document.

Into this social complexity, organisations introduce the safety management system. The system is formal, logical and well structured. It assigns responsibilities, establishes processes and creates verification mechanisms. It is introduced with real commitment, and in the immediate aftermath of its introduction, it changes things. New procedures are followed. New reporting channels are used. People do things differently, at least for a time, because there is organisational attention on whether the new system is working, and that attention is sufficient, temporarily, to shift behaviour.

It is what happens next that I find most revealing about how organisations actually work. What happens to a system when the attention moves on?

Organisations, like all complex adaptive systems, have a powerful tendency toward homeostasis: toward the restoration of the patterns and practices established before any given intervention¹³. This tendency is not resistance. It is not obstruction. It is the natural operation of a social system that has developed efficient ways of doing its work and that will, in the absence of sustained counterpressure, revert to those ways because they are efficient and familiar, and because the new ways require effort the system has not been given the resources to sustain.

Safety management systems are not abandoned. They become virtual because the implementation energy that accompanied their introduction dissipates. Organisational attention moves to other problems, and the system is left to operate in the social environment that preceded it, an environment shaped by different priorities, different habits, and a leadership culture that has not fundamentally changed its relationship to operational risk. The folders remain on the shelf. The certificates remain on the wall. The quarterly updates continue. The work continues to be done in the ways that experience, pressure and culture have shaped it, which may bear a close resemblance to the

documented procedures or may not, depending on factors the safety management system was never designed to detect.

..... major accidents are frequently traced to failures in safety management systems — investigations reveal systems that are little more than sets of manuals occupying metres of shelf space and bearing little relation to what goes on in the workplace.

In more than two decades of reviewing post-incident investigations, I have not read a single major inquiry report that did not arrive at some version of this finding. The system was present. The practice was different. Nobody had noticed the gap until it was too late. When an investigation concludes that an organisation's safety management system was not effectively implemented, this is what it means: not that the system was absent, but that its presence on paper had not translated into changed practice, and that no one in a position to detect the gap had chosen, or been equipped, to do so. The distance between the paper and the practice was the space in which the conditions for a serious event had been developing, unmonitored and unaddressed, for as long as it took.

Culture Beneath the System

Edgar Schein¹⁴, whose analysis of organisational culture remains the most intellectually rigorous framework available, identified three distinct layers of organisational culture — artefacts, espoused values, and basic underlying assumptions — that together explain why a safety management system can be fully documented, regularly audited, and genuinely present while remaining entirely disconnected from the behaviour it is supposed to govern. The system is an artefact of the culture. It exists at the surface — visible, assessable, and maintainable independent of whether the culture beneath it is genuine. The espoused values that surround it — the stated commitments, the declared priorities, the leadership messages about what the organisation stands for — exist at the middle layer, sincerely held by the leaders who articulate them and systematically inconsistent, in most organisations, with the enacted values that actual behaviour reveals. Beneath both of these, invisible to any diagnostic

approach that relies on what can be seen or read or counted, sit the basic underlying assumptions — the taken-for-granted beliefs about why safety systems exist, who owns them, and what they are actually for. These assumptions are the layer that determines everything. The underlying assumptions always win — not because they are stronger than the stated values, but because they are older, deeper, and entirely invisible to the systems designed to assess whether the culture is working. The deeper architecture of those layers, and the diagnostic tools for reading them with precision, is examined in detail within Chapter Four.



Understanding why assumptions win — and why organisations consistently invest in the layers above them while leaving the layer that matters most unexamined — requires a clear account of the specific incentives that make documentation more attractive than risk management.

Why Organisations Choose Paper Over Practice

I am regularly asked why intelligent, well-intentioned organisations consistently choose documentation over real risk management. The answer is not comfortable, but it is straightforward. The preference is not irrational. It is a rational response to a set of incentives that, examined carefully, consistently reward the appearance of safety over its substance.

The first incentive is regulatory. Safety regulation, with a small number of exceptions, is assessed through inspection and audit, processes that are

document-based by necessity, because an inspector cannot be present at the moment of operational decision-making and must instead rely on evidence of what was planned, implemented and recorded. This creates a systematic incentive for organisations to invest in the quality and completeness of their documentation rather than the quality and completeness of their risk management. The organisation with outstanding documentation and poor operational discipline will pass a routine inspection. The organisation with outstanding operational discipline and incomplete documentation will not. Over time, this incentive shapes both organisational investment and organisational culture. Law on paper and law in practice, as observers of international safety regulation have noted, are two very different things¹⁵. You can have highly sophisticated legislative frameworks and almost no translation into operational behaviour, because the enforcement systems are not equipped to see what actually happens in practice.

The second incentive is the one I consider the most insidious, because it operates on good leaders as much as poor ones: real safety management is invisible. The work of engaging regularly with operational reality, understanding how work is actually done, closing the gap between procedure and practice, and creating the conditions in which workers can raise concerns without consequence is time-consuming and ongoing, and its results are invisible by definition. What it prevents cannot be counted. The absence of harm generates no data, no narrative, no report. Documentation, by contrast, is visible, quantifiable and immediately assessable. A corrective action register showing a hundred items opened and ninety-seven closed is a legible story of safety activity. A site visit that reveals how a maintenance supervisor has built his team's understanding of the risk controls on a particular piece of equipment is not legible in the same way. The legible story wins the attention of leaders managing multiple priorities under time pressure. The illegible story, the quiet, continuous work of building safety culture, stays hidden behind a mirage of documentation.

The third incentive is the hardest to raise in a room full of safety professionals, because it implicates all of us, and it has nothing to do with laziness, cynicism or a lack of commitment to safety. It has to do with how the human brain responds to complexity and uncertainty. Completing the documentation feels like doing something. Every procedure updated, every form signed, every audit conducted generates a small but real sense of accomplishment and closure. Real risk management generates discomfort

rather than closure, because it means accepting that the gap between how work is designed and how it is done is unknown, possibly widening, and in need of sustained attention. It requires leaders to sit with uncertainty and invest effort in activities whose results are ambiguous and distant. The psychological preference for closure over ambiguity is well established in behavioural science, and the achievement of closure can reduce the motivation to continue searching for potentially important information¹⁶. Systems that offer the comfort of completion can therefore draw attention away from matters that warrant further thought.

Together, these incentives produce an organisational environment in which the safety management system functions as designed, as a document management system, while the actual management of safety is left to the informal, largely invisible practices of experienced workers and to leadership relationships that may or may not be adequate to the task. In organisations with strong informal safety cultures, carried by experienced workers who have developed sound safety judgement through years of close engagement with the hazards of their work, this arrangement produces acceptable outcomes most of the time. In organisations where that informal culture has weakened, through workforce turnover, the loss of experienced practitioners or the gradual erosion of the margin that experience maintains, the virtual safety management system provides no compensating protection. It is present, and it is irrelevant. Chaos is just around the corner....

The Living System And The Dead One

I can usually tell within the first hour of arriving at a site which kind of system I am dealing with. It is rarely the documentation that gives it away; the documentation is almost always adequate. It is the texture of the place: whether the safety conversation at the worksite has the quality of real engagement or of performed compliance, and whether workers pause before answering a question about how work is actually done or answer without hesitation because the question and the reality are the same thing. It is in small things: whether a supervisor knows the names of the people on the crew, how people talk about the permit system when they think no manager is listening, and what happens when I ask to see the procedure for the job being done in front of me right now. Two organisations can have safety management systems of identical architecture, developed to identical standards and assessed by identical audit

methodologies, and produce radically different outcomes, because what makes a system live is not its documentation but the culture in which it operates and the leadership behaviour that shapes that culture.

I have seen this within a single company. Two sites, the same corporate system, the same procedures, the same audit scores. At the first, I asked a maintenance fitter how a particular isolation was done, and he recited the procedure word for word, a little too quickly. At the second, I asked the same question and the fitter said, “Well, the procedure says this, but we do it slightly differently, and here’s why,” and then explained a sound reason, which his supervisor already knew about and had raised to have the procedure changed. On paper the first site was the more compliant. In practice the second was the one with a living system.

A living safety management system is one in which the procedures reflect how the work is actually done or, where they do not, the gap is known, acknowledged and being actively worked on. A procedure that cannot be followed as written is treated as a problem with the procedure, not as a failing of the worker who found it. It is one in which risk assessments are conducted by the people who do the work, and who have the authority and confidence to describe that work as it is done rather than as it is imagined. Workers can explain why a control exists, not just what it requires. It is one in which near-miss reports are welcomed, investigated and used to improve the system’s understanding of where the gap between design and practice is widening. It is one in which the corrective action register is a record of continuous improvement rather than an administrative mechanism for closing out obligations. And it is one in which leaders at every level engage with operational reality often enough to keep their understanding of the system current, honest and grounded in what they can see and hear for themselves rather than what the dashboard tells them.

A dead safety management system is one in which none of those things are true. It is one in which procedures were written by the safety manager, reviewed by the operations manager and signed off by the general manager without the people who actually do the work being asked whether the procedures describe it accurately. Ask a worker to find the procedure for the task in front of them and it takes ten minutes, a phone call and a search of a shared drive, if it can be found at all. The knowledge that actually keeps people safe lives somewhere else entirely: in experienced heads, private habits and conversations that never reach the system. It is one in which near-miss

Beyond Compliance

reporting is low not because near-misses are infrequent, but because the culture around reporting does not give workers confidence that what they raise will be received without consequence. It is one in which corrective actions are closed by the raising of a further document rather than by verification that the underlying condition has changed. And it is one in which leaders understand safety performance only through the dashboard, with no mechanism for the real state of the system to surface to a level at which it could be seen and acted upon.

The danger of the virtual safety system is not its existence but the confidence it generates. An organisation with no real safety management system knows it has none, and may respond with appropriate concern. An organisation with an extensive, well-documented, regularly audited system that is not embedded in operational practice does not know what it has missed. It believes it has met its safety obligations and built a structure that will protect its workers and its operations. It has built a structure that will satisfy an auditor, and it has mistaken the satisfaction of the auditor for the achievement of the purpose. The audit confirms that the system exists. Nothing in the audit confirms that anyone at the workplace believes in it, uses it, or would turn to it when something starts to go wrong. And the confidence changes behaviour at the top. Leaders who believe the system is working stop looking, at exactly the point where the system most needs them to look. This is why the organisations with the most sophisticated systems are often the most exposed: the better the paperwork, the stronger the confidence. Organisations like this are not surprised by serious events because they were negligent. They are surprised because they were confident, and their confidence rested on the wrong foundation.

..... Safety is known by the absence of events — which makes it a difficult challenge to achieve. In most cases the precursors of serious events were there to see, but the organisations were seemingly unable to recognise them.

What Makes a System Real

The transformation of a virtual safety system into a living one is not primarily a technical exercise. It does not require a better standard, a more sophisticated audit methodology or a more comprehensive document architecture. It requires a change in the cultural assumptions that govern why the system exists, who owns it, and what it is supposed to do. Those assumptions, as Schein spent much of his career demonstrating, cannot be mandated or documented. They can only be shaped by the consistent, visible behaviour of leaders over time.

Making a system real begins with an honest assessment of what the current system actually measures. As Chapter One described, most safety management systems are extraordinarily good at measuring activity and systematically poor at measuring conditions: the actual state of risk controls, the current gap between procedure and practice, and what workers know about their own risk exposure that the system does not. Closing that gap requires leaders to invest in forms of operational intelligence the documentation and review cycle does not generate: direct observation of work as it is done, structured conversations with frontline workers about how they actually manage the hazards they face, and a willingness to receive information that is inconsistent with the story the data currently tells.

It requires that line managers, not just the safety function, own the safety management system in the operational sense: that they understand the procedures governing their areas, can describe the risk controls and explain why they exist, and take personal responsibility for the gap between how work is designed and how it is done. The safety function's role should be advisory and assurance-focused, not operational. It is there to provide technical expertise, measurement frameworks and independent verification that line managers cannot provide for themselves. It is not there to own the system on behalf of the organisation while line managers run their operations in the system's shadow.

It requires that leaders at every level treat the safety management system not as a compliance obligation to be met but as a diagnostic tool: the organisation's best available instrument for understanding where its risk landscape is shifting, where its controls are degrading, and where Work-as-Imagined and Work-as-Done have drifted far enough apart to require intervention. A system used as a diagnostic tool generates different behaviour

Beyond Compliance

from a system used as a compliance demonstration. A diagnostic tool is probed, challenged and tested against operational reality. A compliance demonstration is maintained, audited and presented to external observers. The difference in orientation, subtle at the level of individual interactions and profound at the level of organisational culture, is the difference between a system that lives and one that does not.

Safety management systems do not fail because they are poorly designed. They fail because the organisations that adopt them do not change the cultural assumptions that treat safety management as an administrative function rather than a leadership responsibility. Until those assumptions change — until the organisation's basic, taken-for-granted belief is that safety is owned by leaders rather than administered by specialists, that the system exists to improve real conditions rather than demonstrate due diligence, and that the gap between procedure and practice is a leadership problem rather than a worker compliance problem — the manuals will stay on the shelf, the work will continue as it always has, and catastrophe will never be far away.

PLEASE RESPECT COPYRIGHT

CHAPTER THREE

The Myth of ‘Safety as Priority’

There is a definition of safety culture that appears, in various forms, in almost every major regulatory framework, audit standard and corporate safety policy written since the concept was first formally articulated in the aftermath of the Chernobyl disaster. In the language of the International Nuclear Safety Advisory Group (INSAG)¹⁷ and in the countless derivative formulations that followed, safety culture is that assembly of characteristics and attitudes in organisations and individuals which establishes that, as an overriding priority, safety issues receive the attention warranted by their significance. In its more popular form, safety culture is the corporate atmosphere in which safety is understood to be, and is accepted as, the number one priority.

This definition has shaped four decades of organisational safety thinking. It has been incorporated into regulatory frameworks, embedded in audit criteria, translated into leadership competency assessments, and cited in board presentations as a description of what a safety-committed organisation looks like. It is widely accepted, widely taught, and routinely displayed in the opening slides of safety culture seminars, where senior leaders nod at its familiarity.

I have a real problem with this.

The problem is not with the aspiration. The problem is with the architecture. Safety cannot be a priority, not because the ambition is unworthy, but because the concept of priority is structurally incompatible with the conditions under which safety has to operate. Organisations that have adopted the priority framing have, as a direct and largely unexamined consequence,

created safety cultures that are exactly as safe as circumstances permit. That is to say, not safe enough when circumstances are difficult, which is precisely when safety matters most.

Priorities shift. That is their nature and their purpose. They shift under production pressure and resource constraint, under competitive pressure and regulatory change, and under the accumulated weight of a hundred small decisions, each individually defensible, that together turn an organisation away from the conditions it intended to maintain. In most domains of organisational life this flexibility is useful; a priority that cannot be adjusted is not a priority but a constraint. Safety is not most domains. The organisations that have framed safety as a priority have made it, without realising or intending it, adjustable.

The Priority Illusion

Here is what the priority illusion actually looks like. It does not look like a decision. It looks like a Tuesday. It looks like a production meeting that runs twelve minutes over because the schedule is the problem and safety is the agenda item that everyone already knows the answer to. It looks like a pre-start read off a laminated card in four minutes flat, signed by everyone and understood by no one in particular. It looks like a manager walking past a condition he has seen before and will see again, because the shift is short-staffed and the alternative to walking past it is a conversation nobody has time for today. It looks like a permit signed at the end of a long shift by someone who did not walk the job. It looks like a supervisor who knows the isolation is incomplete, but also knows what will be said if the crane sits idle for another hour. It looks like a job that should have been stopped and wasn't, because stopping it would have been noticed and finishing it would not. It looks like a worker who has learned, not from anything he was told but from everything he has observed, that the organisation's relationship to safety has a price, and that when the pressure is real, the price gets paid.

It looks like a senior leader who visits the site, walks the floor, appears to ask the right questions and leaves without ever reaching the question that mattered, because the right questions, asked fluently and received warmly, produce the appearance of engagement without requiring its depth. It looks like an executive who asks how many days it has been since the last injury, and never asks how many times this month someone was told to keep going. It

looks like a safety moment at the start of a meeting that everyone sits through and nobody remembers by the end of it. It looks like a safety briefing delivered with conviction on a Monday morning and contradicted, not in words but in decisions, by Wednesday afternoon. It looks like a culture that has learned to perform safety for the moments when performance is required, and to manage it, quietly and efficiently, around everything else.

This is the illusion in operation. Not dramatic. Not cynical. Entirely ordinary, and for that reason invisible to the systems designed to detect it. The corrective action register does not record what was not reported. The audit does not capture what was not said. The inspection checklist confirms the guard is fitted; it cannot tell you how often it comes off to clear a blockage. The near-miss that was talked about over the break but never written down does not exist, as far as the system is concerned. Neither does the contractor who was told not to bother with the form. Nobody falsified anything. The system simply recorded what it was given. The incident investigation finds the proximate cause and closes the file. The dashboard stays green while the gap between what the organisation believes about itself and what is actually happening in its operations widens, quietly, in the hidden spaces between the metrics.

The workforce sees all of it. Workers are not passive recipients of the safety culture their leaders intend to build. They are experienced and exact readers of the culture their leaders actually create: through presence and absence, through questions asked and questions avoided, and through the decisions made when safety and something else competed and one of them won. They know which leaders mean it and which perform it. They know what happens to the person who raises a concern at the wrong moment. They know the difference between a stop-work authority that exists on paper and one that has been used without consequence. They know whether a near-miss report will be investigated or administratively closed, and whether the commitment expressed at Monday morning's safety briefing will survive Monday afternoon's production meeting. New starters learn it fastest of all. Within weeks they know which rules are real, because they have watched what happens when someone breaks them. Workers calibrate their behaviour accordingly — what to report, whether to stop work, how closely to follow a procedure the schedule leaves no time to follow fully — with an accuracy that no survey instrument or audit process will ever capture.

In safety, the consequences of that calibration stay invisible until they are catastrophic. Each trade, taken on its own, is small enough to defend. A shortcut taken once becomes the way the job is done, and the way the job is done becomes the standard the next crew inherits. The gap they create accumulates silently across months and years before it announces itself, and the moment of reckoning arrives not as a trend line on a dashboard but as two men walking along a footpath in the grey early light, carrying news that no priority, however sincerely stated, had managed to prevent.

An organisation that has built its safety culture on a priority has built it on something that will move. And it will move, always, at the moment it was most needed to hold.

What Values Actually Are

The distinction between a priority and a value is not semantic. It is structural, and it produces different behaviour under different conditions in ways that are entirely predictable once the distinction is clearly understood.

A priority is an ordering principle. It tells an individual or an organisation which of several competing demands to address first when they cannot all be addressed at once. Priorities are inherently comparative: they exist only in relation to other priorities, and their ranking can, and does, shift as the context shifts. An organisation that prioritises safety above production on a Monday may, under enough production pressure, effectively reverse that ranking by Friday without formally revising its stated commitment. The revision happens not in declared policy but in actual decision-making, because a priority, by its nature, can be overridden by a competing claim that feels more urgent. On a Friday afternoon, even the pull of getting away for the weekend can be enough¹⁸.

A value is a different kind of thing. It is not an ordering principle. It is an identity statement. It describes not what an individual or organisation will choose to do first when choices must be made, but what kind of individual or organisation it is. Values operate at the level of basic assumption, the deepest layer of Schein's model: the taken-for-granted beliefs that shape how the world is interpreted and which decisions are even entertained. A value does not compete with production pressure the way a priority does, because it is not positioned in relation to production pressure. It precedes it. It shapes the frame

within which operational decisions are made, rather than sitting on the decision-making agenda as one item among many^{19,20}.

Organisations that understand this distinction describe it in terms immediately recognisable to anyone who has worked in a safety-committed culture: safety is not something we do; it is something we are. Safety requirements are not barriers to production; they are the conditions within which production is possible. A worker who stops a job because a safety condition is not met is not prioritising safety over production. They are operating from a framework in which a job that cannot be done safely is not a job at all, but an activity their identity will not permit them to undertake, whatever the schedule requires.

The clearest evidence for this comes from the study of High Reliability Organisations: the rare institutions that operate for long periods in conditions of extreme hazard without serious incident²¹. What distinguishes them is not better procedures, though their procedures are often excellent. It is that safety is embedded in their operational identity so deeply that it shapes how they think about what they are doing. The gap between Work-as-Imagined and Work-as-Done is narrower in these organisations not because the work is simpler, but because people at every level are looking for it. Not because they have been told to, but because their professional identity is built on the belief that finding the gap before it becomes catastrophic is what they are there to do. Chapter Twelve examines how these organisations operate in detail.

Building that identity — moving safety from the priority column to the value column — is a leadership challenge most organisations have not yet understood, and it is a different challenge from implementing a better safety management system. It cannot be achieved by a more compelling safety policy, a safety awareness campaign, a zero-harm pledge or an updated leadership competency framework. It requires leaders to model, consistently and visibly, a relationship to safety that is not conditional on operational pressure. And it requires the organisation's response to those leaders, through recognition, accountability, resource allocation and the stories it tells about itself, to reinforce rather than undermine the values they are trying to embed.

When Priority Becomes Pressure

The point at which the priority framing most consequentially fails is not the point of crisis. It is the extended, ordinary middle of organisational life: the

sustained period of operational pressure, resource constraint and competitive tension that characterises most organisations most of the time. It is in this middle ground, not in the dramatic moments of obvious emergency, that safety culture is formed or eroded, and that the practical difference between a safety priority and a safety value becomes visible.

Production pressure is the most pervasive and the most poorly understood influence on safety decision-making in organisations²². It is rarely the crude, explicit pressure of a manager telling a worker to take a shortcut. That version exists, but it is less common and less dangerous than its subtler form: the ambient pressure of an organisation that is behind schedule, under budget constraint or facing competitive threat, communicated not through instruction but through presence, tone, question, and the cumulative signals of what the organisation responds to and what it ignores.

The research on safety climate points in the same direction. When Rhona Flin and her colleagues²³ reviewed the safety climate surveys used across high-hazard industries in 2000, they found the same themes recurring whatever the instrument: management commitment, the competence of supervisors, the priority given to safety over production, and time pressure. Two of those four are about exactly the tension this chapter describes. Workers, asked what shapes safety where they work, keep returning to whether safety holds when production and time push against it.

This is how organisational silence comes to masquerade as alignment. When workers stop raising concerns, when near-miss reporting drops, and when the corrective action register fills with items that have sat open for months, the organisation often reads these as positive signs: conditions are stable and the workforce is engaged and productive. They are nothing of the kind. They are evidence that workers have internalised a set of signals that makes raising a concern feel futile or risky, rather than a set of values that makes it feel natural.

As Chapter Two described, organisations with the most sophisticated safety management systems are often the most exposed to this, because their systems create the appearance of comprehensive coverage. The register is populated, the audits are acceptable, the dashboard is green, and beneath it all workers are still making the same adaptations under the same pressures, guided by cultural signals the system was never designed to detect.

The relationship between production pressure and safety culture failure runs through the major industrial accidents of the past half century. Piper

Alpha²⁴ and Texas City²⁵ both reveal the same pattern: not conscious decisions to sacrifice lives, but the predictable outcome of organisations in which safety was one priority among others rather than a value that preconditioned every operational decision.

Zero Harm, as most organisations deploy it, shows the same problem in its sharpest form. The aspiration to eliminate workplace injury is morally right. But when Zero Harm is adopted as a performance metric, a priority becomes a target, and targets are managed. Chapter Fifteen examines what that does to reporting, to credibility and to belief, and what honest safety leadership does instead.

Safety as a Way of Being

The alternative to safety as a priority is not safety as a higher priority. It is safety as a way of being: embedded in the identity of the organisation and of every leader within it, at a level that precedes prioritisation rather than competing within it. In the terms of the Transformational Safety® Culture Maturity Model examined in Chapter Seventeen, this is what separates the Developing level from the Functional and Transformational levels: not the sophistication of the system, but the depth at which the commitment is held.

Safety requires people to adopt habits of mind and ways of acting that are often counterintuitive and sometimes demanding. Reporting one's own mistakes requires the suppression of a natural self-protective instinct. Stopping a job in the middle of a task, when the schedule is tight and the crew is waiting, requires the courage to withstand social and organisational pressure. Raising a concern about a condition that everyone else appears to accept requires the confidence to trust one's own judgement over the consensus of the group. None of these behaviours come naturally in an environment that has not been deliberately and consistently shaped to support them. They are the particular challenge of leadership: the behaviours that a safety culture built on values must actively cultivate rather than assume will emerge from the implementation of a management system.

The cultivation of these behaviours begins with what leaders model. This is not a metaphor or a motivational observation. It is a description of how cultural norms are transmitted in organisations. Workers learn what is acceptable and what is valued not primarily from policy documents or safety briefings, but from watching what their leaders actually do, particularly when

The Myth of 'Safety as Priority'

doing the right thing is inconvenient or costly. A leader who stops a production activity because a safety condition is not met, at real cost to the schedule, communicates something about organisational values that no policy statement can replicate. A leader who walks past an unsafe condition, noted but not acted on, communicates something equally powerful in the opposite direction. Neither message is verbal, and neither appears in the corrective action register. Both are received, and absorbed into the culture with a permanence that formal communication rarely achieves.

*..... safety requires a deep-rooted set of beliefs, practices and behaviours that guide people's decisions and actions at all times.
Building a strong Transformational Safety® culture is a specific challenge of leadership.*

Of all the things leaders model, the one the workforce reads most closely is how they spend their time²⁶. Time is the scarcest resource a senior leader has, and everyone knows it, so where it goes is the most honest signal of what that leader values. A leader who says safety comes first but fits it around everything else, arriving late to the safety meeting, cutting the site walk short when the day gets busy, rescheduling the conversation about a concern for the third time, is telling their people exactly where safety sits. Workers are answering a simple question, without ever saying it aloud: is safety fitted around this leader's other priorities, or are the other priorities fitted around safety? That is the difference between a priority and a value, measured in hours.

The research on transformational leadership is unambiguous on this point, and Chapters Nine and Ten examine it in detail. Leaders who show individualised consideration, a real concern for the safety and wellbeing of each worker and the behaviour the Transformational Safety® Leadership framework calls the Carer, produce measurably different safety cultures from those who manage safety as a compliance obligation. Leaders who model intellectual stimulation, the Thinker, by asking how work is actually done and challenging the assumptions built into existing procedures, create organisations that are better at finding risk before it becomes an incident. Leaders who show inspirational motivation, the Motivator, by speaking about safety in terms of shared values rather than regulatory obligation, build

workforces more willing to engage with safety beyond the minimum the system requires.

Among the strongest relationships in all of leadership research is the one between Idealised Attributes, the Knight, whose stated safety values are visibly consistent with their actual decisions, and the extra effort followers give. Followers who believe their leaders mean it do not merely comply with safety requirements. They contribute beyond them, in the ways Chapter Eleven examines as safety participation, and that contribution, accumulated across a workforce, is the difference between an operation that is safe and one that is merely compliant.

From Priority to Principle

The journey from safety as a priority to safety as a value is not a communications exercise. It is not achieved by replacing the language of priority with the language of value in the organisation's safety policy, though changing the language is a necessary beginning, because language shapes thought and thought shapes culture. The journey requires leaders at every level to renegotiate their relationship to safety: to understand it not as one of several competing obligations to be managed, but as the foundational principle within which every operational decision is made.

This renegotiation begins, as all real cultural change begins, with an honest reckoning with what the culture currently is. Not what the espoused values say it is. Not what the documentation describes. What the underlying assumptions actually are, as revealed in the daily behaviour of leaders, in the response to near-miss reports, and in the decision made at 4pm on a Friday when the schedule is tight and a condition has been identified that should stop work but probably won't. That reckoning is uncomfortable, and it is meant to be.

When it is honestly undertaken, it produces a particular recognition: that the safety culture the organisation has is the culture its leaders have created, through their behaviour, their decisions and the signals they have sent, in good faith and with real commitment, over the years. This is not an indictment. It is a description of how culture works. And it carries the most useful insight available to any leader who wants to change it: culture is not a legacy that resists leadership but a living system that responds to it, continuously, in the ways Schein's framework predicts.

The Myth of 'Safety as Priority'

Safety as a principle — not a priority, not a target, not a metric, but an orientation that preconditions how the organisation thinks about every aspect of its operations — looks different at every level of the organisation. At board level, it looks like real engagement with the quality of the safety information being received, and healthy scepticism about whether that information tells the full story of what is happening in operations. At executive level, it looks like leaders who engage with operational reality often enough to close the gap between their model of what is happening and what is actually happening. At management level, it looks like the kind of leadership Deming described: not supervision but development, not inspection but coaching, not the measurement of past performance but the creation of conditions for future performance²⁷.

At the frontline, it looks like the thing safety culture research consistently identifies as the most reliable indicator of a healthy safety culture: workers who believe their safety concerns will be heard, acted on, and received without personal consequence. Not because they have been told they will. Because they have seen it happen, consistently, when leaders who could have chosen production over safety chose instead to model the value they say they hold.

What produces that culture is not the safety management system, and not a zero-harm commitment. It is the accumulated effect of the decisions, behaviours and signals of leaders who understand that the culture they create is not incidental to their safety performance but determines it: a culture in which safety is not the highest priority but the ground on which all priorities rest.

If safety is a priority, it will be traded when the pressure is sufficient. If safety is a value, it will not. The gap between those two sentences is the gap between the organisations that protect their people and the organisations that do not.

CHAPTER FOUR

What Culture Really Is: Artefacts, Values, Assumptions

The new general manager spent his first week walking the site. The operations manager and the safety manager accompanied him. The site was clean. Housekeeping standards appeared high. Safety signage was current, clearly legible and appropriately placed at each relevant access point. Workers they passed were wearing their required PPE. At the pre-start meeting they attended, the day's hazards were discussed with a conversational fluency that suggested real familiarity rather than the laboured recitation of a trained performance.

The documentation in the safety office was organised. The training matrix showed no gaps that had not been accounted for. The corrective action register contained nothing older than four weeks without a recorded status update. In the main production area, a board displayed the current safety metrics — days since the last recordable injury, the rate for the preceding quarter, the targets for the year — in a format that was readable, maintained and evidently regarded as important.

By the end of the first week the general manager had formed a view. He told his leadership team at the monthly meeting that the site had a solid, positive safety culture. Not excellent, and there were areas he intended to develop, but solid. A foundation worth building on.

He had spent a week reading artefacts. Artefacts tell a leader what a culture has constructed. They do not tell a leader what a culture believes. The construction can reflect the culture's values or it can exist independently of them. It can be the surface expression of a deeply embedded commitment to safe operations, or it can be the carefully kept facade of an organisation that

learned long ago how to look the part without necessarily being it. Without the capacity to read below the surface, without the diagnostic tools to distinguish the artefact from the assumption it is supposed to represent, a leader cannot tell the difference. And a leader who cannot tell the difference is, in the most consequential sense, managing safety blind.

What Leaders See and What They Miss

The problem with cultural diagnosis in safety is not that leaders lack analytical capability. The problem is structural. The aspects of culture most visible to an external or newly arrived observer are exactly those most open to deliberate management, and therefore most likely to be maintained at a standard that does not reflect the culture's actual health. The aspects of culture that most reliably predict safety outcomes — the assumptions, habits of interpretation and unspoken understandings that govern how people actually think about and respond to risk — cannot be seen, counted or read from a document. Without the capacity to explore the hearts and minds of the people within the culture, a leader has little basis for any defensible conclusion about it. If that is not confronting, I don't know what is.

This structural problem has a specific consequence that safety leaders at every level need to understand: the information an organisation presents to a leader who comes to assess its culture is not random. It is shaped by that organisation's understanding of what leaders are looking for and what they reward. An organisation that has experienced rigorous audit processes will have developed the capacity to perform well in rigorous audit processes. An organisation that has been rewarded, year after year, for demonstrating the visible markers of safety commitment will have developed those markers with care and maintained them reliably. What neither the audit nor the annual recognition process will necessarily have improved is the organisation's actual safety culture: the beliefs, assumptions and habitual ways of interpreting and responding to risk that determine what people do when the auditor has left and the award has been hung in the reception foyer.

The new general manager had no reason to distrust what he had seen. Everything he had seen was real. The housekeeping was maintained. The documentation was current. The pre-start meeting happened. What he had not yet seen, because the tools he used — observation, document review,

conversation with site management — were not designed to reveal it, was the cultural substrate from which those visible practices had either grown or been applied, like a coat of paint, to a surface whose underlying condition was different from its appearance.

Culture is a system of meaning. It is the accumulated, largely tacit understanding that a group of people develop, over time, about how their world works and what is required to function within it. That understanding is not distributed equally or consistently across the organisation. It is transmitted through experience, through observation of what leaders do and what the organisation rewards, through the stories that circulate and the ones that are suppressed, and through the ten thousand small interactions that make up the lived reality of belonging to an organisation. It cannot be transferred by a policy document or captured by a survey alone. It can be described, approximately, by an insider who has developed the capacity for self-reflection. It can be inferred, carefully, by an outsider who knows what to look for. And it can be read, with some accuracy, by a leader who has learned to move deliberately through each of its layers.

Building that capability is the purpose of this chapter. Not as a replacement for the operational instinct and relational intelligence that experienced leaders develop through years of engagement with complex organisations, but as a framework for making explicit what good leaders often do intuitively: the movement from surface observation to deeper inference, from what an organisation shows to what it believes.

The Visible Layer: Artefacts and Their Deceptions

Artefacts are everything you can see, hear or otherwise directly perceive in an organisation. In a safety context, they include the physical environment: the condition of the worksite, the state of equipment, the visibility and currency of safety signage. They include the language people use: the terms that appear in safety conversations, the way risk is described, and whether the vocabulary of the safety system is used with real familiarity or as role-play. They include the observable behaviours of workers and leaders: PPE compliance, participation in toolbox meetings, the way safety concerns are raised and responded to in real time. And they include the visible outputs of

the safety management system: the documents on the shelf, the boards on the wall, the metrics on the dashboard.

Artefacts are the most abundant and most immediately available source of cultural data. They are also, consistently, the most misleading, not because they are false but because they are ambiguous. An artefact tells you that something is present. It does not tell you what that something means to the people who maintain it.

Consider the safety metrics board in the production area, a common artefact in many workplaces, typically displaying the number of days since the last recordable injury alongside various leading indicator targets. Someone decided to install it, someone updates it regularly, and workers see it every day. What does it mean? In one organisation, the board represents a commitment to transparency and continuous improvement: a shared reference point for a team that is invested in the safety trajectory it describes. In another, an identical board is a source of pressure, carrying the unspoken understanding that the number displayed must be kept above a certain level to avoid a conversation with management that no one wants to have. In a third, it is simply part of the furniture: present because it has always been present, updated because the safety manager updates it, noticed by no one in particular.

The board is identical in each case. The cultures are entirely different. No amount of observation of the board will distinguish between them. What distinguishes them is what the board means to the people who work beneath it, and meaning is not an artefact. It is a value or an assumption, and it lives in the layers below.

The same ambiguity applies to every artefact in the safety context. High PPE compliance can indicate a culture of personal safety awareness, an effective enforcement mechanism, or the learned understanding that being caught without PPE is unpleasant. Participation in toolbox meetings can reflect real engagement with the day's risks or the practised skill of appearing engaged while mentally already at the worksite. A corrective action register with no items older than thirty days can indicate rapid and effective response to identified hazards, or a systematic habit of closing actions against documentation rather than against conditions. The artefact is present in every case. The culture it represents varies enormously.

The deception of artefacts is not malicious. Organisations do not, in most cases, construct their visible safety culture with the intent to mislead. They construct it in response to the signals of the environment in which they

operate: the signals of regulators, auditors, clients and internal leadership about what is expected and what is rewarded. When those signals focus on visible compliance, organisations develop excellent visible compliance. When they focus on operational safety, organisations develop both the visible markers and the underlying culture. The difference lies entirely in what the signals reward.

The diagnostic implication for leaders is this: artefacts must be read, but they must be read as questions rather than answers. Every artefact prompts the same inquiry: what belief or assumption does this artefact express, and how would I know whether that belief or assumption is real? Answering that question requires moving to the next layer.

The Middle Layer: Espoused Values and the Gap They Create

Espoused values are what an organisation says it stands for. They are the articulated principles, stated commitments and declared priorities that appear in the safety policy, the values framework, the leadership messages and on the boardroom wall. In most organisations they are sincerely held by the leaders who articulate them. They are not cynical performance or deliberate misdirection, but honest expressions of what those leaders believe the organisation should be and, in many cases, believe it is.

The problem with espoused values is not their sincerity. It is the systematic gap that exists, in virtually every organisation, between what is espoused and what is enacted: between the values the organisation says it holds and the values its actual behaviour reveals. Chris Argyris²⁸, whose work on organisational learning remains one of the most practically useful bodies of research in management science, described this gap as the distinction between espoused theory and theory-in-use: the theory an individual or organisation articulates, versus the theory actually governing their behaviour. The gap is not hypocrisy. It is the natural result of the fact that behaviour is governed by habits and assumptions that formed long before the espoused values were articulated, and that articulating values does not, by itself, change those habits and assumptions.

In a safety context, the gap between espoused and enacted values is one of the most consequential and most commonly underestimated features of organisational culture. An organisation might sincerely espouse the value that

every worker has the right and responsibility to stop unsafe work. Yet if the enacted culture, the theory-in-use revealed by what actually happens when work is stopped, treats stoppages as inconveniences to be minimised and resolved as quickly as possible, then the espoused value is not the operative one. Workers learn this not from being told but from observing what happens when the value is tested. After enough observation, they adjust: they continue to espouse the value when asked, because they know it is the right answer, while their operational decisions are governed by the enacted value, because that is the one with experienced consequences.

The safety culture survey is where this divergence most commonly shows up. Organisations regularly find that workers' stated attitudes to safety are strongly positive — they agree that safety is important, that they would stop unsafe work, that they feel comfortable raising concerns — while observation or incident analysis over the same period reveals behaviour inconsistent with those attitudes. The gap is not lying. It is the coexistence, in the same individual, of a value sincerely held as an aspiration and a different value that governs behaviour under the pressures of operational reality.

Leaders add a further complication. The espoused values of senior leaders carry particular weight in organisational culture: they define the official story the organisation tells about itself, and they shape the communication of everyone below them, who have a rational interest in aligning their visible commitments with those of the people who control their prospects. This means an organisation's espoused safety values can be essentially uniform, with everyone from the board to the frontline articulating the same commitments in the same language, while the enacted values at the operational level are entirely different. The uniformity of espoused values tells a leader that everyone knows what they are supposed to say. It tells them nothing about what people actually believe.

Testing the gap therefore cannot rely on traditional surveys, which measure espoused values by design. It requires close observation of what happens when safety values are tested by operational pressure. When a worker raises a concern that has a production cost, what happens next? When a near-miss is reported that implicates management decision-making, how is it received? When a procedure is found to be inconsistent with how the work is actually done, who is expected to change: the procedure or the practice? The answers, accumulated across many observations, describe the enacted values

of the organisation far more accurately than any number of policy statements, safety pledges or survey responses.

The Invisible Layer: Assumptions That Run Everything

Basic underlying assumptions are the deepest and least accessible layer of organisational culture. They are the taken-for-granted beliefs about human nature, about what the organisation is for, about the relationship between different kinds of value, and about what is real and what is important, that operate beneath conscious articulation in the daily life of the organisation. They were not chosen. They accumulated over time, through the experience of what worked and what did not, and they have been absorbed into the organisation's operating system so thoroughly that most of its members cannot name them, would struggle to recognise them if asked, and would likely dispute that they are beliefs at all rather than obvious features of reality.

The assumption most consequential for safety culture, and most consistently misunderstood by leaders who encounter it, is rarely about safety directly. It is about the relationship between different kinds of organisational value — production, profitability, efficiency, safety — under conditions of tension. An organisation might have an explicit, frequently articulated value that safety is non-negotiable. Beneath that, at the level of basic assumption, might sit a belief that safety is the lens through which operational decisions are reviewed after they are made, not a constraint that shapes how they are made in the first place. The explicit value says: safety is non-negotiable. The underlying assumption says: production decisions are operational; safety decisions are oversight. The two are not in explicit conflict. Under pressure, the assumption governs, and the explicit value supplies the language used to describe the decision after the fact.

Assumptions are not stated. They are enacted, and they are enacted most clearly in three contexts: the way the organisation responds to unexpected events, the stories that circulate in the informal culture, and the gaps between what the organisation says and what it does.

The response to unexpected events is the most direct diagnostic. When a serious near-miss occurs, how does the organisation's response unfold? Not the official investigation process, but the informal response in the first hours

and days. Does the immediate reaction focus on understanding what system conditions produced the event, or on identifying the individual whose behaviour most directly preceded it? The answer reveals the organisation's underlying assumption about causation: whether it believes incidents are the product of systems and conditions, or whether, beneath a systems-thinking vocabulary, it still operates on the assumption that incidents happen because someone did the wrong thing.

The stories that circulate in the informal culture are the second diagnostic. Every organisation has stories: about incidents and how they were handled, about workers who raised concerns and what happened to them, about leaders whose behaviour in moments of pressure revealed what they actually valued. These stories travel through the informal network with a fidelity and persistence that formal communication rarely achieves, because they serve a function formal communication cannot: they teach new members what the organisation actually believes. A worker who, within their first few weeks, hears a story about a colleague who was sidelined after raising a concern has received a precise and actionable piece of cultural information, regardless of what senior leaders and the safety policy say.

The gaps between what the organisation says and what it does are the third, and the richest, diagnostic. Every significant gap between an espoused value and an enacted one is the shadow of an underlying assumption: the belief that makes the gap feel natural rather than contradictory to the people who live with it. An organisation that says it values near-miss reporting but consistently fails to close the loop between what is reported and what is actioned is operating from an assumption, not necessarily articulated or even consciously held, that near-miss reports are a compliance activity rather than a source of operational intelligence. The assumption holds the gap in place, and it will not change until the leadership behaviour that sustains it changes.

How Assumptions Form and Why They Persist

Organisational assumptions do not form arbitrarily. They form through a process that Schein¹⁴ described with precision and that decades of subsequent research have supported: the group learns, through repeated experience, that certain ways of interpreting and responding to its environment

produce successful outcomes. Those interpretations and responses become habitual. The habits become invisible. The invisible habits become the natural order of things: the way things are, rather than the way things have been chosen to be.

In a safety context, this process is particularly consequential because the feedback loop that validates assumptions about safety management is systematically distorted. Safe conditions produce no events, no data and no compelling narrative. The assumption that the organisation's risk management approach is adequate is confirmed, day after day, by the absence of evidence to the contrary. The assumption that near-miss reporting matters less than production is confirmed every time near-misses go unreported and nothing consequential happens. The assumption that the safety management system is working is confirmed by every audit that passes and every certification that is maintained. Because the experience that formed these assumptions keeps producing outcomes that read as confirmation, they deepen and harden, becoming progressively less open to challenge.

The social dynamics of group membership accelerate the hardening. Individuals whose assumptions differ from the group's come under significant pressure to revise or suppress them. The new worker who arrives with a higher standard of near-miss reporting than the existing culture supports quickly learns, through informal signals rather than explicit instruction, that their standard is out of step and that maintaining it carries a social cost. The experienced safety professional who challenges the group's belief that its current approach to risk is adequate meets the same dynamic at a more sophisticated level. Assumptions protect themselves, because the group has a collective interest in the stability and coherence they provide, and because challenges to fundamental assumptions threaten the group's identity and its shared story of success. Chapter Five examines that threat in detail.

This is why changing organisational culture is harder than changing structure or procedure, and why a new safety programme or a revised safety management system so often fails to produce the cultural shift it was meant to generate. Procedures can be changed by authority. Structure can be changed by design. Assumptions can only be changed by experience: the repeated, consistent demonstration, over a period long enough for the group to absorb it, that the assumption is wrong and that a different reading of reality produces better outcomes. That experience depends on leadership behaviour that is

visible enough, and consistent enough, to challenge the existing assumption and provide evidence for a new one.

Reading the Culture: A Leadership Diagnostic

The practical implication of understanding culture in its three layers is that effective safety leadership requires a diagnostic orientation — a continuous, deliberate inquiry into the gap between what the organisation presents and what it believes. This is not the same as scepticism, and it is not the same as distrust of the workforce or the management team. It is the intellectual discipline of a leader who understands that the most important information about the organisation's safety culture is almost never found on the surface, and who has developed the habits of attention and inquiry that allow deeper information to surface.

The diagnostic begins with artefacts, because artefacts are what is available first. It does not stop there. A leader reading artefacts diagnostically asks not what is present but what its presence means, and then how they would know whether that first interpretation is correct. The clean site and the current documentation are noted, and set aside as insufficiently informative. The leader looks instead for the artefacts that are harder to maintain under pressure: the quality of informal safety conversations, the frequency and character of voluntary reporting, and whether the corrective action system responds to actual operational conditions rather than scheduled review dates. These reveal more about the underlying culture because they are unlikely to be sustained without a culture that values what they represent.

..... Culture is not something an organisation has. It is something an organisation is. The distinction transforms what leadership is required to do.

Moving to espoused values means testing what the leader hears against what they observe. Every stated value is a hypothesis about the enacted culture, and every observation of behaviour under pressure is a data point that supports

or challenges it. A leader who is told that workers feel comfortable raising safety concerns tests that claim against the concerns actually raised in their presence: whether they come spontaneously or only when invited, whether they are voiced with the ease of someone whose previous concerns were well received or with the careful framing of someone expecting a defensive response, and whether they are substantive or performative.

Reaching basic assumptions is the most demanding step, because it requires the leader to surface what is, by definition, below conscious organisational thought. The most dependable method is to ask the same question across many contexts: what would the organisation have to believe, at a level it does not articulate, for this behaviour to make sense? When workers do not report near-misses despite a well-publicised reporting system, what would the organisation have to believe for that silence to be rational? When corrective actions are closed administratively rather than operationally, what must the organisation assume the corrective action system is for? When the same safety concern surfaces repeatedly across different reviews and time periods without producing change, what does the organisation implicitly believe about who is responsible for that change?

The answers, accumulated across contexts and checked against both artefact observation and the testing of espoused values, produce a picture of the culture's basic assumptions that is considerably more accurate than any survey, any audit, or any amount of time spent reviewing documentation. They also produce something more useful still: the identification of the specific assumptions creating the largest gaps between what the organisation says it does and what it actually does, and therefore the assumptions that most need to be challenged if the safety culture is to develop.

Culture is not a mystery. It is a system, and like all systems it responds to the discipline of structured inquiry. The discipline required of leaders is the willingness to move through its layers, looking past what is visible and comfortable toward what is hidden, challenging and consequential, and to act, consistently and persistently, on what they find. The leader who can read culture in this way has access to the strongest lever available for improving safety performance. The leader who cannot is confined to managing its surface, which, as the history of industrial accidents consistently demonstrates, is not enough.

PLEASE RESPECT COPYRIGHT

CHAPTER FIVE

Why Culture Fails: Anxiety, Identity, and Misalignment

The investigation took four weeks. Its findings were precise: a worker had sustained serious injuries as the predictable outcome of a known and documented condition that had been reported three times in eighteen months, addressed administratively on each occasion, and never resolved at the level of the underlying operational condition. No evidence of deliberate neglect. No rogue decision. No individual whose behaviour, examined in isolation, was obviously wrong. What the investigation found instead was a system that had processed every concern correctly — received it, categorised it, assigned it, closed it — while the condition those concerns described continued to exist and continued to develop.

What made the finding particularly difficult for the executive team was not the investigation's conclusions. It was what those conclusions implied about the six weeks that preceded the incident. Six weeks before the injury, the organisation had completed the most extensive safety culture survey it had ever conducted: three months to design, six weeks to administer, four weeks to analyse. The results had been presented on a Tuesday morning, in a room with good coffee and a considerable investment of organisational attention. The overall engagement score was positive. Workers reported high confidence in the organisation's safety commitment. They agreed that their supervisors took safety seriously, that they felt comfortable raising concerns, and that safety culture had improved significantly over the preceding two years. The site had not had a recordable injury in fourteen months. The survey had been thorough, rigorous and entirely accurate.

It had also been entirely blind to the condition that injured the worker six weeks later. Not because the survey had failed, but because the layers of culture it could reach were not the layers where that condition lived. Workers believed what they told the survey. The culture the survey described was real. What it was not was complete. Beneath the espoused values and the goodwill, at the level of basic assumption that no survey instrument is designed to surface, a different set of beliefs governed what actually happened when a concern was raised, processed and closed without being resolved.

The survey had measured the culture accurately. It had measured the wrong layer.

Culture as an Anxiety-Management System

The function of organisational culture that is least discussed in safety literature and most consequential for safety outcomes is also the most fundamental: culture exists, at its deepest level, to manage anxiety. This is not a metaphor. It is a description of the psychological function that Schein identified as the driver of cultural formation, and the force that makes cultural change so persistently difficult.

Human beings are, individually and collectively, profoundly uncomfortable with uncertainty. The inability to predict, understand and control the events of one's environment produces anxiety that is psychologically costly and operationally debilitating. Organisations, as social systems made up of human beings, develop the same need to reduce uncertainty and the anxiety it generates. Culture is the mechanism through which that reduction is achieved. By providing a shared framework of assumptions about how the world works, what is to be expected, how problems are to be interpreted and addressed, and what outcomes are within the organisation's capacity to produce, culture gives its members the cognitive stability required to function within complexity.

This function is valuable. An organisation whose members were uncertain about every aspect of their operational environment could not function. The shared assumptions of culture — about authority, process, risk, and the relationship between effort and outcome — allow people to act with

Beyond Compliance

confidence in conditions of real complexity. They provide the psychological foundation from which effective work becomes possible.

The problem, and it is a deeply consequential one for safety, is that the same mechanism that makes culture valuable makes it dangerous. Culture reduces anxiety by providing certainty: by telling the organisation what to believe about its environment and how to interpret the information it receives. That certainty is selective. It covers the things the culture's assumptions address, and those assumptions were formed in a past environment that may no longer describe current conditions. When the environment changes — when new risks emerge, when operational conditions shift, when the gap between how work is designed and how it is done widens further than earlier experience anticipated — the culture continues to provide the same certainty it always provided, about a state of affairs that no longer exists.

In safety, this mechanism operates with particular force. The environment is complex and uncertain, and allowing that uncertainty to be felt in full is psychologically intolerable for leaders who have built their professional identity around managing risk. It would mean accepting that the organisation's understanding of its own risk landscape is incomplete, that its controls may be failing in ways not yet visible, and that the confidence its dashboards produce may be out of proportion to the actual state of safety. Culture relieves that anxiety by asserting, through its assumptions, that the situation is understood, the controls are working, and the organisation is doing what it needs to do. The assertion may be false, but it is deeply and urgently needed, and the culture will defend it with considerable force.

Chapter Two described how the most sophisticated safety management systems generate the most confidence. The anxiety-management function explains why they also generate the most resistance. The more elaborate the safety management apparatus, the more completely it satisfies the culture's need for certainty, and the more threatening any information becomes that suggests the apparatus is not producing what its complexity implies. The system becomes an anxiety-management tool masquerading as a safety management tool. From the inside, the distinction is invisible. It becomes apparent only from the outside, or in the aftermath of an event that the system's elaborate reassurance did not prevent.

When Anxiety-Reduction Becomes Truth-Avoidance

The shift from legitimate uncertainty management to systematic truth-avoidance is not a dramatic event. It does not announce itself. It happens through the same gradual drift that produces every other form of organisational safety failure: a series of small decisions, each locally rational and individually defensible, that together progressively narrow the organisation's capacity to receive and act on information that is inconsistent with its existing assumptions.

The first stage is the filtering of uncomfortable information. All organisations filter information. The volume of data available to a complex organisation is far too large to be processed in full at every level, and filtering is both necessary and rational. The lethal version is the selective exclusion of information that challenges the organisation's current self-assessment. Near-miss reports that reveal systemic conditions are processed as individual events and closed individually. Concerns raised through formal channels are addressed at the level of the immediate issue, without investigation of the underlying conditions they reveal. Audit findings that point to structural problems receive procedural responses. The filtering is not dishonest. It is the natural operation of a system designed to process information within an existing framework of assumptions, a framework that does not contain the categories needed to receive information that falls outside it.

The second stage is the reinterpretation of anomalous data. When information is too prominent or too formal to be filtered out, the culture reinterprets it within its existing assumptions. A cluster of near-misses in one area is attributed to the individual characteristics of the workers there, rather than to the systemic conditions those workers have been navigating. A pattern of corrective actions closed without resolution is attributed to the administrative habits of a particular safety manager, rather than to a cultural assumption about what the corrective action system is for. A declining rate of voluntary safety reporting is read as an improving safety climate, rather than as a workforce that has learned not to bother. In each case, the anomalous data is absorbed by the culture's existing explanations and converted into a form consistent with its current self-understanding. The information is received. The assumption is protected.

The third stage, and the least visible, is the development of an organisational memory that systematically underrepresents the conditions that precede serious events. When near-misses are not reported, they do not enter the organisational record. When reported near-misses are closed without real investigation, the systemic conditions they revealed are not documented. When corrective actions address symptoms rather than causes, the causes accumulate without record. Over time, the organisation's formal memory — its incident history, its corrective action register, its risk assessment library — comes to describe a progressively sanitised version of its operational reality. Leaders whose understanding of safety is built on this formal memory are not receiving a biased account of their organisation. They are receiving an account that is accurate as far as it goes, an account of what was reported and how it was addressed, while being systematically unaware that it omits the most important part of the story.

The Texas City investigation²⁵ described this dynamic precisely. The formal safety record did not contain an accurate account of the process safety conditions that had been deteriorating year by year. That information existed in the informal knowledge of the workforce but had never been recorded in a form the organisation's safety management culture was equipped to process. So it remained invisible to the leaders who needed it most.

Identity and the Refusal of Information

The anxiety that culture manages is not only about operational uncertainty. It is also, and at least equally, about identity. Organisations develop identities — about who they are, what they stand for, what they are capable of — that are as psychologically important to their members as individual identity is to people. The identity of a well-managed, safety-committed organisation is a source of pride, cohesion and meaning, and it is defended with the same force, and by the same psychological mechanisms, as individual identity.

Those mechanisms — denial, rationalisation, selective attention, and the discounting or reinterpretation of disconfirming evidence — are well documented in individual and social psychology^{29,30,31,32}. In an organisation, they operate through the collective processes of cultural transmission: which accounts get repeated and which quietly disappear, how challenging information is framed and absorbed, and the social sanctions applied to

members who threaten the group's self-concept by producing information or interpretations inconsistent with it.

For safety culture, the identity most powerfully defended is usually some version of the organisation that takes safety seriously: the organisation that has made a real commitment, invested real resources and achieved real results. This identity is not false. The commitment is sincere. The resources have been invested. The results, in injury reduction, system development and the other visible markers of improvement, are real. But the identity, once established, creates a cognitive structure that makes certain kinds of information very difficult to receive. In particular, it makes difficult any information suggesting the investment has produced incomplete results: that the commitment, sincere as it is, has not reached the depth of the culture it was meant to transform, and that beneath the visible improvement, conditions are developing that the improvement has not addressed.

This is the mechanism behind what investigators of major accidents frequently describe as the normalcy of the pre-incident period: the absence, in the organisation's experience of itself, of any signal that something is wrong. The organisation is not in denial in the dramatic sense. It is not consciously suppressing uncomfortable truths. It is experiencing its culture accurately at the level of its espoused values and its artefacts, which really are as described. What it is not experiencing is the layer below: the assumptions that have made certain conditions invisible, the interpretive framework that has made certain information unprocessable, and the identity investment that makes questioning that framework feel like a threat rather than an opportunity.

The leaders most vulnerable to this dynamic are, paradoxically, the most committed. A leader who has invested heavily, personally and professionally, in building a safety culture has the most to lose from information suggesting the culture is not what they believe it to be. Their identity as a safety leader is bound up with the culture they have built, and the psychological cost of revising their assessment of it is proportional to the depth of their investment. This is not weakness. It is the natural operation of the same identity-formation process that makes deep commitment possible. But it requires specific countermeasures — structures of accountability, habits of inquiry, and relationships of honest feedback — that most leaders do not develop without deliberate effort.

The Latent Conditions That Misalignment Creates

When the visible layer of an organisation's culture, its artefacts and espoused values, is inconsistent with the invisible layer of its basic assumptions, the consequence is not immediately apparent. The organisation continues to function. Its safety metrics continue to move in the expected direction. Its management systems continue to generate the documentation and compliance evidence that audits require. The inconsistency does not produce an event, because it is not itself the cause of events. It is the condition that allows causal chains to develop undetected, and to accumulate until they reach the threshold of a serious outcome.

This is where Reason's latent conditions, introduced in Chapter One, take hold. Cultural misalignment generates latent conditions because it opens gaps between what the organisation believes about its own safety performance and what that performance actually is, and those gaps persist because the organisation lacks the cultural mechanism to detect them. At first the gap is small: the assumptions are broadly accurate, the artefacts and espoused values are reasonably consistent with enacted behaviour, and the risk controls are substantially working as designed. Over time, through the drift and normalisation described in Chapter One, it widens. Assumptions that were once approximately accurate become progressively less so. Artefacts that once reflected real operational practice become disconnected from it. Risk controls continue to be maintained and certified while the conditions they address evolve in ways the controls were never designed to accommodate.

The event, when it comes, is experienced by the organisation as a surprise: evidence that something went wrong in a way no one anticipated. That experience is not dishonest. The organisation really did not anticipate the specific configuration of conditions that produced the event. But the surprise is itself a product of the misalignment that allowed the latent conditions to accumulate. An organisation whose assumptions were accurate, whose leadership's model of operational reality matched the reality workers actually navigated, would have anticipated those conditions, because it would have been receiving the information that described their development. The surprise is the signature of the gap, not the content of the event.

Cultural Closure and the Loss of Peripheral Vision

There is a particular pathology of well-established cultures that is distinct from the general problem of misalignment and deserves explicit attention: cultural closure, the progressive narrowing of the range of information, interpretations and perspectives a culture is able to receive and process.

Cultural closure is the organisational equivalent of the confirmation bias that cognitive psychology has documented in individual decision-making: the tendency to seek, attend to and retain information consistent with existing beliefs, while filtering, discounting or reinterpreting information that challenges them³³. In individuals, confirmation bias can be partially offset through deliberate effort and structural countermeasures. In organisations, it operates through the social and structural machinery of culture: through the career trajectories of people who challenge cultural assumptions compared with those who confirm them, and through formal systems of information processing that are designed around existing assumptions and therefore systematically inadequate to information that falls outside them.

The result is a loss of peripheral vision: the progressive narrowing of the organisation's awareness to the range of conditions and events its existing assumptions are equipped to register. Events and conditions outside that range — novel risks, familiar risks appearing in unfamiliar combinations, or the particular configuration of conditions that makes up the trajectory of a serious incident — are present in the operational environment but absent from the cultural register. They happen without being seen, because the culture does not have the interpretive apparatus to see them.

Closure feeds on success. A long period of safe operation produces the assumption that the current approach to safety management is adequate to the current risk environment, and the longer the period, the harder that assumption becomes and the narrower the range of information the culture can process without triggering its defences. Why the strongest cultures are often the most exposed to this is the subject of Chapter Seven.

The Conditions for Cultural Failure

Cultural failure — the catastrophic misalignment between a culture's model of its own safety performance and the actual state of safety conditions — does not require negligence. It does not require dishonesty. It requires only the sustained operation of the mechanisms described in this chapter, and what makes them so consequential is that each strengthens the others. Anxiety-reduction supplies the certainty in which identity investment deepens³⁴. Identity investment drives the filtering and reinterpretation through which cultural closure develops³⁵. Cultural closure produces an information environment that sustains the anxiety-reducing assumptions, because if challenging information never enters the culture, its assumptions are never challenged, and the gap between the culture's self-understanding and operational reality widens without ever producing the dissonance that might trigger correction^{36,37}.

The conditions for cultural failure are therefore not exceptional. They are the normal operating conditions of any mature, successful organisation working in a complex environment over an extended period. An organisation that has not developed explicit countermeasures to these mechanisms is not a failure. It is simply an organisation that has not yet met the configuration of conditions that will make the gap between what it believes about itself and what is actually true consequential.

The countermeasures required are not technical. They are cultural and leadership-based, and they operate at each layer of the cultural model. At the artefact level, they require systems that surface information the existing culture does not seek: formal mechanisms for gathering operational intelligence from frontline workers, independent verification of whether risk controls are effective, and audit methods that test the gap between procedure and practice rather than the mere presence of procedure³⁸. At the level of espoused values, they require accountability structures that test those values against enacted behaviour and attach real consequences to the gap between the two. At the level of basic assumptions, they require leadership behaviour that is, day after day and decision after decision, consistent enough with the values the culture is trying to embed that the accumulated experience of that behaviour gradually shifts the assumptions currently governing how the organisation interprets its own safety performance.

None of this happens in the safety management system, and none of it can be delegated to the safety function. It happens in the invisible layer of the organisation, where the assumptions that govern real safety behaviour live, where the latent conditions accumulate, and where the gap between what the culture believes and what is true is either found and closed, or left to grow until the organisation can no longer be surprised by what it produces.

PLEASE RESPECT COPYRIGHT

CHAPTER SIX

The Condition Beneath the Culture: Psychological Safety

She had known for eleven days. Not suspected — known. The pressure reading on the secondary relief valve had been not quite right for nearly a fortnight, appearing and disappearing in a pattern that experienced operators learn to recognise as the signature of a developing condition rather than an instrumentation fault. She had mentioned it once, obliquely, in a morning toolbox meeting, and the conversation had moved on. She had not raised it again.

It was not that she lacked the knowledge. She had twenty-three years of process experience and an understanding of this particular system that no procedure manual could replicate. It was not that she lacked the opportunity. The safety reporting system was accessible, the channels were open, the corrective action process was documented and current. It was not that she was indifferent to the risk. She thought about the valve every shift, checked it every time she passed, and noted the readings in a private record she kept in a notebook no one else would ever see.

What she lacked was the belief that raising the concern would produce anything other than the particular kind of professional discomfort she had learned, over twenty-three years, to avoid. Not punishment. Not discipline. Nothing as dramatic or as actionable as that. Simply the experience of being the person who raised the thing nobody wanted to deal with, in a culture that had developed, without anyone deciding it, the collective habit of not dealing with things that were difficult to deal with. The unspoken understanding that some information, once raised formally, creates obligations that are easier to avoid if it is never raised at all. The look across the table. The slight shift in how her name appeared on the next month's roster. The subtle, entirely

deniable, entirely real signals that told her what everyone in that room already knew: that the right answer was silence.

She kept the notebook. She kept checking the valve. Eleven days after she first noticed the anomaly, the condition it represented produced the event it had been signalling.

The investigation found no evidence of the notebook. It found no evidence that anyone had known. The corrective action register was clean. The toolbox meeting records were complete. The culture, in all its documented dimensions, appeared exactly as it had been designed to appear.

The Condition That Makes Everything Else Possible

The gap between artefacts and assumptions, the filtering of uncomfortable information, the silence that passes for alignment: each of these depends on a single organisational condition, and whether that condition is present decides whether the rest of the safety culture apparatus functions or merely performs. That condition is psychological safety: the shared belief among the members of a team or organisation that the interpersonal environment is safe for taking risks. That concerns can be raised, mistakes admitted, questions asked and challenges made without the expectation of punishment, humiliation, or the quieter forms of social consequence that most organisational cultures have developed to discourage exactly these behaviours.

The idea has a longer history than its recent popularity suggests. Edgar Schein and Warren Bennis³⁹ wrote about psychological safety in the 1960s, as a condition people need before they can learn and change, and William Kahn⁴⁰ later studied it as one of the conditions that allow people to bring themselves fully to their work. It was Amy Edmondson⁴¹ who turned it into a shared property of teams that could be measured, and whose decades of research have produced the most empirically grounded account available of why some teams learn from their experience and others do not;; why some organisations surface the information that would allow them to improve their safety performance and others systematically suppress it; and why the gap between what workers know and what organisations know is, in most industries, one of the least examined features of safety culture failure. The research is extensive,

consistent and practically precise in its implications. It is also, in mainstream safety management practice, substantially underused.

Psychological safety is not a personality trait, and it is not a measure of individual confidence or courage. It is a property of the group: of the interpersonal climate the group has developed through accumulated experience of what happens when someone takes an interpersonal risk. When a worker raises a safety concern and it is received with real interest and followed through, the climate shifts slightly toward greater psychological safety. When a concern is met with defensiveness, dismissal, or the subtle administrative consequences that experienced workers read as clearly as any explicit signal, the climate shifts the other way. The climate is the accumulated product of thousands of such interactions across the history of the team, and it shapes, with a reliability that individual courage cannot consistently override, what people will and will not say in the situations that matter most.

This is the mechanism that produced the operator with the notebook. She was not a coward. She was not indifferent to safety. She was a rational, experienced professional working within a climate that, over years, had made silence the rational response to the kind of concern she was holding. That climate was not visible in the corrective action register, the incident history or the toolbox meeting records. It was visible only in the pattern of what was said and what was not, what was raised and what was managed privately, what the culture rewarded and what it quietly and consistently penalised. The safety management system never detected it. The investigation never found it. Yet it shaped the safety behaviour of every person in that organisation, every shift, more powerfully than any procedure the system had ever produced.

What Psychological Safety Is Not

Psychological safety is now popular enough to be widely misunderstood, and it is. The misunderstandings are worth addressing directly, because they produce organisations that believe they have built psychological safety when they have built something quite different, and they are then surprised when the silence continues.

The first and most common misunderstanding is that psychological safety means comfort: that a psychologically safe culture is one in which people feel good, conflict is avoided and challenge is kept to a minimum in the interests of positive relationships. This is almost the exact opposite of what

Edmondson's research describes. Psychological safety is not the absence of challenge. It is the condition that makes real challenge possible: the belief that disagreement, question and concern can be expressed without becoming a personal liability. A psychologically safe culture is not a comfortable one. It is frequently uncomfortable, because it is a culture in which the difficult things are actually said rather than managed into silence. The comfort is not in the absence of difficulty. It is in the knowledge that difficulty, honestly expressed, will be received as a contribution rather than a threat.

The second misunderstanding is that psychological safety is the same as trust. The two are related but distinct. Trust is a property of a relationship between two people: what one believes about the other's intentions and reliability. Psychological safety is a property of the group climate: what everyone in the group believes about the consequences of interpersonal risk-taking in this particular setting. A worker can trust their supervisor completely and still not feel safe enough to raise a concern in a team meeting, because the climate of the group, shaped by the behaviour of others present, by the history of what has happened when concerns were raised, and by the unspoken norms about what information is welcome, does not support it. Building trust is necessary. It is not sufficient.

The third misunderstanding, and the most dangerous in safety contexts, is that psychological safety is a soft concept: the domain of human resources and wellbeing programmes, valuable but peripheral to the serious business of safety management. Edmondson's research makes the opposite case. Psychological safety predicts learning behaviour, error reporting, information sharing and the willingness to challenge authority, each of which directly determines safety outcomes. In medical teams, higher psychological safety has been associated with higher reported rates of medication error, not because those teams make more errors, but because they report more of them, creating the information environment in which errors can be caught and corrected before they reach the patient⁴². In industrial settings, it is a strong predictor of near-miss reporting^{43,44}. In aviation, crew resource management training was developed to make it possible for a first officer to challenge a captain's decision. Yet when researchers analysed the cockpit voice recordings of 172 real accidents, they found that the crew had spoken up in all but two of them⁴⁵. What failed, again and again, was the listening: the concern was raised, and the captain, or the climate of the flight deck, did not let it change what happened next. Speaking up is not enough on its own. The variable that matters is

psychological safety: whether the flight deck, shaped above all by the captain, makes a challenge both safe to raise and certain to be heard. Pilots themselves say so. Most report that the captain sets the tone of the flight deck, and most say that as first officers they have held back safety information, or been silenced after raising it, because the flight deck did not feel psychologically safe⁴⁶. And accident investigators have repeatedly found that where the authority gradient between captain and first officer is too steep, errors go undetected and uncorrected even by experienced crews⁴⁷.

Psychological safety is not soft. It is among the hardest safety culture conditions to build. The operator with the notebook was not the victim of a programme failure. She was the victim of a psychological safety failure, and the two are not the same and are not fixed by the same interventions.

How the Climate Forms

Psychological safety does not arrive in an organisation because its leaders declare it. It accumulates, or fails to, through the observable, remembered interactions that tell every member of the group what happens when someone takes the interpersonal risks a strong safety culture depends on.

Edmondson described the pattern with precision. Each member of the group holds beliefs about the consequences of interpersonal risk-taking, drawn from their experience of what has happened when such risks were taken in this setting. Those beliefs shape what they are willing to say in situations where taking a risk might be warranted. The collective pattern of what is said and not said then becomes the evidence from which every new member forms their own beliefs about the climate. The climate is self-reinforcing. A low-safety climate produces silence; silence removes the contradicting evidence that might challenge the low-safety belief; and the belief persists and deepens. A high-safety climate produces disclosure; disclosure produces the experience of being properly heard, which reinforces the high-safety belief; and that belief, too, persists and deepens. Both states are stable. Neither changes quickly without deliberate, sustained and visible leadership intervention.

The interactions that most powerfully shape the climate are not the dramatic ones: the formal investigations, the disciplinary processes, the highly visible responses to serious events. They are the ordinary daily interactions through which workers calibrate their sense of what information is welcome. The leader who receives a near-miss report with real interest, and comes back

within twenty-four hours to say what was done with it, builds the climate in a lasting way. The leader who receives the same report with a slightly impatient nod and moves to the next agenda item erodes it just as quickly. Neither interaction announces itself as significant. Both are stored as evidence in the group's running assessment of whether honesty is safe here.

This is the feature of psychological safety that leaders most consistently underestimate. Most leadership development frameworks direct attention toward the visible, high-stakes moments: the public response to a serious event, the all-hands communication about safety culture, the formal launch of a speak-up programme. These are not irrelevant, but they matter far less to the climate than the texture of the ordinary: the quality of a leader's attention in the brief safety conversation at the start of a shift, the warmth or flatness of their response to a concern raised after a toolbox meeting, the follow-through or its absence in the week after a worker mentioned something that needed fixing. The climate is built in those moments. Not in the announcement.

The Silence Beneath the System

Each of the failure mechanisms described in Chapter Five depends on psychological safety as its enabling or inhibiting condition. Where psychological safety is high, those mechanisms are counteracted: uncomfortable information is more likely to surface, anomalies are more likely to be reported as anomalies rather than reinterpreted as background noise, and the organisation's investment in its current view of itself is more likely to be challenged by information that contradicts it. Where psychological safety is low, all of them operate without counterweight, because the interpersonal conditions that would generate the challenging information are absent. The knowledge exists somewhere in the organisation; it simply does not travel. A fitter who notices that a pressure reading has been drifting for a week either says so at the pre-start or decides it is someone else's job to notice. Which of those happens is rarely a matter of individual character. It is a property of the room he is standing in.

The silence that low psychological safety produces is not experienced as suppression by the people within it. It is experienced as normal. Workers in these cultures do not, in most cases, consciously decide to withhold safety information. They make continuous, largely automatic assessments of what is worth raising, guided by their experience of how the culture receives safety

information. That experience is learned the way most workplace judgement is learned: by watching. A worker who once saw a colleague's concern met with a sigh, a joke at the next meeting, or a request to put it in writing that was never followed up does not need to be told what that means. He files it away, and it informs every decision about speaking up that follows. The cost need not be dramatic to be decisive. From that accumulated experience, workers conclude, rationally given the evidence available to them, that the concern they are holding does not justify the interpersonal risk of raising it; or that raising it will produce administrative activity without resolution; or that the person they would need to raise it with will not receive it well. Each of these assessments may be accurate. None of them looks, from within the culture, like a safety culture failure. They look like sound professional judgement about what is worth saying, to whom, and when.

This is why the silence is so hard to detect from outside, and why the formal indicators of safety culture health are such unreliable measures of psychological safety. A culture with low psychological safety will likely show near-miss reporting rates sufficient to satisfy audit criteria, toolbox meeting participation that is consistent and apparently engaged, and survey responses describing a workforce that feels comfortable raising concerns. The reporting rate is sufficient because the threshold for reporting has settled at what the culture will receive without discomfort: the minor, the routine, the uncontroversial. Toolbox participation is consistent because attendance is expected, the content is undemanding, and the questions asked are the ones that have safe answers. The survey responses are positive because the question — do you feel comfortable raising safety concerns? — is answered with reference to the concerns people do raise, not the ones they have already decided are not worth raising. It is also answered by people who have their own views about who will read the results. The instrument measures the surface and reports the measurement as if it described the whole. In some cases the formal indicators look best precisely where the climate is worst, because a workforce that has learned what is welcome will supply it reliably. The leaders reading those indicators are not being deceived. They are being shown exactly what the culture is able to show them.

What it does not capture, and what no instrument relying on voluntary disclosure can reliably capture in a culture with low psychological safety, is the notebook. The private record of observations that never entered the formal system. The conversation that happened in the car park after the meeting

because it could not happen in the meeting. The quiet word to a mate instead of a word to the supervisor. The workaround everyone on the crew uses and nobody has ever written down. The decision, made in seconds and never revisited, that this particular thing was not worth the particular cost of raising it in this particular organisation. Some of it is written down. Most of it is carried in memory, shared selectively, and lost when people move on. These are the data that describe the actual state of the psychological climate, and they are precisely the data the climate makes hardest to collect. When they do finally surface on their own, it is usually in the interview room of an investigation.

..... The silence that low psychological safety produces is not experienced as suppression. It is experienced as normal — as good professional judgment about what is worth saying. That is what makes it so difficult to detect and so consequential to ignore.

Fear, Accountability, and the Crucial Distinction

The most persistent objection to psychological safety in safety management is a reasonable one, and it deserves a direct answer: if workers feel entirely safe from consequence, what prevents unsafe behaviour from going unchallenged? If psychological safety means people will not be held accountable, does it not create the conditions for exactly the risk-taking that safety management exists to prevent?

The objection rests on a confusion that Edmondson is careful to untangle. Psychological safety is not the absence of accountability. It is the condition in which concerns can be raised, mistakes reported and challenges made without fear of interpersonal reprisal³⁶. Accountability — the clear expectation that safety standards will be met and that departures from them will be addressed — is entirely compatible with psychological safety and is, in effective safety cultures, its necessary companion. The two operate in different domains. Psychological safety governs the interpersonal climate: what people will say. Accountability governs the performance standard: what people are expected to do. A culture that has both is one in which workers report their

own near-misses without fear of blame, and in which the standards that govern safe behaviour are clear, consistent and enforced.

The cultures that most reliably fail on safety are not those with too much psychological safety. They are those that have mistaken its absence for the presence of accountability. An organisation in which workers are afraid to report near-misses is not an organisation with high accountability. It is an organisation with high fear, and fear is among the least reliable safety management tools available.

Fear produces compliance under observation and concealment under pressure. It drives safety-relevant information underground, where it accumulates as the latent conditions that serious events require. It creates the environment in which admitting a problem is more dangerous than concealing one, and that environment produces the operator with the notebook, the engineer who knows about the O-ring, and every other instance in which the information that would have prevented the event existed in the organisation and never reached the people who could have acted on it.



The distinction between fear and accountability is not abstract. It shows in observable leadership behaviour. The leader who responds to a reported error by asking what system conditions made it likely sends a different climate signal from the leader who asks who made it. The leader who thanks a worker for raising a concern about a developing condition, specifically and visibly, in a way others observe, communicates something about the interpersonal climate that no policy document can replicate. The leader who then makes sure the concern is properly investigated and resolved, rather than administratively closed, supplies the follow-through that makes the signal credible.

These are not gestures. They are the cumulative, climate-forming behaviours that build psychological safety over time. They require a particular self-discipline of leaders: the consistent willingness to respond to uncomfortable information in ways that reward its disclosure rather than penalise it. Even when the information is inconvenient. Even when the condition it reveals implicates a management decision. Even when the easiest response would be to acknowledge it briefly and move on.

Building It

Building psychological safety in a safety culture is not a programme that can run in parallel with the existing safety management system. It is a transformation of the climate: slow, non-linear, vulnerable to setback, and entirely dependent on consistent leadership behaviour over a period that cannot be fixed in advance. What can be described is the leadership practice that builds it, and the behaviours that most reliably erode it.

The foundational practice is what Edmondson calls framing⁴⁸: leaders explaining explicitly why input, concern and challenge are valued, in terms concrete enough to be believed. A leader who opens a safety discussion by saying they are particularly interested in what is not in the system — the concerns that have not been formally raised, the conditions experienced workers are navigating without formal support, the gaps between how work is designed and how it is actually done — has told the room something specific. The threshold for raising something is lower than the formal reporting system implies, and the interpersonal risk of raising something below it is one the leader is actively inviting rather than merely tolerating. This is not a technique. It is the honest expression of a leader who wants to know, and that honesty is as legible to the workforce as its absence would be.

The second practice is curiosity without judgement in response to what is raised. When a worker shares a safety concern, the leader's response in the first thirty seconds is one of the strongest climate signals available. Curiosity — visible engagement with what was shared, the follow-up question that shows the concern has been heard rather than merely received — builds the climate. Judgement — the immediate assessment of whether the concern is significant enough to warrant attention, the slight impatience with a concern that interrupts the agenda, the response that converts the concern into a corrective action and closes it before the worker has finished speaking —

erodes it, and can destroy it. The difference between the two is often ten seconds and the quality of a leader's attention. Its effect on the climate is lasting.

The third practice is closing the loop: consistently, personally, and in a timeframe that shows the concern mattered. The worker who raises a concern and never hears what happened has learned something precise about the culture: that their contribution entered a system and disappeared. The worker who, within a week, receives a specific account of what was investigated, what was found and what was done has learned something different. The loop does not always close with the resolution the worker wanted. Sometimes the investigation finds the condition is understood and managed. Sometimes the fix takes longer than the worker would like. Closing the loop is not the resolution. It is the communication that the concern was received and acted on, and that communication, made consistently across every concern raised, is the accumulated experience from which the climate is built.

What erodes it is simpler and faster, because the accumulation is asymmetric: positive interactions build the climate slowly, and negative ones erode it quickly. A single instance in which a worker who raised a concern is subsequently treated differently — given fewer desirable shifts, left out of conversations they were previously part of, found to have made a technical error that attracts unusually close scrutiny — will travel through the informal culture within days and reshape the climate for people who were never in the room. The story does not need to be accurate in every detail. It needs only to be plausible, and in a culture that already has a low-safety climate, plausibility is easy to achieve. A leader who wants to build psychological safety must therefore attend not only to their own response to disclosed concerns but to the wider organisational signals about what happens to the people who disclose them. The climate is shaped by the story the culture tells about itself, and that story is told not by the leader's intentions but by the accumulated evidence of the organisation's behaviour.

The Ground Beneath Everything

Psychological safety is not another layer to be added to the architecture of culture described in the preceding chapters. It is the ground on which that architecture stands or collapses.

Much of what the rest of this book argues for depends, in operational reality, on psychological safety in the teams that make up the organisation. The High Reliability Organisations examined in Chapter Twelve show why. Their preoccupation with failure is impossible in a culture where reporting failure carries personal cost. Their sensitivity to operations, which keeps the gap between Work-as-Imagined and Work-as-Done in view, is impossible where what workers see cannot be safely shared with the leaders responsible for it. Their deference to expertise, which lets the person closest to a developing condition act on what they know, is impossible where the interpersonal risk of acting exceeds the apparent risk of the condition itself.

The safety management system, however well designed, cannot compensate for the absence of psychological safety. It can provide the channels through which safety information could travel. It cannot create the interpersonal conditions that make workers willing to use them. That is the work of leadership: not the leadership of safety management, but the kind of leadership that builds the human climate in which safety management either functions or merely performs.

The operator with the notebook did not fail the safety system. The safety system failed her. Not because it lacked the channels she could have used, but because the organisation had never built the climate that would have made using them feel safe. Her knowledge, her experience, her concern for the people who worked alongside her were all present, and all unavailable to the organisation, because the ground beneath the culture had never been laid.

Laying that ground is not the work of a programme. It is the work of leaders in the daily, granular, apparently unremarkable interactions through which the climate of every team is continuously made and remade. It requires them to show, again and again, that the kind of honesty safety depends on — the honesty that admits uncertainty, reports error, raises concern and challenges the established view when the established view is wrong — is not only safe in this organisation but welcomed.

When that has been shown, consistently and over time, the culture changes. Not dramatically. Not all at once. In the way cultural change always happens: through the accumulation of individual experiences that shift the shared assumption about what this organisation is and what it values. The notebook disappears. Not because the system improved, but because the person who would have kept it no longer needs to

CHAPTER SEVEN

The Anatomy of Weak and Strong Cultures

He had spent nearly forty years at sea. By his own description, it had been an uneventful career. Not in the sense of being without incident, because there had been winter gales and storms and fog, but in the sense that, through those four decades, he had never been in any accident worth speaking about. He had never seen a wreck. He had never been in any predicament that threatened to end in disaster.

This was not false modesty. It was an accurate description of a career defined by professional mastery, accumulated experience, and the kind of confidence that is not arrogance but the legitimate product of sustained competence in a demanding and dangerous environment. He was the kind of leader whose culture — the culture of an experienced maritime professional, and of the ships he commanded — was truly strong. It was coherent, widely shared among those who sailed with him, grounded in real experience, and resistant to the anxiety that lesser experience would have produced.

He was not especially concerned about icebergs. He was, in his own words, more afraid of fog.

The great icebergs, he had observed, were visible. You could see them, navigate around them, manage them with the tools available to a professional navigator of his era. What he feared was the hidden thing: the submerged reef extending beneath the surface of a melting iceberg, invisible to any instrument, capable of tearing the bottom from a vessel before the crew had registered the contact. The thing you could not see, because the thing you were trained to see, the visible mass above the waterline, did not reveal it.

The culture that produced that observation was a strong culture. It was also, in one specific and consequential respect, insufficiently afraid of the right things.

The Seduction of Cultural Strength

The language of safety culture almost universally treats strength as a virtue. A strong safety culture — widely shared, deeply embedded, resistant to erosion under operational pressure — is the aspiration toward which safety culture development programmes are directed, and against which maturity models measure progress. The language is intuitive. A strong culture seems, obviously, more likely to produce safe outcomes than a fragile or inconsistent one. And in most respects, the intuition is correct.

Strength, in culture as in engineering, describes a structure's resistance to force. It does not describe the direction in which the structure is pointed. A strong culture is difficult to change, difficult to challenge, and difficult to deflect from its current course. If that course is toward safety, if the values embedded in the culture are accurately aligned with what real safety requires, then strength is an unambiguous advantage. If the course is toward something else, strength is the mechanism by which the misdirection is preserved.

The distinction matters because it is invisible from the inside. A strong culture in which the wrong things are deeply embedded does not feel, to its members, like a culture pointed in the wrong direction. It feels like a culture. The shared assumptions, the common interpretations and the consistent behaviours all produce the sense of coherence and correctness that marks a strong culture. The members of an avoidant culture that is widely and coherently shared do not experience themselves as members of an avoidant culture. They experience themselves as members of a culture that has learned, through accumulated experience and repeatedly validated outcomes, how to understand and respond to its operational environment. That what they have learned may be systemically dangerous is invisible to them, because the culture is strong enough to make the dangerous things feel obviously right.

This is the seduction of cultural strength: it produces certainty. Certainty, as Chapter Five examined, is what anxiety-management requires. A strong culture gives its members a complete, coherent and authoritative account of how their world works and how they should respond to it. In a well-directed culture, that completeness and coherence is a real virtue. In a strong but

misdirected one, it is the mechanism by which the misdirection becomes immune to challenge.

This is why progress on a maturity model, the subject of Chapter Seventeen, cannot be read simply as growing cultural strength. An organisation can have a strong Lethal culture, in which the assumptions that generate dangerous conditions are widely shared, deeply embedded and rigorously defended by every member of the group. It can have a weak Functional one, in which the right values are held by some members but not yet broadly embedded, and in which resistance to challenge is low enough that those values can be displaced by pressure. Strength and virtue are independent properties. A culture can have either, both, or neither.

Strength Without Direction

The most instructive examples of strong cultures producing catastrophic safety failures share a common feature: the culture's strength was not incidental to the failure but fundamental to it. Its coherence, its resistance to challenge, and its capacity to make some things feel obviously true and others obviously irrelevant were what allowed the conditions for failure to develop undetected.

Challenger is the defining case. Chapter One introduced Vaughan's account of normalised deviance; what matters here is the kind of culture in which it took hold. NASA in 1986 was not a weak culture but an extraordinarily strong one: coherent, deeply shared, and grounded in an identity of technical excellence that had carried the programme through decades of extraordinary achievement. That strength is what stopped the engineers' O-ring concerns from changing the launch decision. The culture did not ignore the information. It received it, discussed it, and reinterpreted it within assumptions too deeply embedded to be revised by a single anomalous data point. The launch decision was not made by people who did not care about safety. It was made by people whose strong, coherent, widely shared culture had placed certain assumptions beyond the reach of challenge. That is not weakness. That is strength, pointed in the wrong direction, doing exactly what strength does.

Deepwater Horizon provides the parallel from a different industry. BP's offshore culture was not weak but extensively structured, and its strength was oriented toward visible personal safety at the expense of process safety. The

strength of that orientation made invisible the engineering-level conditions the April 2010 blowout exposed.

These are not exceptional cases. They are the expected outcome of strong cultures whose strength is not matched by the direction and adaptive capacity that make strength safe. The implication for safety leaders runs counter to most of the discourse on safety culture development: the goal is not a strong culture. The goal is an accurately oriented culture, one that is strong in the direction of safety, which requires not only widely shared safety values but the continuous alignment of those values with the actual conditions of the operational environment.

The Anatomy of Weak Culture

A weak safety culture is diverse: fragmented across different parts of the organisation, inconsistent in its expression, and variable in how it applies the values it nominally holds. A weak culture does not mean an organisation that does not care about safety. It means an organisation whose safety culture has not achieved enough coherence and consistency to shape behaviour reliably across all contexts and levels.

Weak cultures usually reflect weak safety leadership. Not the absence of committed safety leaders, but the absence of leadership behaviour that is consistent, visible and sustained enough to create the shared understanding that culture requires. Where leaders show safety commitment inconsistently, seriously engaged in some settings and disengaged in others, the workforce builds an accurate picture of the inconsistency and calibrates its own behaviour to match. Where leadership messages about safety are clear at executive level but diluted or contradicted at management level, the culture frontline workers experience is not the culture the board intends. Where different parts of the organisation operate under different leadership styles, accountability expectations and operational pressures, the safety culture of each part develops independently, and the result is a diverse culture that resists the consistent application of any single safety standard.

The vulnerabilities of weak culture are different from those of strong, misdirected culture. Where strong misdirected culture fails by making the wrong things feel obviously right, weak culture fails by not making the right things feel reliably important. Workers in a weak safety culture understand the organisation's espoused values. They often know what the system requires.

They are capable of behaving safely under direct observation, and of departing from safe behaviour when observation is absent and competing pressures are sufficient. The culture does not provide enough consistent counterpressure to make those departures reliably less likely than safe behaviour.

Yet the weak culture is also more open to directed change than the strong misdirected one. Because it has not embedded its assumptions deeply enough to make them hard to challenge, new leadership behaviour, new accountability structures and new cultural signals have a better chance of taking hold. The weakness that makes the culture vulnerable also makes it malleable, and the leader who treats a weak culture not as a failure but as a development opportunity has a point of leverage that the leader of a strong misdirected culture may struggle to create.

The practical challenge of a weak culture is that the diversity it produces is hard to diagnose without the kind of systematic cultural inquiry described in Chapter Four. An organisation with a weak safety culture will show variation in its diagnostic data that is more informative than the aggregate scores usually reported: strong engagement in some areas and weak in others, strong near-miss reporting in some teams and weak reporting in the team next door, strong procedure compliance on some shifts and variable compliance on others. The variation is the diagnostic information. It reveals which parts of the organisation have developed stronger sub-cultures, and therefore which leaders, at which levels, are already creating the conditions the organisation needs more broadly. A weak culture, carefully diagnosed, frequently contains the blueprint for the stronger culture it needs to become.

When Strong Means Rigid

The form of cultural strength that most often precedes serious industrial accidents is not the strength of a culture oriented in the wrong direction. It is the strength of a culture that was once oriented correctly, that developed its assumptions in an environment where they were accurate, and that has become rigid as the environment changed around it.

Rigidity is strength without adaptive capacity: the cultural equivalent of a material that is strong under static load but brittle under dynamic stress, not unlike a carbon bicycle frame. A culture that is strong and adaptive responds to changes in its operational environment by updating the assumptions that are no longer accurate, integrating new information into its interpretive

framework, and revising its espoused values when they prove insufficient to the conditions they are meant to address. A culture that is strong and rigid responds by maintaining its existing assumptions, filtering out or reinterpreting information that challenges them, and defending its espoused values by reframing difficult conditions as edge cases or temporary aberrations.

The difference shows most clearly in how organisations respond to near-misses and to the gradual erosion of safety margins. A strong culture that is still learning treats each near-miss as diagnostic information: evidence that the gap between Work-as-Imagined and Work-as-Done may be widening, and that its assumptions about the adequacy of current risk controls need testing. A rigid culture treats each near-miss as an anomaly: a departure from the normal state of affairs, explained by the circumstances of the event, requiring no revision of underlying assumptions that are too deeply embedded to be revised by a single event.

This is why a high score on a maturity assessment, taken at a single point in time, can mislead. It may reflect a culture that has developed real adaptive capacity. Or it may reflect a culture that has become exceptionally skilled at presenting the artefacts and espoused values of maturity without having embedded the assumptions that make maturity real. The difference is invisible to any assessment that does not probe below the surface, and the surface of a strong culture is a very effective place to hide what lies beneath it.

What Effectiveness Looks Like

An effective safety culture is not simply a widely shared one. A culture can be strong and wrong. Strength describes how widely a culture is held; effectiveness describes whether it is worth holding. An effective culture is a widely shared culture pointed in the right direction, maintained by leaders whose behaviour is consistent enough to hold that direction under pressure, and adaptive enough to update itself as the operational environment changes. Effectiveness, in this sense, has three components that can be examined separately and that, together, produce the kind of safety culture associated with exceptional safety outcomes.

The first is coherence: the degree to which the culture's artefacts, espoused values and basic assumptions are consistent with one another. In a coherent culture, the visible markers of safety commitment reflect the values leaders articulate, which in turn reflect the assumptions that actually govern

operational decisions. Coherence is not automatic, even in a strong culture. It requires the culture to have been built from the inside out rather than the outside in, with artefacts that express real values rather than perform expected ones. It can be assessed by examining the gap between what the culture presents and what it actually produces: between the near-miss reporting system and the quality of investigation it drives, between leaders' stated safety commitment and their behaviour when that commitment is tested, between the procedures on the shelf and the work done in the field. Where coherence is low, the workforce notices first. People learn to read the artefacts as decoration and to take their cues from decisions instead.

The second is alignment: the degree to which the culture's assumptions accurately describe the actual conditions of its operational environment. A culture can be coherent, internally consistent in its assumptions and practices, while being misaligned with its operational context. The maritime culture that placed its confidence in surface observation was coherent. Its practices were consistent with its assumptions, and its assumptions had been validated by decades of experience in which surface observation had been adequate to the risks encountered. It was misaligned with the conditions of that April night, because the risk it faced was below the surface, and its assumptions had never been formed to accommodate risks of that kind. Misalignment is the hardest of the three to see from inside, because a coherent culture produces confidence, and confidence is the condition least likely to go looking for evidence that it is misplaced.

The third is adaptability: the culture's capacity to revise its assumptions when the operational environment changes in ways that make them inaccurate. This is not the adaptation described in Chapter One, in which work practices drift away from procedure without anyone deciding they should. It is the opposite: a deliberate revision of what the organisation believes, made openly, in response to evidence. Adaptability depends on the psychological safety examined in Chapter Six, combined with a leadership orientation toward learning rather than confirming. A culture with this capacity actively seeks information that challenges its assumptions, because it understands that unchallenged assumptions in a changing environment are a source of risk, not of stability. It treats the anomaly that does not fit as something to be examined, not explained away.

Coherence, alignment and adaptability are the anatomy of effective safety culture. They are also the dimensions along which safety leadership operates. The behaviours that build coherence are different from those that maintain alignment, which are different again from those that foster adaptability. A leader who works hard on the wrong dimension can make a culture more confident without making it any safer. Understanding the difference is what allows a leader to diagnose their culture accurately and to direct their effort to the dimension that most needs it.

The Maturity of Culture

Here is what cultural maturity actually looks like, stripped of the language that has accumulated around it. It does not look like a score on an assessment. It does not look like a certificate on the wall or a maturity level declared in a board presentation. It looks like a workforce that raises concerns before they become incidents, because the culture has made raising concerns feel natural rather than risky. It looks like a leader who visits the worksite not to inspect but to understand, and who leaves knowing something they did not know when they arrived. It looks like a near-miss report investigated with the same rigour as a recordable injury, because the culture understands that the near-miss is the incident that has not yet happened, and that the difference between the two is narrower than the metrics suggest.

Cultural maturity is not the achievement of a strong culture. Strength, as this chapter has argued, is not a virtue when it is pointed in the wrong direction or maintained after the environment has changed beneath it. The mature culture is not the one most resistant to challenge. It is the one most capable of receiving challenge: the one that has built into its basic assumptions the understanding that its current model of its own safety performance is always incomplete, always open to revision, and always more confident than the evidence fully warrants.

This is the distinction most safety culture development programmes miss, and missing it is consequential. Organisations invest in building strong cultures when what they need are accurate ones. They invest in embedding values when what they need is the capacity to test whether the values they have embedded are producing what they believe. They celebrate the absence of incidents as evidence of cultural health when that absence is, at best, evidence that the gap between the conditions that exist and the conditions needed to produce an

incident has not yet closed, and at worst, evidence that the reporting culture has matured faster than the safety culture, and that incidents are occurring but not appearing in the data.

Nothing about this is hidden from the people doing the work. Workers know a culture's maturity with a precision no external assessment can match, because they experience it daily: in the decisions leaders make under pressure, in the fate of the concerns they raise, and in the gap between what the organisation says about itself in formal settings and what it shows in the operational moments those settings never capture. A worker who has experienced real cultural maturity, in an organisation whose assumptions about safety are accurate, whose leadership behaviour is consistent, and whose response to challenge is curiosity rather than defence, carries that experience as a standard. They know what it feels like. And they know, with equal precision, when the culture they are in now falls short of it.

Cultural maturity is not a destination. It is a direction and a discipline: the ongoing, never-finished work of keeping the culture's assumptions aligned with the operational reality they are supposed to describe. The organisations that achieve it do not do so by building a culture and maintaining it. They do so by building a culture capable of examining itself, one that has developed, in its leaders and its workforce alike, the habit of asking whether what they believe about their safety performance is actually true, and the courage to act on the answer when it is not.

PLEASE RESPECT COPYRIGHT

CHAPTER EIGHT

Why Safety Leadership Matters More Than Safety Management

He had been in the industry for twenty-two years. He understood the chemistry of the processes he managed with a precision the safety auditors who periodically assessed his site could not match. He knew the equipment: not just its functional specifications but its history, its quirks, and the particular ways it behaved under different loads and ambient temperatures. He knew the regulatory framework in three jurisdictions. He had completed every safety leadership programme his organisation offered, and completed them well. His safety metrics were good. His audit results were consistently acceptable.

Over those twenty-two years he had been promoted for these things. Not for these things alone; he was also organised, technically articulate, and politically competent in the ways organisations value in the leaders they advance. His safety credentials had been a significant part of his progression, because the organisations he worked for regarded safety management competence as a meaningful leadership differentiator, and he was truly competent at safety management.

What he had never been asked to do, in twenty-two years of professional development and advancement, was show that he could change the assumptions of a safety culture. What he had never been developed to do was create the conditions in which workers at every level of his organisation would value safety as a principle embedded in their professional identity, one that governed their behaviour whether or not anyone was watching. He had been developed to manage safety. He had never been developed to lead it.

The difference was invisible in the technical vocabulary of safety management, invisible in the audit criteria, invisible in the training records. In the culture of his organisation, it was entirely apparent.

The Management Default

W. Edwards Deming²⁷, whose work on quality management transformed manufacturing in the second half of the twentieth century and whose insights into organisational systems remain among the most rigorous available, made an observation about management so widely cited that its force is often lost in the familiarity. The job of management, he argued, is not supervision but leadership. The transformation management needs requires that managers become leaders.

..... focus on outcome — management by numbers, management by objectives, work standards, zero defects, performance appraisal — must be abolished and leadership put in its place.

Deming was writing about quality, but he was describing a structural problem that applies with equal force to safety: the systematic confusion, at every level of organisational design, of management activity with leadership activity, and the resulting investment of organisational energy in the first at the expense of the second. Management activities — developing and maintaining systems, monitoring metrics, enforcing compliance, administering corrective processes — are necessary. They are not sufficient. In organisations that treat them as sufficient, the culture that leadership could create, and that safety actually requires, is never built.

The management default is not accidental. It is the rational outcome of a regulatory and cultural environment that has, for most of the history of modern industrial safety, defined safety obligation in terms of system implementation and compliance demonstration. The frameworks that govern safety — management system standards, regulatory requirements, audit methods — are built to assess management activity. Do you have the procedures? Do you conduct the reviews? Do you maintain the records? Those are the questions the frameworks are designed to answer. The leadership

questions are not on the list. Do your people believe safety matters? Do they expect their concerns to be heard and acted on? Do they make safe decisions when no one is watching, because safe decisions are consistent with who they are?

So organisations learn to speak the language of safety management fluently and to perform its activities competently, without developing the leadership capability that transformational safety culture requires. Leaders measured on management activity develop management capability. Leaders measured on cultural outcomes develop leadership capability. In most safety management systems the measurement is overwhelmingly focused on management activity, which is why the gap between safety management and safety leadership is so widespread, so persistent and so costly.

None of this is a new discovery. In 1982, a United States National Research Council committee reviewing safety in underground coal mines concluded that while legislation had improved safety, compliance alone would not produce an outstanding safety record. The initiative had to come from management, and the committee added a warning that could stand as a summary of this chapter: commitment, it said, legislative compliance “can create only an illusion of accomplishment unless it results in action⁴⁹.” Sixteen years later, Lisa Ronald’s review of the research on successful safety programmes, prepared for the Workers’ Compensation Board of British Columbia, reached the same conclusion: active safety leadership and management commitment played an integral role in safety success, and quite possibly the most crucial one⁵⁰.

The safety intervention literature is instructive here. A landmark Campbell Collaboration systematic review, drawing on one hundred studies covering nearly thirty-two million individuals, found that interventions operating at the organisational and group level consistently outperform those targeting individual behaviour or administrative compliance⁵¹. Audit-based and documentation-heavy approaches are the weakest performers.

A parallel meta-analysis of ninety-five studies across fifteen countries found the same gradient within safety training⁵². Participatory, behaviourally engaged methods produce meaningfully greater injury reductions than passive, lecture-based ones. The method matters as much as the message.

The pattern is consistent across both bodies of evidence. Interventions that reach the level of culture, changing the conditions in which behaviour occurs rather than documenting the conditions that already exist, produce the

greatest safety returns. Those that stop at the level of system and documentation produce the least. And changing culture is a leadership function, not a management one.

Working Within the System and Working On It

The most fundamental distinction between management and leadership, in any domain, is the distinction between working within an existing system and working on it. A manager, by functional definition, operates within a system: implementing, maintaining and enforcing the processes and structures it comprises. A leader works on the system: creating, changing and shaping the conditions within which it operates. Both are necessary. Neither is sufficient without the other.

In safety, this distinction maps onto the gap the previous chapters have been describing. Safety management, working within the system, produces the administrative completeness that audits assess and compliance frameworks require. It keeps the system current and ensures it generates the documentation and operational discipline its design calls for. Done well, it is a real contribution to safety. But it cannot, by itself, create the cultural conditions that make the system real rather than virtual, because those conditions lie outside the system's boundaries. They live in the assumptions workers hold about whether the system matters, in the values leaders enact when safety and production compete, and in the organisation's basic orientation toward safety as a principle rather than a programme.

Working on the system is what changes those conditions. It is the work of examining the system's assumptions and asking whether they are accurate, identifying the gap between what the system says it does and what it actually does, and creating the conditions in which workers at every level can engage with safety as a real obligation rather than a compliance performance. The manager enforces rules; the leader adapts them to circumstances, understanding which rules serve their purpose in which contexts and which have drifted away from the conditions they were designed for. The manager controls the risks the existing system recognises; the leader goes looking for the ones it has not yet identified, because the system's blind spots are where the next serious event is developing. The manager gives clear instructions and

expects execution; the leader coaches, builds independent thinkers, and expects people to lead by example, because the safety decisions that matter most are made by workers with no instruction at all, in the moment of operational pressure, guided only by the values and assumptions the organisation has built into them.

These are not differences of style, personality or leadership philosophy. They are structural differences in what the two activities accomplish and in the outcomes each can produce. An organisation managed well by a safety manager will have a functional, compliant safety management system. An organisation led well by a safety leader will have a valid trajectory toward a transformational safety culture. The first is a necessary condition of the second. It is not a substitute for it. Organisations that have taken this seriously describe it in the same terms. Writing in 2008 about the Bristow Group's experience in offshore helicopter operations, Evans and Parker⁵³ argued that management and leadership are fundamentally different when it comes to safety, and that improving safety culture is work that sits alongside the safety management system rather than within it. Leadership, they concluded, was "the secret ingredient".

The Competency Gap

The research on management and leadership competencies in safety gives a clear picture of how differently the two are built. The competencies of safety management are working within existing systems, responding to events as they arise, enforcing established rules, following direction from above while giving clear instruction below, coordinating effort across teams, and providing the operational predictability that complex industrial organisations need. The competencies of safety leadership overlap only in part: working on systems, creating opportunities for improvement rather than waiting for problems, providing a vision of safety that people want to pursue rather than a set of instructions they follow under observation, and motivating people by connecting safety to what makes it meaningful to them rather than by the transactional exchange of compliance for reward.

The leadership competencies depend heavily on emotional intelligence, the capabilities examined in Chapter Eleven, and those are precisely the capabilities least developed in the technical and managerial pathways through which most safety leaders progress. The engineer who has spent a decade

building technical mastery and another decade building managerial competence has usually had few structured opportunities to develop them. The oversight is not malicious. It reflects an assumption, embedded in most professional development frameworks, that leadership capability is either innate or will develop through exposure to leadership responsibility: that it happens, as it were, by osmosis.

The osmosis assumption is one of the least examined beliefs in leadership development. Consider the typical career in industry. A professional in engineering, operations, finance or another technical discipline develops deep expertise, shows competence in managing people and processes within that domain, and is progressively promoted to broader responsibility. At each promotion, the implicit expectation is that the leadership capability needed at the next level will develop through the experience of operating there. The professional is given more responsibility and left to work out, by trial and error, how to discharge it. For some people, in some settings, this produces real leaders. For many, it produces managers with larger portfolios: people whose repertoire for engaging with the human and cultural dimensions of their organisations has not kept pace with the scope of their responsibility.

Warren Bennis⁵⁴, whose work on leadership development remains among the most practically useful available, was blunt about this. He called the belief that leaders are born the most dangerous myth in leadership, dismissed it as nonsense, and argued the opposite: leaders are made. He also observed that failing organisations are usually over-managed and under-led. The implication for safety is direct. The leadership capability safety culture requires is not a natural endowment some leaders have and others lack. It is a developed competency that needs as deliberate and sustained an investment as any technical skill, and in most organisations it receives nothing close to that investment.

The Path-to-Promotion Problem

The path-to-promotion problem is the specific mechanism through which the osmosis assumption reproduces the management default at every level of an organisation. In most industries, the route to senior leadership runs through technical and managerial competence in a particular domain. Engineers become engineering managers, then operations managers, then general managers. Safety professionals become safety managers, then safety

directors, then heads of health, safety and environment. The competencies assessed and rewarded at each stage are overwhelmingly technical and managerial: understanding and managing the domain, delivering results within existing systems, managing relationships and resources.

The competencies safety leadership requires — reading and influencing culture, creating the conditions for real engagement, challenging assumptions at the level of basic organisational belief, showing safety values in the moments of operational pressure that workers use to calibrate their own behaviour — are rarely assessed explicitly at any stage of that progression. They are assumed to develop. They rarely do, at least not at the pace and depth optimal safety culture requires.

The result is that organisations are consistently led by people whose safety competence lies mainly in management rather than leadership, promoted into positions where leadership is what safety needs. This is not a criticism of those people. It is a systemic failure of organisational design: a failure to take seriously, in leadership development and succession planning, what the leadership task of improving safety culture actually demands. Nobody builds a roof without checking that the numbers add up, because if the calculation is wrong, the roof falls. Safety culture development needs the same discipline: the calculation has to include the leadership competencies culture requires, not only the management competencies the existing system can assess.

Correcting the path-to-promotion problem needs deliberate intervention at the design level. Leadership development for safety has to stop being an add-on to technical and managerial development and become the central challenge of safety capability building, with the same resources, structure and sustained attention that technical training receives. The assessment of safety leadership capability — not safety management system knowledge, but the behaviours that create safety culture — has to become part of how people are selected for and assessed in leadership roles. Organisations have to accept that this kind of development takes time, needs specific expertise to design and deliver, and will not happen through exposure to responsibility alone.

What Safety Leadership Actually Requires

If safety leadership cannot be developed by osmosis, what does its development require? The research on leadership development, combined

with the demands of safety culture transformation, points to three capabilities that distinguish the safety leader's impact from the safety manager's.

The first is self-awareness: the ability to understand, accurately and without defensiveness, one's own safety leadership behaviour and its effect on the culture. This is not a soft skill or a personality trait. It is the precondition for every other safety leadership capability, because a leader who does not know how their behaviour is being received cannot know which of their signals are building the culture they intend and which are building something else. Jim Collins⁵⁵, researching the leaders who took companies from good to great, found that his team had initially been told to downplay the role of the leader, to avoid the simplistic credit-the-leader or blame-the-leader framing common in leadership writing. Despite those instructions, they found that the leaders who made the most significant and sustained difference combined fierce professional will with personal humility, and shared a particular habit of attribution: looking in the mirror when things went wrong and out the window when things went right. The self-aware leader takes personal accountability for the culture they have built, because they understand it is built through their behaviour, not their intentions.

The second is a model of human motivation more sophisticated than the transactional exchange underlying most safety compliance systems. Leaders who understand why people do what they do, at the level of basic psychological need rather than rational self-interest, can engage with safety in ways that connect to what workers value rather than to what the system requires them to perform. Safety, at its deepest level, is about protecting human life and human dignity. Leaders who speak about it in those terms, and connect safety requirements to what workers care about most — their families, their colleagues, their professional identity, their sense of agency and competence — produce a different quality of engagement from leaders who speak in the language of compliance obligation and performance management.

The third is credibility: the kind built only through visible consistency between stated values and actual behaviour. Nothing else in the safety leadership repertoire compensates for its absence. A leader who espouses safety values and then, under production pressure, shows a different set of values in practice has told the workforce something no later message can retract, because workers calibrate their behaviour on what they see the leader do, not on what the leader says. Building credibility is therefore a behaviour

challenge, not a communication one. Chapter Fifteen examines how it is built and how quickly it is lost.

Leadership as the Multiplier

Research on leadership and engagement gives a sobering account of what is at stake. Studies of the relationship between leadership effectiveness and employee engagement consistently find that the quality of a person's direct leader is among the strongest determinants of their engagement with their work, outweighing pay, working conditions and the quality of the organisation's strategy⁵⁶. Poor leaders create disengagement in people who were motivated when they started. Good leaders sustain it. Great leaders multiply it.

For safety, this means the leader is not one contributing factor among several. The leader is the multiplier: the element whose quality determines how well every other element performs. A high-quality safety management system in the hands of a safety manager produces compliant behaviour. The same system in the hands of a safety leader produces the discretionary contribution to safety, examined in Chapter Eleven as safety participation, that no system can mandate and no audit can measure, and that separates an adequate culture from an exceptional one.

The Transformational Safety® Culture Maturity Model, examined in full in Chapter Seventeen, reflects this directly. Movement from the Lethal level to the Developing level is achievable through management competence. Movement from Developing to Functional and Transformational is not. It requires the leadership capabilities this chapter describes, which is why safety culture transformation is a leadership problem rather than entirely a systems problem.

The man with twenty-two years of experience, the good metrics and the acceptable audits was not a poor safety manager. He was an excellent one. What his organisation needed from him, and had never asked of him, was the work no system can do on its own behalf: finding and closing its own gaps. That work belongs to leadership, deliberately and without delegation.

PLEASE RESPECT COPYRIGHT

CHAPTER NINE

The Full Range Leadership Model® and Its Safety Implications

There is no shortage of leadership frameworks. In the decades since leadership became a serious subject of academic inquiry, researchers and practitioners have produced a proliferation of models, taxonomies and competency frameworks describing the behaviours of effective leaders with varying degrees of precision and empirical grounding. Most are useful in a particular context, for a particular purpose, and most are made less useful than they might be by the lack of any rigorous basis for choosing between them.

The Full Range Leadership Model® is different. Its categories are not drawn from the intuitions of practitioners or the theoretical preferences of researchers. They come from systematic empirical investigation of how the entire range of leadership behaviour, from the most passive and avoidant to the most actively transformational, shapes the motivation, performance and wellbeing of the people within its influence. And its safety implications, examined directly in a growing body of safety-specific research, are among the most consistently demonstrated findings available to anyone who wants to understand the relationship between what leaders do and what their organisations' safety cultures become.

A Century of Research

The Full Range Leadership Model®, developed principally through the work of Bernard Bass and Bruce Avolio⁵⁷ synthesises more than a century of investigation into what leadership is and how it works. The early decades of

The Full Range Leadership Model® and Its Safety Implications

that investigation were dominated by the trait approach: the attempt to identify the innate characteristics that distinguish effective leaders from ineffective ones. It produced considerable insight into the personal characteristics associated with leadership effectiveness, but it could not account for how effectiveness varied across contexts, or for the evident fact that different situations call for different leadership responses.

Behavioural and situational theories added what the trait approach lacked: the understanding that effective leadership depends not only on who the leader is but on what the leader does, and that what works in one context may not work in another. These produced richer and more practically useful models, but they still lacked the empirical breadth to support confident conclusions about which behaviours, across the full range of leadership contexts, produce which outcomes.

The Full Range Leadership Model® addressed that gap by mapping the entire spectrum of leadership behaviour, from the most avoidant and disengaged to the most actively and visibly engaged, and examining the empirical relationship between each category and a broad set of organisational outcomes: effectiveness, extra effort from followers, job satisfaction and, crucially, relationships with higher levels of the organisation. The result is not a prescriptive framework telling leaders what to do in particular situations. It is an empirically grounded map of what different categories of leadership behaviour produce across a wide range of contexts, and therefore a useful basis for any leader who wants to understand how their behaviour is shaping their organisation's culture and performance.

For safety leadership, the model offers something of exceptional value: a framework that connects observable leadership behaviours to measurable organisational outcomes. It turns the abstract idea of safety culture into the concrete question of what leaders actually do, and what follows from it. The behaviours it identifies are not theoretical ideals. They are descriptions of actual leadership practice, drawn from the observation and measurement of real leaders in real organisations, and their relationship to safety outcomes has been examined in safety-specific research with results that should reshape how every safety leader understands their role.

Julian Barling, Catherine Loughlin and Kevin Kelloway⁵⁸ showed in 2002 that safety-specific transformational leadership predicted lower occupational injury rates, working through the safety climate and safety consciousness it created. Dov Zohar's⁵⁹ research on group-level safety climate showed how

strongly supervisors' behaviour shapes the climate of their own teams. Reviewing the industrial evidence, Rhona Flin and Steven Yule⁶⁰ concluded that transformational leadership behaviours such as intellectual stimulation and individualised consideration, together with contingent reward, were associated with stronger safety performance

The Passive-Avoidant Spectrum

At one end of the range sits *laissez-faire* leadership: the most passive form and, empirically, the most harmful. *Laissez-faire* leadership is not simply relaxed or non-directive. It is the consistent avoidance of leadership responsibility: failing to make decisions when they are required, to respond to concerns when they are raised, to give direction when it is needed, and to engage with safety issues as they develop. The *laissez-faire* leader is, in effect, absent: physically present in the organisation but functionally disengaged from the responsibilities their position carries.

The consequences for safety culture are severe and well established. *Laissez-faire* leadership is negatively correlated with extra effort from followers. People who watch their leader avoid responsibility do not increase their own discretionary effort; they reduce it to whatever level the leader's signals suggest is appropriate⁶¹. Its correlation with organisational effectiveness is among the lowest across the full range, and its correlation with job satisfaction is negative. In safety terms, this produces the cultures described in Chapter Six: cultures in which concerns are not raised because experience has shown that raising them produces no response, in which near-miss reporting decays to the compliance minimum because the information is never acted on, and in which the distance between procedure and practice widens undetected because the leadership that should detect it has disengaged from the operational reality it is supposed to understand.

Laissez-faire safety leadership is rarely the product of indifference. More often it is the product of the competency gap described in Chapter Eight: leaders promoted into positions that require safety leadership capability, who lack it, and who manage the discomfort of that gap by avoiding the situations in which it is most exposed. The leader who does not visit the worksite is rarely someone who does not care what happens there. More often they do not know how to engage productively with what they would find if they did, and have learned, through a history of unrewarding or uncomfortable worksite

encounters, that their time is better spent on the administrative work at which they are competent.

Passive management by exception sits a step above laissez-faire but shares many of its consequences. The passive management-by-exception leader monitors the system for evidence of deviation — reviewing metrics, examining reports, conducting reviews — and responds when deviations become apparent. Its passivity lies in its orientation: backward-looking, deficit-focused and entirely reactive. This leader never engages with safety conditions before they produce evidence of failure. They wait for the evidence, then respond. In most environments, by the time the evidence is clear enough to prompt a response, the conditions behind it may have been accumulating for months or years, and the response may be too late to prevent the event that makes them visible.

Active management by exception improves on the passive form by monitoring more closely and intervening earlier: setting clearer standards, checking more systematically for deviations, and stepping in sooner. It produces better outcomes than laissez-faire or passive management by exception, but it remains transactional at heart. It defines success as the absence of deviation from standard rather than the presence of the cultural conditions that make safe behaviour natural. In safety culture terms, active management by exception produces the Developing level: a system-driven culture that manages visible compliance effectively without addressing the deeper conditions that determine what workers do when no one is monitoring.

The Transactional Range

Contingent reward is the most effective of the transactional behaviours and the one that most clearly bridges the transactional and transformational ranges. The contingent reward leader works through an explicit or implicit exchange of reward for performance: do what is required and you will receive what you value. In safety terms, contingent reward produces safety compliance: the core safety activities the system requires, examined in Chapter Eleven.

Contingent reward is not without value. Safety compliance is a necessary foundation of safe operations, and leaders who consistently and fairly exchange recognition for safe performance get better compliance than those who manage through correction and discipline alone. The research shows

positive correlations between contingent reward and extra effort from followers, and with both unit and organisational effectiveness. Within the full range, it is among the most reliably positive of the transactional behaviours^{62,63}.

Its limitation is structural rather than a matter of execution. A contingent reward culture is, by definition, a compliance culture: people do what is required because the exchange makes it rational, not because they have internalised the values that make safety matter to them. The worker in a contingent reward safety culture does the right thing when the reward is present and the monitoring is visible. What they do when those conditions are absent is governed by a different calculation, because the exchange has not created the internal motivation that makes safe behaviour consistent and independent of external incentive.

Dealing is, as the practitioner literature on safety leadership has observed, the basis on which most safety management systems are built. The system offers workers something they value — continued employment, a clean safety record, freedom from the uncomfortable consequences of non-compliance — in exchange for the safety behaviours it requires. The deal works, within its limits. It works for compliance. It does not produce the discretionary engagement that exceptional cultures depend on, and it cannot create the conditions in which a worker would stop a job, raise a concern, or challenge a supervisor on safety grounds without the backing of the exchange.

The transactional range, from laissez-faire through management by exception to contingent reward, describes with considerable accuracy the majority of safety leadership practice in most organisations. It is the range within which safety management, properly executed, operates. And it is the range within which, as the research consistently shows, safety culture outcomes are capped at compliance and average performance: necessary, but insufficient for exceptional safety performance.

The Transformational Dimension

Transformational behaviours operate at a different level from transactional ones. Where transactional leadership manages the exchange between leader and follower, reward for performance and correction for deviation, transformational leadership changes the follower's motivation itself. It does not offer better incentives for the existing motivation. It creates a new kind: commitment to shared values, identification with a compelling vision,

The Full Range Leadership Model® and Its Safety Implications

belief in the importance of the collective effort, and an engagement with the leader's goals that produces effort and performance beyond anything a transactional exchange could.

Inspirational motivation is the transformational behaviour that most directly addresses the challenge of safety culture. The inspirationally motivating safety leader speaks about safety in terms that connect to what people care about: not compliance obligations and performance metrics, but the shared values of a professional community that believes its work matters and that the safety of everyone who does it is an expression of those values. They describe what safe operations look like, not as a statistical abstraction or a regulatory requirement, but as a lived reality in which every worker goes home to the people they love. Workers can tell the difference between that conviction and a performance of it, because they watch what the leader decides when the pressure comes. Inspirational motivation that builds safety culture is not rhetorical skill. It is the authority that consistent behaviour earns.

Intellectual stimulation is the transformational behaviour that addresses the adaptive deficit of the Developing culture. The intellectually stimulating safety leader actively challenges the assumptions of the existing safety system, not to undermine it but to keep it accurate. They regularly ask whether the procedures reflect how the work is actually done. They create forums in which workers can describe what the system does not see, and they treat those descriptions as diagnostic information rather than as deviations. They encourage their teams to find better ways of managing the risks they face, rather than treating the current approach as the final word. And through their own willingness to question received wisdom, they show that the culture they are building values critical thinking about safety rather than suppressing it.

Individualised consideration is the transformational behaviour that most directly reaches the human dimension of safety culture, the dimension management by exception consistently misses. The leader who practises it knows the safety concerns of each person on their team as individual concerns, not as data points in an aggregate picture. They know which workers are managing personal circumstances that affect their safety, which are working with equipment or in conditions that create particular exposure, and which have concerns they have not yet felt comfortable raising through formal channels. They create the conditions for real one-to-one conversations about safety, turning the leader–follower relationship from a management relationship into a developmental one.

Idealised influence, in both the attributes the leader shows and the behaviours they model, describes how leaders transmit values through example rather than instruction, and for that reason it bears most directly on safety culture. The leader who shows idealised influence in safety walks the talk: their safety behaviour matches their stated values in every context workers observe, including the inconvenient, costly or professionally uncomfortable ones. They do not merely advocate safety; they practise it, visibly, creating the model the workforce uses to calibrate its own behaviour. And they speak about their commitment to safety as personal belief rather than professional obligation, making the culture's values their own and showing, through that ownership, that the values are real.

What the Numbers Actually Say

The most striking figures come from a 1996 study by William Brown, Elizabeth Birnstihl and Daniel Wheeler⁶⁴, of work groups in the United States Cooperative Extension Service. Using an earlier, four-factor version of the Multifactor Leadership Questionnaire, they found correlations with followers' extra effort of 0.95 for idealised influence, 0.98 for inspirational motivation, 0.92 for intellectual stimulation and 0.97 for individualised consideration.

Those figures should be read for what they are: the results of a single study, in one setting, and considerably higher than most leadership research reports. The wider evidence tells the same story in more measured terms. Timothy Judge and Ronald Piccolo's 2004 meta-analysis⁶², which combined the results of many studies across many settings, found transformational leadership to be the strongest predictor of leadership effectiveness across the full range, followed closely by contingent reward, with laissez-faire leadership negatively related to almost every outcome. The relationships it reports are more modest than in the 1996 study. The order is the same, and for safety leaders it is the order that matters.

The contrast with the transactional and passive-avoidant behaviours is just as instructive. In the 1996 study, management by exception showed no relationship at all with followers' extra effort, and almost none with organisational effectiveness, although it had a modest positive relationship with job satisfaction: some people are content to be managed that way, even when it achieves little. Laissez-faire leadership was negatively related to extra effort and, most strongly, to organisational effectiveness. The data describes,

The Full Range Leadership Model® and Its Safety Implications

in numbers, the difference between leadership that builds engagement and leadership that erodes it.

For safety leaders, what matters is less the size of any single correlation than what the pattern reveals about how leadership shapes safety outcomes. The mechanism is not instruction or enforcement: management by exception, the category most closely tied to monitoring and enforcement, produces the weakest positive outcomes in the transactional range. It is not incentive: contingent reward, the most explicitly incentive-based category, produces outcomes that are positive but bounded. The mechanism is motivation at the level of values and identity: a change in what workers care about, and therefore in what they will do without instruction, enforcement or incentive.

The safety implications follow directly. A leader whose behaviour sits mainly in the laissez-faire and management-by-exception range will produce variable compliance: workers comply when monitoring is present and incentives are active, and not otherwise. A leader whose behaviour sits mainly in the contingent reward range will produce reliable compliance: workers do what is required because the exchange makes it rational. A leader whose behaviour sits mainly in the transformational range will produce something different in kind: safety participation, in which workers at every level contribute beyond the minimum because they have made the values behind safe behaviour their own.

From Research to Practice

Applying the Full Range Leadership Model® to leadership development starts with recognising that the model is a map, not a prescription. It describes what different categories of behaviour produce, not what any particular leader should do in any particular situation. Its practical use is not to push every leader to the transformational end of the range regardless of context. It is to build conscious awareness of where a leader's behaviour currently sits, what that behaviour is producing in the culture, and what shifts would produce the outcomes the safety culture needs.

The ultimate goal is for leaders to be able to operate across the full range, using different behaviours in different contexts and in response to different demands. The evidence consistently shows that leaders whose dominant profile sits in the transformational range, while retaining the ability to use transactional behaviours when the context calls for them, produce the

Beyond Compliance

strongest safety outcomes^{65,66}. The combination is more effective than either alone: transformational behaviour builds the values and engagement that drive discretionary participation, and transactional behaviour maintains the compliance foundation on which participation builds.

Developing that capability requires, at a minimum, three things that most safety leaders' current development pathways do not provide. The first is precise self-awareness about one's current leadership profile: not a general sense of leadership style, but an accurate understanding of which behaviours in the full range are dominant, which are underused, and what the current profile is doing to the safety culture. The second is the deliberate development of the behaviours that are most underused, through practice in real leadership situations rather than in the training room alone. The third is sustained accountability for cultural outcomes, not only for safety management metrics, so that developing leadership behaviour feels consequential rather than optional.

The Transformational Safety® Leadership framework translates the Full Range Leadership Model® into the constructs and behaviours of safety leadership practice, and provides the structure for that development. It begins with a diagnostic map of where safety leaders actually sit across the full range: not where they believe they sit, and not where their organisations have rated them, but where the workers they lead actually experience them.

PLEASE RESPECT COPYRIGHT

CHAPTER TEN

The Transformational Safety® Leadership Constructs

The Transformational Safety® Leadership framework organises the full range of safety leadership behaviour into nine constructs that describe, with increasing specificity, how leaders actually show up in the safety cultures they inhabit and shape. The constructs are not personality types. They are behavioural descriptions: accounts of what particular patterns of leadership behaviour look like in practice, what those patterns produce in the safety culture around them, and what distinguishes each from its neighbours on the leadership spectrum.

The constructs are not a parallel system to the Full Range Leadership Model® examined in Chapter Nine. They are its safety-specific expression, each one a translation of a Full Range behaviour category into the particular demands, pressures and relationships of safety leadership. The Invisible Man is laissez-faire in operational terms. The Fireman and The Policeman occupy management-by-exception territory, passive and active respectively, reframed around the dynamics of safety monitoring and enforcement. The Dealer is contingent reward as safety leaders actually encounter it: the points system, the recognition programme, the implicit exchange that produces compliance and nothing beyond it. The five transformational constructs — The Motivator, The Knight, The Thinker, The Carer and The Missionary— correspond to inspirational motivation, idealised attributes, intellectual stimulation, individualised consideration and idealised behaviours, each recast in the language of what those behaviours look like when the domain is safety and the stakes are the lives of the people being led. The research in Chapter Nine is not background to what follows. It is the grounding of every construct in this chapter.

The Invisible Man®

In safety, invisibility is not a leadership trait to aspire to. It is among the most dangerous conditions a safety culture can experience, and it is more common than the organisations experiencing it are usually willing to admit. The Invisible Man construct describes the safety leader who is functionally absent from the safety life of their organisation: rarely at the worksite, rarely engaged with the operational safety questions their team faces, rarely available when a concern needs the authority of leadership to resolve it, and rarely visible in the moments that would tell workers that safety matters to the person responsible for the conditions in which safety decisions are made.

The Invisible Man is rarely absent on purpose. More often, the invisibility is the product of the management default described in Chapter Eight: a leader whose competence and confidence sit in the administrative side of the role — reviewing safety metrics, approving corrective action registers, conducting formal safety reviews — and who has never developed the capability or the habit of engaging with safety at the operational level where it actually lives. The safety briefing is read. The monthly safety report is reviewed. The corrective action register is signed off. The worksite, where the safety culture the leader is responsible for is actually experienced, goes unvisited. Not because the leader does not care about it, but because they do not know how to engage with it productively.

The cultural consequence is the slow erosion of workers' confidence that senior leadership is interested in the safety reality they navigate. When workers see, time after time, that their leader appears only when something has gone wrong, and is absent through the long ordinary periods when the work of building safety culture is actually done, they adjust their expectations. Near-misses are not raised informally, because there is no informal channel to leadership. Concerns that need leadership support to resolve are handled within the team, at whatever level of resolution is available there. The culture develops the self-sufficiency of an organisation that has learned not to rely on its leaders, which is a very different thing from the empowered independence a transformational safety culture produces.

Leaders who recognise the Invisible Man in their own behaviour have a practical and manageable development path: consistent, purposeful, well-prepared presence at the worksite. Not surveillance. Not inspection. Engagement: asking real questions about how work is actually done, listening

to answers that may be uncomfortable, and showing, through that engagement, that the leader's interest in safety goes beyond the data in formal reports. Regularity matters as much as quality. Workers judge a leader's interest by accumulated experience, and a single highly visible visit does not undo the signal that a long absence has already sent.

The Fireman[©]

The Fireman is present — visibly, urgently, energetically — when something goes wrong. They respond to incidents quickly and with real engagement. They are at the scene of a near-miss. They are accessible and decisive in the immediate aftermath of an event that needs a leadership response. Their commitment to safety in moments of crisis is real and unambiguous.

Their limitation is the one the metaphor implies: they are oriented toward fire, not fire prevention. They sit in the organisational equivalent of the firehouse, busy with the administrative and managerial side of their role and ready when the alarm sounds, but absent from the ongoing, unglamorous work of building the conditions that would make the alarm sound less often. The issues they deal with after an event have often been present in the operational environment for months, issues that a leader engaged with the worksite might have found and resolved before they became events requiring a firefighter.

The cultural consequence is a particular learned expectation. Workers understand that leadership attention is available after events but not before them, and they manage their concerns accordingly. Issues below the threshold of an event are unlikely to get leadership attention, so they are either handled within the team or left to persist as tolerated risk: normalised deviance, in Vaughan's terms, accumulating in the gap between the leader's event-focused attention and the operational reality the team works in. The Fireman fights fires with real commitment. What they do not do is notice the conditions that let the same fires start.

Moving beyond The Fireman means deliberately shifting attention from reactive to proactive: from events that have happened to conditions that are developing. In practice, it means the same worksite presence that addresses the Invisible Man, directed explicitly at conditions and concerns that have not yet become events. What has this team been living with that has never been formally raised? What adaptations to the standard procedure have developed

because the procedure does not match the work? What is the team uneasy about that has not reached the threshold of a report? Asked consistently, and received with real interest, these questions begin to turn the leader from someone who responds into someone who prevents.

The Policeman®

The Policeman is the safety leader whose primary orientation is compliance: detecting deviations from the established rules and correcting them through enforcement. The Policeman's attention is on breaches: whether procedures are followed, PPE worn, permits properly completed, and the administrative requirements of the safety management system met. They monitor. They inspect. They correct. When they find deviations, they respond with the authority their position gives them.

The Policeman is not without value. Some monitoring, inspection and enforcement is necessary in any safety system. The compliance foundation that transactional leadership provides is a real contribution to safety performance, and a leader entirely unwilling to enforce safety standards is abdicating a real responsibility. Organisations need leaders who can make sure safety requirements are understood and met, and the Policeman, within appropriate bounds, does that well.

The limitation shows when this orientation dominates. In strongly autocratic policing cultures, where the Policeman style is widespread and deeply embedded, workers comply with safety requirements when they feel watched and deviate when they do not. The compliance is real but conditional: produced by external enforcement rather than internal commitment, and gone when the pressure is removed. More seriously, a Policeman culture suppresses the initiative and independent thinking that safety leadership should be cultivating. Workers in a culture of surveillance learn to comply visibly and to think mainly about whether their behaviour will pass inspection, not about whether the conditions they are working in are safe. Safety becomes a performance directed at the Policeman® rather than an orientation toward protecting people.

Moving beyond the Policeman means deliberately adding inspiration to enforcement: moving from monitoring for deviations to engaging workers in how to make safety better, and from detecting non-compliance to building the values that make compliance natural. The Policeman who keeps the capacity

for appropriate enforcement while developing the transformational behaviours above it on the spectrum becomes an effective safety leader. The one who does not stays confined to a compliance culture that will never deliver the discretionary engagement safety needs.

The Dealer[©]

The Dealer is contingent reward in practice, and Chapter Nine has already described the basic deal: the safety system offers workers something they value in exchange for the behaviour it requires. The Dealer safety leader works through that exchange, rewarding safe performance and attaching consequences to unsafe performance. Safety points systems, recognition programmes, bonus structures linked to safety metrics, and the many formal and informal reward-for-compliance mechanisms across the safety management landscape are all the Dealer at work.

The Dealer is not a cynical archetype. The exchange is real and produces real results: more consistent compliance, higher formal safety engagement, and better metrics in the categories the exchange is designed to improve. The Dealer understands motivation at the level of self-interest and rational exchange, and uses that understanding effectively. Where a safety culture is not yet developed enough to sustain compliance through intrinsic motivation, the Dealer's transactional scaffolding supplies the external structure compliance needs while the deeper conditions are built.

The ceiling is the one Chapter Nine described: exchange produces compliance, not commitment. Its particular vulnerability is change. When the reward is withdrawn, or when the cost of safe behaviour exceeds the value of the reward it would earn, the compliance the Dealer has bought evaporates. A Dealer who does not add the behaviours that build intrinsic motivation is running a safety culture on borrowed time, and will find, when conditions change, that they have built a structure without a foundation.

Like policing, dealing can also be overdone. When everything becomes an exchange, the culture turns mercenary: people begin to ask what each safety contribution will earn them, and the contributions that earn nothing quietly stop. The culture also becomes demarcated. People do what they are rewarded for and no more, and the boundaries between roles harden. Those boundaries matter for safety, because hazards do not respect them. The condition that sits

between two people's responsibilities, noticed by both and owned by neither, is exactly the kind of hole a demarcated culture leaves open.

The development challenge for the Dealer is to widen their motivational repertoire beyond the transactional: to connect safety to values and identity, to inspire rather than incentivise, and to create the conditions in which workers want to contribute to safety beyond what the exchange requires. That does not mean abandoning the transactional tools. It means adding what the Dealer's repertoire currently lacks.

The Motivator®

The Motivator is the first of the transformational constructs: the first profile oriented toward changing what workers care about rather than managing what they do in response to the motivations they already have. The Motivator is dynamic in how they communicate about safety. They connect it to what workers value, describe a vision of safe operations that people find compelling, and engage with safety in a way that is infectious, generating an energy the administrative safety culture cannot produce.

The Motivator operates at the level of inspirational motivation, the behaviour Chapter Nine showed to be most strongly associated with extra effort from followers. In practice, that association looks like workers who, because their leader has connected safety to what they care about, contribute beyond the minimum compliance requires and beyond what any incentive could buy. They raise concerns that have never been formally reported, because raising concerns now feels natural and welcome. They coach colleagues on safe practice, because they have made that practice part of their professional identity. They challenge unsafe conditions even when it is uncomfortable, because the Motivator® has built the conviction that safe operations are worth protecting.

The Motivator does not simply talk about safety with enthusiasm. Enthusiasm without substance produces the short-term spike in engagement and long-term cynicism that safety communication campaigns so often generate. What distinguishes the real Motivator is that the enthusiasm rests on credible engagement with operational reality: a detailed understanding of the risks the team faces, a real relationship with the people who face them, and a record of decisions that confirms the stated commitment. The Motivator's

energy about safety is believed because it is consistent with everything else people see about how they lead.

The Knight[®]

The Knight is the safety leader who models the values they espouse, and who walks the talk in the observable, consequential ways that build the credibility safety culture depends on. In Full Range terms, the Knight aligns with Idealised Attributes: the determination, integrity and sense of purpose that followers see in a leader they trust, experienced as both an aspiration and a standard.

Knights have historically been known for honour, ethics and the protection of those in their care. The safety Knight shows exactly these qualities: a deep personal commitment to acting consistently with stated safety values, a willingness to make decisions that are personally costly because they are consistent with those values, and a real orientation toward protecting every person within their sphere of responsibility. The Knight is the second of the transformational constructs and, in many ways, the most foundational, because without the credibility the Knight provides, the Motivator's energy risks becoming rhetoric.

The Knight's most important contribution is what they do when it is difficult. Any leader can show safety commitment when safety and convenience point the same way. The Knight's credibility is built when they do not: when stopping a job for a safety concern will delay the schedule, when raising a concern will make a hard meeting harder, when backing a worker who has reported a near-miss will create political discomfort. In those moments the Knight does the difficult thing, consistently, and the culture watches. What it sees is how the culture's working assumption about safety leadership is formed.

Knights seek honest feedback about their own safety behaviour. They share personal stories and convictions about safety in ways that make the culture personal rather than institutional. They acknowledge their own errors and what they learned from them. And they involve workers in safety decisions in a way that goes beyond the lip service "consultation" too often represents in safety contexts. The Knight does not involve workers as a communication exercise. They do it because the people who do the work have safety

knowledge leadership cannot get through any other channel, and using that knowledge is both an ethical and a practical requirement.

The Thinker®

The Thinker is the construct that most directly addresses the loss of adaptive capacity that Chapter Seven identified as the mark of strong but rigid cultures. The Thinker actively promotes a culture of learning, at both individual and group level, in which continuously improving safety knowledge and practice is treated as a leadership priority rather than an administrative obligation.

Innovation in safety is about thinking. It is the encouraged, culturally supported questioning of existing safety practice and the search for better practice. Not in ways that undermine the compliance foundation the system provides, but in ways that improve the fit between the system's design and the operational reality it is meant to address. The Thinker regularly challenges their team to develop solutions to safety issues themselves, or with colleagues, rather than waiting for instruction from above or a corrective action from the safety function. They trust in the intellectual development of their people, and show it by creating forums for real safety thinking rather than for performances of safety compliance.

The Thinker's particular contribution is the message it sends about who owns safety knowledge. Where the safety function is the main source of safety intelligence, and improvements are developed by experts and handed to workers to implement, workers see their role as receivers and implementers. Where the Thinker's style is present, workers understand that their experience of the safety reality of their work is valued, that their ideas for doing it more safely are sought, and that the organisation's safety knowledge is built by everyone who works in it, not by a specialist function alone.

That understanding changes what workers do. That understanding changes what workers do. In teams led by a Thinker, workers are more likely to identify hazards no one specifically trained them to recognise, because thinking about safety, rather than simply complying with requirements, has become a habit. They are more likely to propose improvements to procedures that do not reflect how the work is actually done, because they understand the organisation wants that gap closed. And they are more likely to share safety

learning with colleagues, because the culture treats safety knowledge as a shared resource rather than a private compliance obligation.

The Carer[®]

The Carer shows the individualised consideration that the Full Range Leadership Model identifies as one of the most powerfully transformational leadership behaviours. The Carer knows the people they lead as individuals: not as positions on an organisational chart or as performers against safety metrics, but as people with particular circumstances, concerns, vulnerabilities and aspirations. And they engage with those people in a way that makes that knowledge felt, communicating through the quality of the relationship, rather than through formal safety communications, that they care about the safety and wellbeing of each person in their charge.

The Carer should not be mistaken for softness. Caring safety leaders are trusted precisely because workers know that a leader who cares about their wellbeing will tell them the truth about safety. They will not play down a risk to avoid discomfort, will not go along with the comfortable consensus of a team that has normalised a hazardous condition, and will not let a concern go unaddressed because addressing it means a difficult conversation. The Carer's credibility comes from the combination of concern and honesty: the experience of a leader who cares enough to be straight with you.

The cultural consequence is psychological safety of the practical kind described in Chapter Six: not an abstract sense that the organisation values openness, but the experience of workers who have raised a concern with this leader and found it received with interest, thanks and follow-through. Once established, that experience changes the reporting culture in ways no formal near-miss system can replicate. Workers who trust their leader report concerns in the informal, ongoing, conversational way that gives the organisation its best leading indicator of risk: real-time intelligence from a workforce that feels valued and is therefore engaged.

The Carer does not play favourites. Their individualised consideration is equitable: workers at every level feel their safety concerns are heard and their contribution to the team's safety is valued. That equity matters in cultures where leadership attention is distributed by hierarchy, where senior workers get more attention than junior ones and experienced workers are engaged while newer ones are managed. The Carer's[®] even-handedness communicates

that safety is everyone's obligation and everyone's right, regardless of tenure, seniority or performance.

The Missionary®

The Missionary is the apex of the Transformational Safety® Leadership constructs. It combines the credibility of the Knight, the energy of the Motivator, the curiosity of the Thinker and the relational depth of the Carer into an orientation toward safety culture transformation that works at the level of the organisation's deepest values rather than its visible behaviours. The Missionary safety leader is not merely committed to safety. They are defined by it: their professional identity is so fully integrated with their safety values that the distinction between their leadership and their safety leadership has dissolved.

In Full Range terms, the Missionary is Idealised Behaviours: the leader whose sense of mission is expressed so consistently that it produces unusually high levels of trust and commitment in the people around them, even in the most demanding circumstances.

The Missionary's impact is the most lasting of all the constructs, because it works at the level of basic assumption rather than artefact or espoused value. The Missionary® does more than model safe behaviour, communicate safety values or engage with individual concerns. They create the conditions in which the whole culture's relationship to safety shifts, in which the underlying assumptions about what safety is, why it matters and who is responsible for it move gradually toward the Transformational level of the Transformational Safety® Culture Maturity Model.

Missionary leadership involves everything the constructs below it require, and adds a particular kind of conviction: a personal belief, visible in every interaction, decision and relationship, that safety is not a management obligation or a regulatory requirement but an expression of what the organisation is and what it is for. That conviction is not manufactured or performed. It is the product of leaders who have, over the course of their development, taken safety into their professional identity, moving from compliance through engagement and commitment to the point where safety is simply who they are.

The path to the Missionary is not a training programme. It is a developmental journey that requires the sustained accumulation of the habits,

relationships and self-awareness each construct below it represents. Leaders who aspire to Missionary leadership do not start there. They start where they are: with an honest assessment of their current construct profile, a clear understanding of what developing each construct requires, and the disciplined commitment to the behaviour the culture they are building needs.

Reading the Constructs as a System

The nine constructs are most useful not as individual descriptions but as a system: a map of the whole range of safety leadership behaviour that lets leaders locate themselves, understand what their current profile is producing in their culture, and identify the development priorities that will move them toward the constructs that matter most for safety culture transformation.

The four constructs below the transition from transactional to transformational are not equal. The Invisible Man and the Fireman have no value to a safety culture. One is absence; the other is attention that arrives only after the harm is done. The development task is to leave both behind. The Policeman and the Dealer are different. They describe behaviours every safety leader needs to some degree, the enforcement of standards and the transactional scaffolding compliance requires, and the challenge is not to eliminate them but to supplement them with the constructs above: the Motivator, the Knight, the Thinker, the Carer and the Missionary, which produce the discretionary engagement and cultural depth safety leadership needs.

Where a leader sits across the nine is not a matter of self-report. The most accurate measure is the experience of the people they lead: the picture of the leader's behaviour that workers have built from watching real decisions under real pressure. That is why the key individual 'Transformational Safety® Leadership assessment is a 360-degree instrument, setting followers' perceptions alongside the leader's own. A leader who believes they are operating as a Knight while their people experience a Policeman has learned something important, and Chapter Seventeen examines what that gap reveals and what to do with it.

PLEASE RESPECT COPYRIGHT

CHAPTER ELEVEN

Leadership, Engagement, and Discretionary Effort

The Towers Perrin global workforce study⁶⁷, one of the largest studies of employee engagement undertaken, covering more than 85,000 people in large and mid-sized organisations across sixteen countries on four continents, produced a finding so consequential, and so consistently ignored by mainstream safety management, that it deserves to be stated plainly at the outset. Only fourteen per cent of the workforce was highly engaged in their work. Sixty-two per cent was moderately engaged. A quarter was genuinely disengaged: not merely uninvested, but in a state of active disconnection from the organisation and its goals.

For safety leaders, these numbers carry an implication almost no safety management system is designed to address. If the workforce engages with safety as it engages with work in general, then three workers in four are contributing to safety at less than their available capacity, and one in four is disengaged in a way that makes their relationship to the organisation's safety culture unreliable at best and harmful at worst. A safety management system that assumes a uniformly motivated workforce, working in good faith within its designed boundaries, is not describing the workforce that actually exists. It is describing an idealised one that the culture of most organisations has never created.

The central question is not whether this is acceptable. It is not. The question is why it persists in organisations that have invested heavily in safety culture, and what leadership behaviour is needed to change it. The engagement research answers both questions, consistently, across industries, cultures and organisations of every size. Safety leaders ignore those answers at their peril.

The Engagement Reality

Examined in more detail, the engagement data reveals a pattern more troubling than the aggregate figures suggest. Engagement is not evenly distributed across organisational levels. Senior executives are the most likely to be highly engaged, and the least likely to be disengaged. Moving down the hierarchy, the proportion of highly engaged workers falls, the moderately engaged proportion stays relatively constant, and the disengaged proportion rises. It is highest at the labourer and plant operator level: among the people closest to the physical work that safety systems are designed to protect.

This gradient is not accidental. It follows the leadership gradient that accompanies it: senior leaders are more likely to work for leaders whose behaviour generates engagement, and less likely to experience the management practices that create disengagement, than frontline workers are. As Chapter Eight noted, the research is clear that the quality of a worker's direct manager or supervisor is among the strongest determinants of their engagement. Not the organisation's safety culture. Not the quality of the safety management system. Not the senior team's commitment to safety. The person who manages and leads them every day.

The safety implication is direct and uncomfortable. The engagement deficit is concentrated exactly where safety performance matters most: at the point where the safety management system meets the physical work it is supposed to govern, and where frontline workers make the operational decisions the system's risk controls are designed to shape. A workforce disengaged at the operational level has withdrawn its discretionary contribution to the organisation's goals, and in safety, the discretionary contribution is not a luxury. It is the difference between a culture that identifies hazards before they become incidents and one that does not; between one that reports near-misses and one that does not; between one that challenges unsafe conditions and one that tolerates them^{68,69}.

The research also says something important about the direction of causation. Workers do not become disengaged because the work is inherently unrewarding or because they lack motivation. The Towers Perrin study found that most workers want to make a bigger contribution to their organisations. The potential for higher engagement exists, and the factor most consistently responsible for failing to realise it is the behaviour of managers and leaders who, usually without intending to, put obstacles in the way of the engagement

their workforce is offering. The disengagement is not an individual failure. It is a leadership failure.

What Engagement Actually Measures

Engagement, in the research that has most influenced our understanding of its effect on organisational outcomes, is not a single thing. It has two distinct components, which interact in ways that matter for safety and call for different leadership responses.

Rational engagement is the cognitive component: how well workers understand how their work contributes to the organisation's goals, how their unit's performance relates to its overall direction, and how their own effort connects to meaningful outcomes. Workers who are rationally engaged know why their safety behaviour matters. They understand the relationship between what they do and what the organisation is trying to achieve, and they have enough information about its safety goals to direct their effort well. Rational engagement in safety is not the same as technical knowledge of safety requirements. It is a deeper understanding of purpose: why the requirements exist, how they connect to real risk, and what their consistent practice produces across the whole culture.

Emotional engagement is the affective component: how proud workers are to be part of the organisation, how far they feel it brings out their best work, and how much they care about its future. Emotionally engaged workers are the ones who would recommend their organisation as a safe place to work, not because they are asked to, but because their experience of its safety culture supports it. They are personally invested in their team's safety performance, and they experience that investment not as a professional obligation but as an expression of who they are and what they care about.

Safety management systems are built to address rational engagement alone, through the information, procedures, training and metrics that give workers a cognitive framework for what safe performance requires. What they are not built to address, and what safety leadership must, is emotional engagement: the sense of belonging, pride and personal investment that makes a safety culture self-sustaining rather than dependent on the continuous administration of compliance incentives.

The research on the main drivers of engagement shows what produces each. The rational drivers include clarity of role and contribution, the integrity

of senior management, input into decisions, and being managed by someone who connects one's contribution to the organisation's goals. The emotional drivers include challenge and growth in the work, real care from the manager, recognition and appreciation, and belonging to an organisation whose direction and values one can support. Every one of these drivers can be addressed through the Transformational Safety® Leadership constructs described in Chapter Ten, and everyone has a direct equivalent in safety culture that leadership is uniquely placed to create.

The Sins of Omission

Management's most consequential failures in safety engagement are, with remarkable consistency, sins of omission rather than commission. They are not the result of leaders actively damaging their workforce's engagement. They are the result of leaders failing to do what engagement requires: failing to recognise, to connect, to include, to develop, and to say the things workers most need to hear from the people responsible for leading them.

The research on recognition and appreciation is among the most consistent in organisational psychology⁷⁰. Employees who receive meaningful recognition are substantially more engaged and less likely to leave than those who do not. A competitive pay and benefits package gets workers through the door. Recognition and appreciation are what help engage them once they are inside.

In safety, the omission of recognition is particularly costly, because the things that most need recognising in a safety culture are exactly the things that are hardest to see and easiest to take for granted. A worker who stops a job because a safety condition is not met, at real cost to the schedule and with real personal courage, has done something that deserves acknowledgement. Not because acknowledgement was what they were after, but because its absence tells them, precisely, that the organisation does not regard what they did as significant. Workers who receive that message adjust their future behaviour to match.

The pattern of omission extends beyond recognition to a cluster of behaviours the engagement research consistently identifies as powerful generators of disengagement. Not providing a vision, never explaining why the safety culture matters or where it is trying to go, leaves workers without the sense of purpose emotional engagement needs. Saying important things only

once, at the annual safety day, without reinforcing them through the daily, informal channels where culture actually changes, makes the message look like performance rather than conviction. Not holding people accountable, and so treating safety compliance as optional, signals that the stated commitment is not real. And not spending enough one-to-one time with workers on safety — on their individual concerns, their learning needs and their experience of the work — is the costliest omission of all. It is the absence of the individualised consideration the Carer represents, and the engagement research identifies that consideration as the primary driver of real engagement.

Safety Compliance Is Not Safety Participation

The distinction between safety compliance and safety participation, first drawn by the Australian researchers Mark Griffin and Andrew Neal⁷¹, has appeared briefly in earlier chapters. This is where it belongs in full. Safety compliance describes the core activities workers need to perform to maintain a safe workplace: following standard work procedures, wearing personal protective equipment, completing required safety checks, and engaging with the formal elements of the safety management system as its design requires. Safety participation describes something different: behaviours that do not directly contribute to an individual's own safety but help develop an environment that supports safety. These include taking part in voluntary safety activities, helping colleagues with safety issues, attending safety meetings beyond the required minimum, raising concerns through channels that create work for the system to process, and challenging unsafe conditions that someone else could be expected to deal with.

Compliance and participation line up with rational and emotional engagement, and with transactional and transformational leadership. Compliance is what rational engagement and transactional leadership produce. Participation is what emotional engagement and transformational leadership produce. And the research on their relative contribution to safety outcomes is clear: it is participation, the discretionary, voluntary, beyond-the-minimum contribution, that is the leading indicator of safety culture maturity and the strongest predictor of serious incident prevention.

Research from the offshore oil and gas industry shows how the two separate in practice. One study of 231 workers on six offshore installations measured compliance with safety rules and safety initiative — the voluntary contribution that goes beyond the rules⁷². Transformational leadership, management commitment to safety and management involvement in safety activities were associated particularly with greater safety initiative, while a people-focused supervisory style predicted rule compliance. Site managers exerted greater influence on discretionary safety initiative, while frontline supervisors had greater influence on day-to-day compliance. Transformational leadership, in other words, strengthens the compliance foundation as well as building on it. But the discretionary contribution depended on belonging, on the emotional engagement described earlier in this chapter. People go beyond the rules for an organisation they feel part of.

The reason is not hard to see. Compliance, by definition, produces behaviour within the boundaries of what the system requires, and what the system requires is designed to address the risks the system knows about. The risks it does not know about — emerging hazards, developing conditions, the adaptive practices that have grown up in the space between procedure and practice — are outside the compliance requirement because they are outside the system's view. The worker who merely complies does what the system requires about the risks the system knows about. The worker who participates does what their values, engagement and professional responsibility lead them to do about all the risks they encounter, including the ones the system has not yet identified.

Exceptional safety cultures, those whose safety performance consistently exceeds what their formal management systems would predict, are exceptional because they have created the conditions for widespread participation. Their workers identify hazards the system has not seen. They report near-misses the system would not otherwise capture. They challenge conditions the system has normalised. They coach colleagues in safety practices the system has documented but not embedded. They do all of this not because the system requires it or the incentive structure rewards it, but because their relationship to safety, built through the accumulated effect of transformational safety leadership, has made these contributions a natural expression of who they are.

The question for every safety leader is not whether their workforce is complying. It is whether their workforce is participating. The answer depends not on the quality of the safety management system but on the quality of the

safety leadership that has built, or failed to build, the emotional engagement participation requires.

The Discretionary Safety Leader

The same distinction applies to leaders themselves. The discretionary safety leader is one whose own engagement with safety goes beyond the requirements of the role, contributing to the organisation's safety culture in ways no job description specifies and no performance review measures. They are not simply doing their job well. They are doing it from a commitment that makes the job's boundaries irrelevant. They raise safety concerns that are not formally theirs to raise, because they regard all safety concerns as their responsibility. They engage with workers on safety issues outside their direct span of control, because they see their responsibility to the safety culture as extending as far as their responsibility to the organisation. And through their own discretionary engagement, they model the participation they are trying to build in the people they lead.

The discretionary safety leader is produced by the same conditions that produce discretionary participation in workers: emotional engagement with the organisation's safety purpose, a sense of personal ownership of the safety culture, and the experience of being led by people who show, through their own behaviour, that discretionary engagement with safety is valued. So the quality of safety leadership at every level is both a product and a cause of the engagement culture. Leaders who are themselves engaged with safety create the conditions for engagement in the leaders below them, and the quality of the leadership culture cascades through the organisation in exactly the way the safety culture it is trying to build cascades through the workforce.

Emotional Intelligence as the Foundation

Emotional intelligence has come up in earlier chapters. Here it takes its proper place: as the foundational competency connecting the transformational constructs to the practical capacity to create worker engagement.

In the framework developed by Daniel Goleman⁷³ and refined through extensive empirical research, emotional intelligence has four domains: self-awareness, self-management, social awareness and relationship management.

Each has a direct safety leadership application. Self-awareness, an accurate understanding of one's own safety leadership behaviour and its effect on the culture, is the precondition for developing the transformational constructs at all; a leader who does not know how their behaviour is being received cannot develop what their culture needs from them. Self-management, the ability to regulate one's own emotional responses and stay engaged and clear-headed under the operational pressures and interpersonal tensions safety leadership involves, is what allows the transformational leader to be consistent, especially in the moments of pressure when consistency is most tested and most closely watched.

Social awareness, the capacity to read accurately what workers are experiencing in relation to safety and respond with real understanding, is what makes the Carer's individualised consideration practically possible. A leader who cannot read how their workers are experiencing the safety culture cannot engage with them in the way transformational leadership requires. Relationship management, the ability to build and sustain the trust-based relationships through which the transformational constructs operate, is the practical expression of everything the Knight, the Motivator, the Thinker, the Carer and the Missionary represent.

Research on the business effects of emotional intelligence, working through the climate leaders create, shows that the climate senior leaders create among their direct reports predicts the performance of the whole unit, and that the relationship holds across industries, organisation sizes and cultures⁷⁴. The safety equivalent is that the safety climate leaders create, the emotional and relational environment in which safety decisions are made, predicts the safety participation and performance of the whole organisation. The quality of that climate depends on the emotional intelligence of the leaders who create it.

Emotional intelligence is a developable competency, not a fixed trait. That is what makes the engagement figures at the start of this chapter a leadership problem rather than a fact of working life. Three workers in four are contributing to safety at less than their capacity, and one in four has withdrawn altogether, not because of who they are, but because of how they are led. Leadership can be developed, and what can be developed can be changed.

CHAPTER TWELVE

What Safe Organisations Actually Look Like: High Reliability Organising

The USS Carl Vinson is a Nimitz-class aircraft carrier: a nuclear-powered vessel carrying almost five thousand people and running one of the most dangerous industrial environments on the planet⁷⁵. On its flight deck, aircraft weighing thirty tonnes land at two hundred kilometres per hour on a deck barely three hundred metres long, while fuel lines run live, munitions are moved, engines are running, and dozens of people work within metres of each other amid noise, physical hazard and operational complexity no safety management system has ever fully anticipated. In a single day, the flight deck of a carrier conducts hundreds of operations of a kind that, in a conventional industrial setting, would each need its own risk assessment, its own permit to work and its own pre-start meeting.

The rate of serious incident on a modern United States carrier flight deck is, by any available comparison, extraordinarily low. Not zero: the conditions are dangerous and the margin for error is narrow. But the rate at which operations of this complexity, at this pace, by crews of this size, produce catastrophic events is lower than the rate at which far less complex, less pressured and better-documented operations in other sectors produce them. The flight deck is not safe because someone wrote a better procedure. It is safe because the culture of the people who run it has been built, over decades of deliberate organisational development, into something most organisations have never come close to: a High Reliability Organisation.

The concept is not idealism. High Reliability Organising, the field of research and practice that examines how organisations sustain safety performance in conditions of extreme hazard, is among the most demanding traditions in contemporary safety science. It offers no programme that can be implemented in a quarter. It describes a set of cultural and leadership conditions that, when they are present, produce the kind of organisational awareness that makes serious events less likely than the complexity of the operation would predict. And it begins with a recognition that challenges the premise of most safety management practice: that the path to safety runs not through eliminating error, but through building an organisational capacity to detect, contain and learn from error before it becomes catastrophic.

The Organisations That Fail Least

Karl Weick and Kathleen Sutcliffe⁷⁶, whose work on High Reliability Organising (HRO) remains the most practically useful account of what these organisations do and how they do it, began from a deceptively simple question. What distinguishes organisations that operate for long periods in conditions where errors are highly consequential, without suffering the catastrophic outcomes the probability of error would predict? The question is not why some organisations are safer than others in general. It is more specific and more demanding: why do some organisations, in the most hazardous conditions there are, produce safety outcomes their operational complexity does not seem to justify?

The answer, developed through close study of carrier flight decks, nuclear power plants, air traffic control centres, wildfire crews and other organisations working at the edge of what complexity and hazard permit, is not what most safety management frameworks would predict. It is not better procedures, though their procedures are often excellent. It is not more comprehensive audits, though their audits are thorough. It is not a more sophisticated safety management system, though their systems are well designed. What distinguishes High Reliability Organisations is a particular quality of collective attention: a culture in which awareness of what is actually happening in the operational environment is more complete, more accurate and more continuously maintained than in organisations with less reliable outcomes.

Weick and Sutcliffe describe that quality of attention through five principles. They are not rules or procedures. They are cultural orientations,

ways of looking at the world that have been built into the organisation's basic assumptions through sustained leadership behaviour over time. In the terms of the Transformational Safety® Culture Maturity Model, they describe the Transformational level in concrete operational terms: what safe organisations actually do. Understood properly, as descriptions of what the organisations that fail least look like from the inside, they answer the question that follows naturally from the chapters on cultural failure: if that is what failure looks like, what does success look like, and what does it require?

Mindful Organising

Weick and Sutcliffe gave that quality of attention a name: mindfulness. They did not mean it in the sense the word has since acquired in wellbeing programmes. They meant a rich, continuously refreshed awareness of the detail of one's situation, combined with the capacity to notice when that situation is departing from what was expected. Its opposite is mindlessness: working from expectation, category and routine, seeing what one expects to see, and carrying on with the plan after the conditions have moved away from it.

Mindlessness is not stupidity, and most of it is efficient. Routine lets people get complex work done without re-deciding everything from first principles. The trouble is that routine and expectation are exactly the conditions under which weak signals go unnoticed. A mindless organisation is not a careless one. It is one in which the plan has quietly replaced attention to the work.

Weick and Sutcliffe also divided the five principles into two groups, and the division matters more than it first appears. The first three, preoccupation with failure, reluctance to simplify and sensitivity to operations, are about anticipation: noticing that something unexpected is developing, early enough to act. The last two, commitment to resilience and deference to expertise, are about containment: stopping the unexpected from becoming a catastrophe once it has begun, however it began.

Most organisations invest heavily in the formal equivalents of anticipation: risk assessment, hazard identification, planning. Far fewer invest in containment, because containment means accepting in advance that anticipation will sometimes fail. An organisation strong in anticipation and weak in containment will see most problems coming and have no answer for

the one it does not. An organisation with both can afford to be wrong about what will happen, because it has built the capacity to respond when it is.

Preoccupation With Failure

The first and most counterintuitive principle is preoccupation with failure. HROs treat every near-miss, anomaly, departure from expected conditions and operational surprise as a signal needing serious attention. Not because they expect catastrophe, but because they have learned, from the history of their industries, that catastrophe announces itself through exactly these signals long before it arrives. They are, in the specific sense of the term, worried organisations. Not anxious to the point of paralysis, but continuously alert to the possibility that the next event is developing somewhere the formal indicators have not yet reached.

This is the opposite of the confidence examined in Chapters One and Five. Where most cultures draw reassurance from the absence of recorded harm, from improving incident rates, passed audits and closed corrective actions, HROs draw concern from those same indicators. Not because improvement is unwelcome, but because they understand what most cultures do not: that the absence of recorded harm is not evidence of safe conditions. It is evidence that the conditions that produce harm have not yet lined up with the other elements needed to produce a recorded event. Those conditions may be developing perfectly well.

Preoccupation with failure shows up in behaviours that are easy to recognise once you know what to look for. Near-miss reports in HROs are not filed and closed. They are investigated as seriously as actual incidents, because the HRO's basic assumption is that a near-miss is an incident that happened to miss: an event that reveals a condition in the system and deserves the full attention of one that did not. The investigation asks not which procedure was not followed but what system condition made the decision before the near-miss the rational choice in that moment. The question is always: what does this tell us about the system that we did not already know?

Reporting in HROs is different as a result. Voluntary near-miss reporting runs higher, not because the HRO has a better reporting system, though it may, but because the culture has established, through sustained leadership behaviour, that reporting is welcomed and used. Workers report because they have seen, repeatedly and consistently, that what they report changes

something. The engagement-communication loop examined in Chapter Fourteen, which builds or erodes safety culture depending on whether communication produces a visible response, runs in the positive direction in HROs with a consistency most organisations never approach.

What preoccupation with failure asks of a leader is clearest in how they handle a near-miss. Three questions do most of the work.

1. What did this person see that made their decision reasonable at the time?
2. Where else could the same conditions exist right now?
3. What would we have had to be doing differently for this to have been reported a week earlier?

None of them asks who was at fault. Each treats the near-miss as information about the system, and each tells the person who reported it that their report started something. A leader who asks those questions consistently, and who is visibly more interested in near-misses than in the injury rate, teaches the organisation where to look. A leader who congratulates the site on a clean month teaches it the opposite, however sincerely it is meant.

Reluctance to Simplify

The second principle is reluctance to simplify: resisting the seductive compression of complex conditions into simple explanations, simple categories and simple solutions. HROs are intellectually demanding organisations. They treat the complexity of their operations as a real feature of reality that calls for an equally complex response, and they are suspicious of any explanation of an event that is clean, simple and fully satisfying, because experience has taught them that clean, simple, fully satisfying explanations of complex events are almost always wrong.

Most organisations simplify by design. The information systems through which safety is managed, the incident categories, corrective action types and audit finding classifications, are built to compress operational complexity into manageable categories. This is not irrational. The volume of operational information in a complex organisation is too large to process in full, and categorisation is a necessary feature of any information system. The damaging version is the loss, in the act of categorising, of exactly the information most relevant to understanding what happened and why.

The classic form of this is the root cause analysis that identifies the proximate cause of an event, the broken component, the missed step, the distracted worker, and stops there. It satisfies the investigation's formal requirements without reaching the systemic conditions that made the proximate cause predictable. The simplification is comfortable. It produces a finding, an action and a closure. What it does not produce is an understanding of the conditions that will produce the next event, which the action taken against the finding usually leaves untouched.

HROs counter this tendency through specific cultural and structural mechanisms. They invest in the quality of investigation rather than its speed, resisting pressure to close events quickly in favour of understanding them fully. They use several analytical frameworks rather than one, because different frameworks surface different aspects of complex causation, and the picture from any single framework is reliably incomplete. They actively look for disconfirming evidence, the facts that challenge the emerging explanation, because their basic assumption is that a comfortable explanation of a complex event should be treated with suspicion rather than relief.

For a leader, reluctance to simplify is tested at the point of sign-off. The investigation report arrives, the finding is a missed step, the action is retraining, and the whole thing can be closed today. The leader who sends it back asks what made the missed step likely: what else was happening on that shift, what the crew had been living with, and whether the step had been missed before without consequence. That is slower and less comfortable, and it will sometimes lead back to a decision the leader made. It is also the only way the organisation learns anything the finding did not already assume. The discipline is not to distrust every investigation. It is to distrust the ones that come back neat, and to make it known that “we don't fully understand this yet” is an acceptable thing to bring to a senior leader.

Sensitivity to Operations

The third principle most directly addresses the gap between Work-as-Imagined and Work-as-Done set out in Chapter One. Sensitivity to operations is the capacity to maintain, at every level of the hierarchy, an accurate and current understanding of what is actually happening in the operation: not what the system says is happening, not what the metrics suggest, but what the people doing the work are experiencing right now.

Beyond Compliance

In most organisations, what reaches senior leadership about operational conditions has passed through several layers of summarising, categorising and optimistic presentation. Chapter Five described this as the filtering and reinterpretation of anomalous information. By the time the operational reality of a complex site reaches the executive safety dashboard, it has been compressed, averaged, categorised and compared against targets in ways that strip out the specific, granular, time-sensitive information most relevant to whether conditions are developing toward a serious event.

HROs counter this filtering through what Weick and Sutcliffe describe as a continuously updated picture of the state of the system: an awareness of current conditions maintained less through formal reporting than through the quality of the relationships between leaders and the operational workforce. In practice, sensitivity to operations looks like leaders who are present in the operation, who ask questions that reveal real conditions rather than confirm existing assumptions, and who have built enough trust with their workforce for operational intelligence to travel upward without the defensive filtering most cultures apply.

The flight deck of the Carl Vinson does not run safely because the admiral has a good safety dashboard. It runs safely because every person on that deck, from the most senior officer to the most junior deckhand, shares a continuously updated understanding of what is happening around them, what the risks are in the current moment, and what needs to happen to keep the next operation safe. That shared situational awareness is not produced by a reporting system. It is produced by a culture in which passing on operational intelligence is natural, expected and acted on immediately.

For a leader, sensitivity to operations is decided at the worksite, and mostly in the first ten minutes of a shift walk. A leader who arrives with a clipboard, looks at what they came to check, and asks whether everything is under control will be told that it is. A leader who asks what has changed since yesterday, what the crew is working around, and what they would fix first if it were up to them will be told something they did not know. The difference is the question, and the history behind it: workers answer openly when their previous answers were received well and acted on. This cannot be delegated to a reporting system, because the most useful information is exactly what people will not write down but will say to someone they trust, standing next to the job.

Commitment to Resilience

The fourth principle, commitment to resilience, is the most often misunderstood. In common usage, resilience has come to mean bouncing back from adversity, and HROs are not primarily concerned with bouncing back. In the HRO sense, resilience means something more specific and more demanding: the organisation's developed capacity to detect, contain and recover from errors before they become catastrophic, including in conditions the system's design never anticipated.

The distinction reveals how differently HROs understand the relationship between system design and operational reality. Most safety management systems are designed to prevent errors: to create, through procedure, training and control, conditions in which the errors most likely to cause harm do not occur. That is necessary and valuable. In complex operations it is also insufficient, because the conditions that produce errors in complex systems cannot be fully predicted by the people who design those systems, and, as Chapter One argued, the gap between Work-as-Imagined and Work-as-Done is a permanent feature of the relationship between any designed system and the variable world it operates in.

HROs accept this. They invest heavily in prevention: their procedures are rigorous, their training comprehensive, their risk controls well maintained. But they do not treat prevention as enough. They also invest in the capacity to catch and contain the errors their prevention systems will not stop, because they know those errors will occur. That investment looks less like a formal programme than like a particular kind of operational alertness: the trained capacity of workers at every level to recognise when something is not quite right, and to act on it before the situation crosses a threshold the prevention systems were never designed for.

Building that capacity requires something most safety management systems actively discourage: acknowledging that things will go wrong. In organisations whose safety culture is built around a Zero Harm metric, where any adverse event is treated as failure, the natural response to developing adverse conditions is concealment rather than containment. Chapter Fifteen examines that dynamic in full. Its effect here is simple: the rapid, honest surfacing of developing conditions that resilience depends on is replaced by their suppression.

What commitment to resilience asks of a leader is most visible when a problem is reported while it is still unfolding. A worker calls to say something is going wrong on the plant, and it is partly their doing. The leader's first words decide whether the next person calls early or late. If the first question is how this happened, the organisation has signalled that accountability comes first, and the next person will try to fix it themselves before anyone finds out. If the first question is what the worker needs right now, the organisation has signalled that containment comes first and the review comes later. The review still happens, and it still examines what went wrong. But the order matters, and workers learn it quickly. People surface problems early only to leaders who have shown that early is welcome.

Deference to Expertise

The fifth principle, deference to expertise, is the most challenging to hierarchical cultures, and therefore the most important to understand clearly. It does not mean authority is irrelevant in HROs, or that decisions are made democratically, or that the most senior person present lacks final authority in a real emergency. It means that decision-making authority in a given operational situation moves to the person with the most relevant knowledge of that situation, whatever their position in the formal hierarchy.

On the carrier flight deck, the most junior person has the authority, and the organisational expectation, to stop an operation if they observe a condition they believe makes it unsafe. They do not need to find their supervisor, who does not need to find a senior supervisor, who does not need to consult the officer of the deck. The authority to stop rests with the observation of the condition, not the position of the observer. This is not a safety initiative the Navy implemented. It is the cultural expression of an understanding that the person closest to a condition has information no one further up the hierarchy has, and that the value of that information depends entirely on whether the culture allows it to be acted on immediately.

Deference to expertise addresses the failure behind some of the most consequential industrial accidents of the modern era: the person closest to a developing condition knows something is wrong, and the culture makes it effectively impossible for that knowledge to influence decisions being made further away and higher up. The Challenger engineers and the Texas City

workers both held information the decision-makers never received. The information existed. The culture did not let it travel.

Its foundation is the psychological safety examined in Chapter Six: the belief that raising a concern, challenging an assumption or stopping an operation will not bring personal punishment or humiliation, and that knowledge shared honestly will be received with the seriousness the HRO culture gives every piece of operational intelligence, whatever its source.

The real test for a leader comes the first time a junior worker stops a job and turns out to be wrong. The condition was fine. The stop cost half a shift, a client is unhappy, and everyone is watching to see what the senior person says. If they say, however mildly, that next time the worker should check with someone first, the authority to stop has just been withdrawn, whatever the procedure says. If they thank the worker, publicly, for acting on what they saw, and treat the cost as the price of a system that works, the authority has been confirmed. Deference to expertise is not proven by the stops that turn out to be right; anyone can support those. It is proven by the ones that turn out to be unnecessary, because those tell the workforce whether the authority is real.

The Condition Beneath All Five

Psychological safety was introduced in Chapter Six as the condition on which the rest of safety culture stands or falls. In High Reliability Organising it is more than a foundation. Each of the five principles depends on someone being willing to say something uncomfortable, and each fails in the same way when they are not.

Preoccupation with failure depends on people reporting small failures, including their own. Reluctance to simplify depends on someone in the investigation being willing to say that the neat explanation is wrong, often to a more senior person who has already accepted it. Sensitivity to operations depends on workers telling leaders what is actually happening, including what they are working around and why. Commitment to resilience depends on a problem being raised while it is still small, which usually means someone admitting that it is partly theirs. Deference to expertise depends on the person with the knowledge being willing to use their authority in front of people who outrank them. Take psychological safety away and all five fail together, quietly, while the documentation goes on describing an HRO.

This is why High Reliability cannot be built by adding the five principles to a culture that punishes candour. The principles do not produce psychological safety; they consume it. A leader who announces a preoccupation with failure to a team that has learned to keep bad news quiet will get more reassurance, not more reports.

The relationship runs the other way too, and this is where leaders have their leverage. Every leadership practice described in this chapter is also an act that builds psychological safety. Asking what made a decision reasonable rather than who got it wrong. Sending back an investigation that stopped at the worker. Asking a crew what they are working around, and then fixing something they mentioned. Asking what the worker needs before asking how it happened. Thanking the person whose stop turned out to be unnecessary. Each is an interaction of the kind Chapter Six described, stored by the team as evidence about whether honesty is safe here. Practised consistently, the principles and the climate strengthen each other: people speak up because it has been safe to, and it stays safe because leaders keep responding in ways that make speaking up worthwhile.

Amy Edmondson's research with hospital surgical teams showed this loop at work in one of the steepest hierarchies there is⁷⁷. The teams that learned a demanding new procedure fastest were not the ones with the most senior or most technically gifted surgeons. They were the ones whose leaders made it safe for everyone in the room to speak up, and then used what they heard. That is High Reliability at the scale of a single team, and it is where every organisation that wants to become an HRO has to start.

High Reliability Beyond the Flight Deck

It is easy to read the HRO literature and conclude that it belongs to aircraft carriers and nuclear plants: organisations with military discipline, large budgets and hazards so obvious that no one needs persuading. Most of the organisations I work with are none of those things. They are mines, refineries, hospitals and construction sites, working under commercial pressure, with contractors, turnover and schedules. The principles apply to them, but they look different.

Consider a refinery in the middle of a major turnaround: well over a thousand people on site, most of them contractors who arrived this month, major maintenance taskings packed into a few weeks, and parts of the plant

still live around them. Preoccupation with failure, here, looks like a daily turnaround meeting that opens with the previous day's near-misses and small surprises rather than the schedule, and a leader who asks what each one says about tomorrow's work. Reluctance to simplify looks like refusing to put a cluster of permit errors down to "contractor quality" before anyone has looked at how the permit system is coping with ten times its normal volume. Sensitivity to operations looks like area leaders who spend most of each shift on the work sites rather than in the planning office, and a site where a contractor supervisor can tell the owner's representative that a job is not ready without it being heard as an excuse. Commitment to resilience looks like agreed responses to the situations the plan did not foresee: who decides, with what authority and how quickly, when a live line turns up where the drawings showed none. Deference to expertise looks like the operator who has run that unit for fifteen years being the person whose view on breaking containment outweighs the schedule.

Healthcare offers a different kind of evidence. Hospitals have adopted the language of high reliability more widely than almost any other sector, and their experience shows both its promise and its difficulty⁷⁸. Surgical teams that brief before a procedure and invite anyone in the room to speak up, and wards that treat a medication near-miss as information rather than as an individual failing, are practising HRO principles where hierarchy is steep and time is short. Where adoption has stalled, the reasons are the ones this book has been describing: professional hierarchies that make deference to expertise feel like a threat, and leaders who endorse the principles in policy while responding to errors in the old way.

None of these settings has a carrier's discipline or budget. What they can have is leaders who practise the five principles in the conditions they actually face.

What HRO Requires of Leaders

The five principles are not independent. They form a system of mutually reinforcing orientations. Preoccupation with failure provides the motivation for sensitivity to operations. Reluctance to simplify provides the analytical discipline that makes sensitivity to operations accurate rather than reassuring. Commitment to resilience provides the capacity to respond to what sensitivity to operations surfaces. And deference to expertise provides the distributed

authority that resilience needs in order to act at the speed developing conditions demand.

The system works when all five are present, and it fails, sometimes catastrophically, when any one is absent or compromised. An organisation preoccupied with failure but unwilling to defer to frontline expertise will be alert to developing conditions without the means to surface and act on what those conditions reveal. One that is sensitive to operations but not committed to resilience will have excellent situational awareness and too little capacity to act on it. One preoccupied with failure but lacking psychological safety will be alert to failure and unwilling to report it. This is what makes building HRO culture different in kind from implementing a better safety management system. A management system can be assembled element by element. The five principles cannot, because each depends on the others to work.

The leadership challenge is therefore to develop all five together: not in sequence, but as one coherent cultural orientation, whose coherence is itself the product of leadership behaviour consistent with all five in every interaction. The leader who shows preoccupation with failure in how they respond to near-misses, reluctance to simplify in how they approach investigations, sensitivity to operations in their presence at the worksite, commitment to resilience in how they respond to developing conditions, and deference to expertise in the quality of their listening, is building HRO culture through the accumulated experience of their behaviour. There is no shortcut, and no substitute.

HRO and the Transformational Safety® Leadership Framework

The relationship between High Reliability Organising and the Transformational Safety® Leadership framework is not incidental. The framework was developed in direct engagement with the HRO research tradition, and its nine constructs describe, from the leader's side, the same territory the five principles describe from the organisation's side. They are two accounts of the same thing: one describing what the organisations that fail least look like, the other describing what the leaders who build them actually do.

The passive-avoidant constructs, the Invisible Man and the Fireman, most directly undermine HRO culture. The Invisible Man's absence from the

operation destroys sensitivity to operations. The Fireman's reactive orientation, present after events and absent before them, is the leadership form of a culture that engages with failure only once it has become undeniable. Between them, these constructs produce the confident miscalibration described in Chapter One: the signature of organisations surprised by events their safety systems did not prevent.

The transactional constructs, the Policeman and the Dealer, produce cultures capable of compliance but not of the organisational awareness HRO requires. The Policeman's enforcement orientation is incompatible with deference to expertise: an organisation in which the most important safety authority belongs to the person with the most formal power, rather than the most relevant knowledge, cannot build the distributed decision-making HRO depends on. The Dealer's exchange produces a compliance culture, in which people do what is required because the exchange makes it rational, rather than the safety identity that preoccupation with failure and commitment to resilience both need.

The transformational constructs build each principle directly. Preoccupation with failure grows from the Knight's consistency between stated values and operational decisions, and from the Carer's engagement with the concerns workers are most likely to notice and least likely to report formally. Reluctance to simplify is the Thinker in organisational form: intellectual stimulation, practised by leaders who let workers question the official account of an event without personal risk and who show curiosity about the conditions behind an event rather than the behaviour immediately before it. Sensitivity to operations comes from the opposite of the Invisible Man: a leader who is consistently and purposefully present, backed by the Knight's credibility, which tells workers that what they share will be acted on, and by the Carer's relationships, the informal channels through which the most important operational intelligence travels. Commitment to resilience rests on the Motivator's treatment of safety as a value rather than a metric, and on the Knight's credibility in making it safe to report developing conditions. Deference to expertise depends on the Carer and the Knight together. And the Missionary embeds all of this as cultural norm rather than individual practice, so that deference to expertise stops being an act of deliberate humility by senior leaders and becomes simply the way the organisation makes decisions.

The practical implication is that developing Transformational Safety® Leadership capability at every level of an organisation is the most direct path to High Reliability culture. Not the only path: HRO also needs systemic conditions, structural support and sustained organisational investment that leadership development alone cannot provide. But leadership is the necessary condition. Without the leadership behaviours that build the five principles, no HRO-oriented system or structure will work as designed. With them, system and culture reinforce each other in the way the principles require.

The Case Against, and What Leadership Adds

High Reliability Organising has some serious critics, and any book that holds it up as the standard should take them on.

The oldest challenge comes from Charles Perrow⁷⁹, whose Normal Accident Theory argued that some systems are so interactively complex and so tightly coupled that accidents in them are inevitable, however well they are organised. In such systems, failures combine in ways no one can foresee and spread faster than anyone can intervene. For Perrow, the right answer to some technologies was not better organising but not operating them at all. Scott Sagan's study of United States nuclear weapons safety sharpened the point, documenting a history of near-catastrophes in organisations that looked, on paper, like model HROs.

A second line of criticism is more practical. Andrew Hopkins, whose investigations of Longford and Texas City are required reading for anyone in this field, has pointed out that the HRO literature struggles to say clearly which organisations qualify, and that its principles are easier to describe than to put into practice⁸⁰. Telling a mine manager to be preoccupied with failure does not tell them what to do on Monday.

Both criticisms deserve to be taken seriously. Perrow is right that no organisation can make a sufficiently complex, tightly coupled system perfectly safe, and HRO does not claim to. Its claim is narrower: that organisations can fail far less often than their complexity would predict, and that some demonstrably do. Hopkins is right that the principles are hard to put into practice, and that is the point this book has been making from the other side. The principles do not put themselves into practice. Leaders do: in the

questions they ask about a near-miss, the investigations they send back, the way they walk a shift, how they respond to a problem in progress, and what they say after an unnecessary stop. Without leadership, HRO is a description. With it, it becomes a practice.

Why Many Organisations Struggle to Achieve It

The honest consequence of this chapter is that most organisations, in most industries, working under the conditions that characterise contemporary work, are not High Reliability Organisations and cannot become one quickly. That is not a counsel of despair. It is a realistic account of what HRO requires, and of the distance between most safety cultures today and the Transformational level at which HRO culture becomes possible.

That distance is mainly a leadership distance. HRO is not out of reach for organisations that lack the right procedures or systems. It is out of reach for organisations that lack the sustained, consistent, culturally intelligent leadership that builds the five principles into the workforce's basic assumptions. That leadership takes time to develop, time to apply, and time to accumulate into the assumptions HRO requires.

The implication for safety leaders is uncomfortable but important. HRO is not a target that can be declared and pursued through a programme. It is a direction that can be committed to and pursued through daily, sustained leadership behaviour. The organisation that commits to that direction, building the leadership capability, developing the cultural conditions and sustaining the investment over the time required, will find itself moving toward the High Reliability end of the spectrum. The organisation that declares HRO a goal and implements it as a programme will find itself, over time, with an excellent set of HRO documentation and a culture that has not changed at all.

HRO Theatre

There is a particular failure worth naming, because it is becoming more common as the language of high reliability spreads. Call it HRO theatre: adopting the vocabulary without the practice.

Its signs are easy to recognise. The five principles appear on posters in the crib room and on the back of lanyard cards. Preoccupation with failure becomes a standing agenda item, discussed for five minutes before the production report. Near-misses are celebrated in the monthly newsletter, while the people who reported them notice that nothing on the plant has changed. Every leader has been through a half-day HRO workshop. Deference to expertise is written into the stop-work procedure, and survives until the first stop that costs real money, when it quietly gives way to a conversation about “using judgement”. Investigations use the new terminology and reach the old conclusions.

HRO theatre is no more cynical than the virtual safety systems described in Chapter Two. It is what happens when an organisation adopts a cultural idea through its management system: the idea is turned into artefacts, because artefacts are what management systems produce. The test is the one this book has applied throughout. Look past the poster to what leaders do when the principles are inconvenient. If the behaviour has not changed, nothing else has either.

..... High Reliability Organising is not a destination an organisation declares. It is a direction a leadership culture commits to — and the distance between commitment and arrival is measured in years of consistent, credible, values-aligned behaviour.

The Standard Worth Holding

There is a risk, in presenting High Reliability Organising as the aspiration of optimal safety culture development, of allowing the distance between current reality and HRO-level performance to function as an excuse rather than a challenge. If HRO is several years of sustained leadership development away from where most organisations currently sit, the comfortable inference is that nothing meaningful can be done in the short term, and that the gap between current reality and aspiration is too large to be worth taking seriously.

This inference is just plain wrong, and it is worth being precise about why. The principles of High Reliability Organising are not binary — organisations

do not either have them or not have them. They exist on a spectrum, and every movement along that spectrum in the direction of HRO culture produces genuine safety improvements. An organisation that develops greater preoccupation with failure — that begins investigating near-misses with more seriousness, that creates more conditions for near-miss reporting, that asks harder questions about what anomalous data reveals about system conditions — is producing better safety outcomes in the process, even if it is not quite yet an HRO.

The standard is worth holding not because it is immediately achievable but because it is the most accurate description available of what organisations that keep people safe actually look like. The Transformational Safety® Leadership framework, the maturity model and every argument this book has made point in the same direction: toward organisations in which safety is embedded in identity, operational awareness is continuous and accurate, expertise is respected regardless of its source, and the gap between how work is designed and how it is done is visible, manageable, and continuously narrowing.

That is what the carrier flight deck has built, and what the best nuclear operators have built: organisations that have made the right thing to do the easy thing to do, not by eliminating difficulty but by embedding the values, awareness and decision-making capacity that make safe choices natural rather than effortful.

The TSL framework has now been examined in full: its research foundations, its nine constructs, and the engagement and discretionary effort that transformational leadership produces. The HRO principles in this chapter are what that framework is building toward. The question that remains is organisational: what conditions, structures and strategies allow those leadership capabilities to take root across an entire culture, rather than residing in individual leaders? That is the territory the chapters that follow examine.

CHAPTER THIRTEEN

The Subjective and Objective Worlds of Safety

The programme had taken eighteen months to design and six months to implement. It covered every part of the organisation's safety management system: revised procedures, updated risk assessments, new training modules, a refreshed incident reporting platform, a leadership safety-walk programme, a revised corrective action process, and a comprehensive set of safety metrics the executive team reviewed monthly. The budget had been substantial. The commitment had been real. The programme had been delivered on schedule, within scope, and to a standard of completeness the external reviewers assessed as exemplary.

At the end of its first full year, the incident rate had improved by twelve per cent. Training completion stood at ninety-four per cent. The safety-walk programme was running on schedule. The corrective action register was being kept with a discipline it had never had before. On every objective measure available to the organisation, the programme had worked.

The frontline workers, asked in a structured review what had changed, described something different. The procedures were clearer, they said. The paperwork was better. The training was more consistent. But the way safety actually felt — what it was like to work somewhere safety was valued, where a concern raised was heard and acted on, where the pressure to hold the schedule was balanced by a real commitment to the people doing the work — was essentially unchanged. The programme had improved the objective

infrastructure of safety management. It had not reached the places where safety culture lives.

The organisation had built a better system. It had not yet begun the transformation.

The Two Worlds of Safety

Every safety transformation programme operates in two worlds at once, and the failure to understand and address both is why so many programmes produce system improvements without culture change. The first is the objective world: the visible, measurable, manageable domain of safety systems, structures, processes and behaviours. The second is the subjective world: the invisible, experiential domain of safety beliefs, values, attitudes and intentions, and of the cultural meanings through which every objective element is interpreted and acted on. Both worlds are real. Both are necessary. Each needs a fundamentally different kind of leadership attention to change.

The objective world is where most safety transformation programmes spend most of their resources. It is the world of the safety management system, the training programme, the audit process, the corrective action register and the metrics dashboard: visible to external assessors, measurable with established instruments, and changeable through the authority and budget that formal leadership controls. When an organisation decides to improve its safety performance, its first instinct is almost always to invest here: better systems, procedures, training and data. That investment is necessary, and its results are real. Better systems do produce better outcomes, all else being equal. The trouble is that all else is never equal in an organisation, and the objective world cannot, by itself, decide what the rest will be.

The subjective world is where safety culture lives. It is the world of intentions and values: the internal orientations that determine how workers interpret the safety requirements placed on them, and how they respond in the operational moments the system cannot fully anticipate or control. It is the world of attitudes: the predispositions workers have formed, through accumulated experience of how the organisation actually behaves toward safety, about whether it means what it says, whether its safety system is real or performed, and whether their own safety is truly valued by the people who manage them. It is the world of decisions: the continuous, mostly unconscious choices workers make about how to read ambiguous situations, how to balance

competing demands, and how to act when the system gives no explicit guidance.

The relationship between the two worlds is not the simple one most programmes assume, in which improving the objective world automatically improves the subjective one. Better systems do not automatically produce better values. Better procedures do not automatically produce better attitudes. Better training does not automatically produce better intentions. The objective world is the medium through which leadership speaks to the subjective world, and what that message says depends not on the quality of the infrastructure but on whether what the infrastructure claims matches what the organisation does under pressure.

The Individual Dimension

The subjective world has two dimensions that need separate treatment. The first is individual: the internal world of each person in the workforce. This is the domain of personal intention: how far each worker has taken the organisation's espoused values as their own, developed the attitudes safe behaviour requires, and built a way of deciding that produces safe choices without external instruction or monitoring.

Individual intention is not the same as training completion or knowledge. A worker who has completed every required module, can recite the relevant procedures and behaves compliantly when observed may still, in the unobserved operational moment, make the decision their own accumulated values and attitudes produce, which may or may not be the one the training intended. The gap between what workers know they should do and what they actually do in that moment is a gap in the individual subjective world, not in the training infrastructure. It can only be closed by interventions that reach individual values and attitudes, not by interventions aimed at knowledge and behaviour alone.

This is why the interventions that produce real reductions in incidents, rather than improvements in compliance metrics, are the ones that engage workers at the level of personal meaning: what safety means to them, how it connects to what they value, and how their professional identity relates to their safety behaviour. These are subjective questions. Better procedures and more comprehensive training do not answer them. They are answered through the quality of the leadership relationships in which workers form their

understanding of what safety is for and why it matters, and through the accumulated experience of an organisation that has shown, consistently and visibly, that its commitment is real.

Responsibility is central to the individual dimension. A worker's real acceptance of personal accountability for safe behaviour is not produced by assigning responsibility in the safety management system or writing it into a position description. It is produced by creating conditions in which responsibility feels possible: where workers have the knowledge, resources, authority and psychological safety to make safe decisions and raise concerns, and where the organisation's response to their doing so is consistently positive. Where those conditions exist, responsibility develops. Where they do not, assigning responsibility on paper produces a performance of compliance rather than accountability.

The Person Making the Decision

Safety management tends to treat the worker at the moment of decision as a constant: a trained, competent person applying known procedures to known hazards. The person making the decision is not a constant. They are a human being whose capacity to notice, judge and act varies from day to day and hour to hour, shaped by how much they have slept, what pressure they are under, what their body is doing, and what is happening in the rest of their life. A worker who is exhausted, frightened, in pain or preoccupied is not deciding from the same internal conditions as the same worker at their best. The procedure is the same. The person applying it is not.

This is not an argument for a new category of human error. It is the opposite. The internal state of the person at the sharp end is shaped, far more than most organisations recognise, by conditions the organisation creates: rosters, workload, the pressure of a schedule, the way people are treated when they are struggling. ISO 45003, the international guidance on managing psychological health and safety at work within an ISO 45001 system, treats these as psychosocial hazards to be identified and controlled like any other. That framing matters. It moves the psychological state of the workforce out of the domain of personal resilience and into the domain of risk management, where leaders are accountable for it.

Fatigue is the clearest case. Sleep loss and circadian disruption degrade attention, slow reaction, narrow the range of options a person considers and

make risk harder to judge. After long enough awake, performance falls to a level comparable with a blood alcohol concentration that would be unlawful for driving in many countries⁸¹. The most dangerous feature of fatigue is that people are poor judges of their own impairment: the faculties needed to recognise that one is too tired to work safely are among the first to go. A fatigued worker who says they are fine is usually sincere. This is why fatigue cannot be managed by asking people to monitor themselves. It is managed through rostering, shift length, the sequence of night work, and leaders who notice what the worker cannot. The technician in Chapter One was in the third week of a shutdown. Nobody asked how he had slept.

Some years ago I was asked to review a serious near-miss at a processing plant. An experienced fitter had isolated the wrong pump before starting work on it. It was one of two identical units side by side, and the tag was visible on the other. He caught it himself, at the last moment, when the line he was about to open was warm to the touch. When I sat with him afterwards, he could not explain it. He had done that isolation dozens of times. What emerged, slowly, was the rest of the picture. He was on the ninth day of a twelve-day run of twelve-hour shifts, the last four of them nights. His sleep between shifts was broken by the heat and by a young child at home. On top of all that, nobody had given a moment's thought that he had a half hour journey to/from work. Asked on the morning of the incident whether he was "good to go", he had said yes, and meant it. He did not feel impaired. That was the problem. The investigation had already recorded the cause as "failure to verify isolation". Nobody had looked at the roster.

Stress works on attention in a more specific way. Under pressure, attention narrows onto the immediate task, and the peripheral cues that would signal something is wrong drop out of awareness. Psychologists have understood this narrowing of cue utilisation for more than half a century, and it explains something safety management usually attributes to attitude⁸². Production pressure does not only tempt people to take shortcuts. It physically reduces their capacity to see the weak signals, the unusual sound, the reading that is not quite right, that the preoccupation with failure described in Chapter Twelve depends on³⁵. A leader who increases pressure on a crew is not only raising the incentive to cut corners. They are reducing the crew's ability to notice what is going wrong.

Pain is the state safety management most often misses. Around one adult in five lives with chronic pain, and in physically demanding industries the

figure among the workforce is likely to be higher. Pain competes for the same attentional resources that safe work needs. It disrupts sleep, which compounds fatigue. It changes how people move and what risks they will accept to get a task finished. Also, the medications used to manage it can dull alertness and reaction time. Most workers in pain do not disclose it, because they fear being judged unfit, losing shifts or losing their job, so they keep working and compensate as best they can. Very few leaders have any idea how much pain is present in their workforce on a given day.

Consider this - a man I will call Mark came to see me through his employer's injury management process, which was itself unusual, because he had been at pains to keep his back out of his employer's sight for nearly three years. He was a crane operator, highly regarded, with a reputation for never missing a shift. He had lived with persistent low back pain since a fall outside work, and he managed it the way many workers do: prescription painkillers from his doctor, a heat pack in his work bag, and a set of small adjustments nobody else noticed. He climbed into the cab less often, stayed in it through breaks, and avoided the ground-level checks that meant bending and twisting. On bad days he told me the pain took up "about half my head". It was on one of those days that he misjudged a lift and swung a load closer to a work party than he should have. Nobody was hurt — it scared the life out of him though. If Mark had not shown the initiative and bravery, to come and have a conversation about his pain management, a disaster may have only been a sleepless night away? What struck me was not the event but the three years before it, during which the organisation had one of its most experienced operators working in pain and medicated, and had no idea.

Then there is the rest of life. Relationship breakdown, financial strain, illness in the family, caring responsibilities: none of these stays at the gate. A person carrying them arrives at work distracted, short of sleep and with less capacity to spare for vigilance. The system cannot see any of this. The only part of the organisation that can is a leader who knows their people well enough to notice that someone is not themselves, which is the practical meaning of the individualised consideration the Carer represents.

One of the best supervisors I have worked alongside ran a maintenance crew on a large mine site. She had a habit I came to recognise as the heart of her leadership. At the start of each shift, before the toolbox meeting, she spent a few minutes with each person, not about the work but about them. One morning a normally talkative member of her crew answered in single words

and would not meet her eye. She did not press him in front of the others. She rostered him onto a lower-risk task, found five minutes with him alone, and learned that his marriage had ended the night before and he had not slept. He had come in because he could not afford to lose the shift and did not want to be alone at home. She sorted the shift, found him support, and got him home safely. None of this appeared in any report. When I asked her why she had noticed, she looked at me as though the question made no sense. “Because I know what he’s like when he’s right,” she said.

What this asks of leaders is not clinical skill. It is attention, and the conditions that make honesty possible. The brief check-in at the start of a shift that asks how people are, and means it. Rosters built around what human beings can sustain rather than what the schedule would like. Decisions about pressure made with an understanding of what pressure does to attention. Most importantly, a culture in which a worker can say “I’m not feeling right today” and be met with support rather than suspicion. That sentence is one of the most important pieces of safety information a worker can offer, and in most organisations it is never said, because saying it feels more dangerous than working through it.

When Harm Has Already Happened

Every serious incident leaves people behind: the workers who witnessed it, the colleagues who responded, the crew who worked alongside the person who was hurt or killed, the supervisor who signed the permit, and the worker whose action, or inaction, was part of the chain. Safety management attends to these people as witnesses and sources of evidence. It rarely attends to them as people who have been harmed. Yet the psychological harm an incident leaves behind is real, often lasting, and directly relevant to safety in the months that follow.

The immediate reactions to a traumatic event at work are well understood: shock, disbelief, intrusive images, disturbed sleep, difficulty concentrating, irritability, a sense of unreality. For most people these reactions ease over the following weeks, particularly where they feel supported. For a significant minority they do not, and they develop into persistent conditions: post-traumatic stress, depression, anxiety, or increased reliance on alcohol and other substances to cope. What happens in the first days and weeks, and

especially how the organisation and its leaders respond, has a real bearing on which path people take.

Two groups deserve particular attention. The first is the worker who may have been involved in the event chain - somewhere. Healthcare has named this person the second victim: someone who carries not only the shock of what happened but the belief that they were responsible for it, whatever an investigation later finds. Second victims are at heightened risk of lasting harm, and they are often isolated at exactly the moment they most need support, because the organisation's first instinct is to establish what they did. The second group is those whose exposure is repeated: emergency response teams, first aiders, and the supervisors and managers who deal with the aftermath of one event after another. Cumulative exposure does not announce itself. It builds over years, and it is easily mistaken for burnout, cynicism or a change in character.

There is also a harm that does not fit neatly into clinical categories. Moral injury, first described in military settings, is the damage done when someone witnesses, fails to prevent, or takes part in something that violates their own sense of what is right. A worker who raised a concern that was ignored, and then watched it produce the event they feared, can carry exactly this kind of injury. So can a leader who knows the pressure they applied was part of the story. Moral injury is less about fear than about guilt, shame and a broken trust in the organisation, and it is directly corrosive to safety culture, because it teaches people that speaking up does not matter.

I was asked to see a supervisor some months after a fatality on his site. He had signed the permit for the job on which a young worker was killed. The investigation had found no fault with his decision; the causes lay well upstream, in design and in a maintenance backlog he had repeatedly raised. By every formal measure he had been cleared. He was not sleeping. He had stopped going out with his crew, stopped coaching his son's football team, and was drinking more than he admitted. At work he had become, in his manager's words, "a nightmare" about permits: rechecking everything, refusing jobs over trivial issues, arguing with planners. His manager saw a performance problem. What I saw was a man who had been part of the chain that killed someone. Who believed, whatever the report said, that he should have stopped it, and who had been interviewed three times in the first week but never once asked how he was. He had been treated as a witness. No one had treated him as a person who had been harmed.

Beyond Compliance

Trauma matters to safety in its own right, not only as a welfare issue. Workers struggling after a serious event are often hypervigilant in some situations and avoidant or detached in others. Their concentration and sleep are disrupted. Some withdraw from the team; some take risks they would not previously have taken. The period after a serious incident is, for these reasons, a period of elevated risk, at exactly the time the organisation believes it has just learned its lesson. It is also the period in which the culture learns most sharply what the organisation really values. How leaders treat the people left behind is watched by everyone, and it shapes trust for years.

What the evidence supports in the aftermath is practical and unglamorous. Immediate attention to safety, basic needs and contact with family. Calm, visible leadership that acknowledges what has happened. Practical support in the days that follow: time, flexibility, information, and help with the things people cannot manage while in shock. Psychological first aid, which offers comfort and connection without pressing people to relive the event. Watchful monitoring over the following weeks and months, by leaders and peers who know what to look for, with ready access to professional help for those whose reactions persist.

Two further practices make a difference. The first is the timing and manner of the investigation. Interviewing people while they are in acute shock produces poorer evidence and adds to their distress; a short delay, a supportive setting and clear information about what the investigation is for serve both the person and the inquiry. The second is the treatment of the second victim. The organisation that stands beside the worker involved, while still examining honestly what happened, tells the whole workforce that people will be supported after something goes wrong. That message does more for future reporting than any policy statement, because it answers the question every worker is silently asking: what will happen to me if it is my turn?

Leaders are not immune to any of this. The managers who carry the news, sit with families, face investigations and live with their own part in what happened are exposed too, and in most organisations there is no one whose job it is to notice when they are struggling. A leader who is not coping will find it very difficult to provide the steady, visible leadership their workforce needs. Caring for the people left behind includes caring for those leading them.

The Collective Dimension

The second dimension of the subjective world is collective: the shared domain of values, norms, stories and meanings that make up the organisation's safety culture. Earlier chapters examined this domain as the place where culture forms and fails. In the context of transformation, one part of it needs particular attention: collective identity.

Collective identity, the shared understanding of who we are as an organisation, what we stand for and what kind of people work here, is among the deepest forces in organisational culture. It outweighs any single leader's behaviour, any safety system or programme, and any set of stated values or commitments, because it is the accumulated product of all of them: shaped over the organisation's history by the stories it tells about itself, the experiences its members have shared, and the character it has shown in the moments that most tested what it believes. Chapter Five examined how that identity is defended. The question here is where it leaves a transformation starting from.

For safety transformation, the identity question is fundamental: what does this organisation believe about itself in relation to safety? Not what it says, in the mission statement, the safety policy or the leadership messages, but what it believes, at the level of shared assumption, about whether it is the kind of organisation that keeps people safe. The answer shapes what a transformation can achieve and how long it will take. An organisation that believes it is safe, and is wrong, faces a different challenge from one that knows it is not safe enough and is committed to changing that. Both face a different challenge again from an organisation that has never seriously examined its safety identity at all.

Shared values are where the shift from priority to value, the core argument of Chapter Three, actually happens. Shared values are not what an organisation declares but what it collectively does in the situations that matter: when safety and production compete, when a concern is raised that implicates management, when the schedule is tight and conditions are not quite right. Changing them takes sustained, consistent leadership behaviour that changes the collective experience of what this organisation does when its stated values are tested. That experience accumulates slowly and needs consistency across the whole leadership team, and the change it produces lasts because it was built through experience rather than assertion.

Why Organisations Default to the Objective

The consistent preference of safety programmes for the objective world over the subjective is not accidental. Chapter Two described the incentives that make documentation more attractive than real risk management. Transformation programmes face three further pressures that push investment toward the visible and measurable and away from the experiential.

The first is accountability. Leaders are held accountable for what can be measured, because measurement is how accountability is exercised. Safety systems can be measured: documents counted, completion rates calculated, audit scores produced, corrective actions tracked. Safety culture can be assessed, but it resists the simple numbers accountability systems run on. So organisations invest in what they will be held to account for, the objective infrastructure, while the subjective culture, equally important and harder to measure, receives proportionally less deliberate investment.

The second is tractability. The objective world responds to authority and budget. A leader with both can implement a new safety management system, deliver a new training programme or install a new reporting platform, and each of these can be designed, resourced, implemented and assessed within a set timeframe. The subjective world responds to something else: the accumulated experience of leadership behaviour over time. That cannot be put into a project plan. It needs sustained, consistent behaviour over a period no one can fix in advance, and the change it produces builds gradually and has no completion date. For organisations used to managing improvement as a series of projects, this is deeply uncomfortable.

The third is expertise. The objective world is where safety professionals are trained and where they are most confident. Safety management systems, audit methods, risk assessment frameworks and training design are the technical core of the profession, and safety professionals are generally skilled at using them. The subjective world, the domain of cultural diagnosis, psychological safety, values development and collective identity, needs a different kind of expertise that most safety professionals have not developed and most leadership programmes do not provide. The default to the objective is, in part, a default to what is known.

Together, these pressures produce a consistent pattern: investment biased toward the domains where intervention is measurable, manageable and

familiar, and away from the domains where culture change most needs it. Given the incentives, the pattern is not irrational. It produces exactly the outcome in this chapter's opening: programmes that improve the objective infrastructure without reaching the subjective world, and that build better systems without the transformation those systems need to become real.

Integrating the Four Domains

Taken together, the two worlds and two dimensions give four domains: the individual subjective (intentions and values), the individual objective (behaviour and action), the collective subjective (culture and shared meaning) and the collective objective (systems and structures). Real safety transformation needs simultaneous, deliberate and sustained work across all four. None can be transformed by working only on the others, and none can be neglected without opening a gap that undermines the work done elsewhere.

Integrating the individual subjective and individual objective domains is the most demanding of these tasks, because it means treating the connection between what workers believe and what they do as a leadership responsibility rather than a training or compliance outcome. Individual safety behaviour does not change durably by targeting the behaviour itself. It changes when the values and intentions behind it change: when workers develop a relationship with safety that makes safe decisions feel natural rather than enforced. This is exactly the territory the Transformational Safety® Leadership constructs in Chapter Ten are designed to reach: the Carer, who engages with each worker as someone whose safety beliefs matter; the Knight, whose consistent example gives those beliefs a visible standard; and the Thinker, who creates the conditions for workers to examine their own safety thinking rather than simply follow someone else's. Where leadership works at this level, change in the individual subjective domain sustains change in behaviour long after any formal intervention has ended, because the motivation for safe behaviour has shifted from external pressure to internal conviction.

Integrating the collective subjective and collective objective domains is the challenge most leadership development programmes have failed to address. Improving safety systems without developing culture produces the virtual safety management systems described in Chapter Two: excellent infrastructure, hollow culture. Developing culture without improving systems

produces the opposite: an organisation driven by strong values but without the administrative structure to turn them into consistent practice.

Integrating across the individual and collective dimensions is the deepest challenge of all, because the work on individual values and behaviour has to connect with and reinforce the work on collective culture and systems, and the reverse.

Workers who are developing their safety values need to find a collective culture that confirms and sustains them. A collective culture that is developing a real safety commitment needs individuals who live it consistently

enough to model it for others. The transformation is individual and collective at once, and it goes best when leadership gives both dimensions equal, deliberate attention.



The Transformation Imperative

The choice of the word transformation, rather than change, is deliberate, and it matters more than it might first appear. In organisational life, change carries a particular meaning: disruption, loss and threat to existing arrangements, which is why it meets the resistance every safety improvement programme eventually encounters. Transformation implies something else: not the disruption of what exists but a fundamental reorientation of what the

organisation is. Transformation does not ask an organisation to give something up. It asks it to become something more.

The distinction matters in practice because it shapes the leadership approach. A change programme is managed: it has a design, a timeline, a budget and an end date, and its success is measured by whether it delivers its specified outputs. A transformation is led: it has a direction, a set of values and a sustained leadership commitment, and its success is measured by the depth and durability of the cultural shift it produces. Managed change suits the objective domains, the implementation of new systems and the development of new behaviours. The subjective domains need led transformation, and so does the integration of all four.

For safety leaders, this means the transformation of safety culture is not a project they manage. It is a commitment they live. Every interaction they have with the culture, every question asked or not asked, every decision under pressure, every signal about what the organisation values and what it will trade, either builds the transformation or undermines it. The objective world of systems and the subjective world of culture are both shaped, continuously and unavoidably, by the behaviour of leaders. The difference between transformation and change is the difference between leaders who understand this and act on it, and those who keep managing safety at the surface while the culture they need to build stays, unaddressed, below.

CHAPTER FOURTEEN

The Communication Gap: Why Safety Messages Don't Land

In 2005 the IABC Research Foundation published one of the largest studies of its kind: *Best Practices in Employee Communication*, by Christie Gay, Michelle Mahony and Jamen Graves⁸³. It drew on case studies, interviews with twenty-two leading communicators, and a survey of communicators in 472 organisations across four continents, and its findings about the state of communication inside organisations should have confronted every senior leader who read them.

More than four in five of the communicators surveyed named aligning employees with the organisation's strategy as one of their two biggest challenges. Almost half named the next: getting leaders and frontline managers to take up their role in communicating with the workforce.

These organisations were not indifferent to communication. Their own communicators had identified it as a priority and named its hardest problems, and the two they named are the same two this chapter is about: a workforce that does not share the organisation's direction, and managers, the most powerful communication channel any senior leadership has, who are not carrying it. These are not marginal problems. They are the conditions under which safety culture development either succeeds or collapses.

The Alignment Gap

In organisational communication, the word alignment has a precise meaning that frequent use tends to wear away. Alignment is not agreement; it does not require every worker to endorse every decision the organisation makes. It means that workers understand what the organisation is trying to achieve, how their own work contributes to it, and what values govern the decisions the organisation makes along the way. An aligned workforce can work through the ambiguities and competing pressures of operational life confident that the direction it is moving in is the direction the organisation is asking for.

In safety, alignment means that workers at every level understand the organisation's safety direction, believe it is real, and have taken its values on clearly enough to act consistently with them without explicit instruction. That is a much higher standard than compliance. Compliance requires that workers know the rules and follow them when observed. Alignment requires that they have absorbed the reasoning behind the rules — the values they express, the risks they address, the outcomes they are meant to produce — deeply enough to apply that reasoning in situations the rules never anticipated.

The alignment gap, the distance between leadership's understanding of the organisation's safety direction and the workforce's, is one of the most consequential and most underestimated features of safety culture failure. It is consequential because it is exactly where the decisions most likely to produce serious events are made: by workers trying to do the right thing, who do not understand clearly enough what the right thing looks like in the ambiguous, pressured conditions of their actual work to be confident in their judgement. It is underestimated because the usual tools for assessing alignment, the safety survey, toolbox meeting attendance and training completion rates, measure whether communication happened, not how well it was absorbed.

The same study found that most of the variation in alignment between organisations came down to a small number of factors, chief among them whether employees understood how their own role contributed to the organisation's goals, and whether they trusted their leaders and managers. In safety, two of those factors matter most, and most organisations address them poorly or not at all. The first is the quality of the direct manager's communication: not senior leadership's messages cascaded down through the hierarchy, but the conversation a worker's own manager has with them, one

to one, about safety and what it means. Chapter Eleven showed that the direct manager is among the strongest influences on a worker's engagement. The same is true of alignment, and the quality of that transmission depends less on what senior leaders say than on how far the direct manager understands and believes what they are being asked to pass on⁸⁴. The second factor is consistency between what is communicated and what is shown. Workers align with what they observe as much as with what they are told, and when the two conflict, the observation wins⁸⁵.

Why Communication Fails

Organisations fail at safety communication despite real intent to succeed, and the reasons are neither mysterious nor hard to address once they are clearly understood. The IABC study identified four main barriers to leader and manager communication, and later research has found the same four persisting, with remarkable consistency, across industries, organisation sizes and cultures.

The first is inconsistent messages. When different leaders at different levels speak about safety differently, with different emphasis, framing and apparent conviction, workers receive a fragmented picture of what the organisation actually believes. The senior leader who talks about safety as a value, the middle manager who talks about it as a compliance obligation, and the supervisor who talks about it as paperwork are each transmitting a different version of the organisation's safety culture. Workers in that environment have no coherent picture of what safety means to the organisation. They have several partly inconsistent versions, and they build their own interpretation, usually giving most weight to the version from the person with the most direct influence on their working day.

The second is leaders who disappear when things get difficult. Workers do not judge the organisation's safety commitment by what leaders say during normal operations. They judge it by what leaders do when operations are hard: when the schedule is tight, the budget constrained, the competition fierce, and keeping the stated commitment starts to hurt. Leaders who are visible and engaged on safety in normal conditions but less visible when conditions are difficult are saying something precise about the depth of their commitment. The workforce notices the change, reads it accurately, and adjusts its own engagement to match.

The Communication Gap: Why Safety Messages Don't Land

The third is a lack of trust in leadership. Trust here is not a general quality of the relationship between leaders and workforce. It is a specific judgement, built from experience, about whether the organisation does what it says about safety. Workers who have seen safety commitments made and not honoured — concerns raised and not acted on, corrective actions closed without resolution, leadership messages contradicted by the operational decisions that shape their day — develop a distrust of safety communication that lasts long after the conditions that produced it have changed. That distrust is not irrational. It is the rational product of experience, and better communication alone will not fix it. Only sustained change in behaviour, which gradually rebuilds the experience trust is built from, can do that.

The fourth is managers who do not understand their role in safety communication, or lack the skills and tools to carry it out. The direct manager is where the gap between senior leadership's safety intent and the frontline's safety experience is either closed or widened. Managers who do not realise that their communication behaviour, not only their formal safety messages but their daily, informal engagement with safety, is continuously shaping their team's safety culture are not failing on purpose. They are failing because no one has built in them the capability their role requires.

The Cascade Effect

In most organisations, safety communication is designed as a cascade. Senior leadership develops the safety vision and the values behind it, communicates it to the management team, and expects it to flow down through successive layers until it reaches the frontline workers whose behaviour the culture is meant to shape. The model has an intuitive logic. It mirrors the formal authority structure, it makes efficient use of senior leaders' time, and it allows the message to be tailored at each level to the concerns and context of the audience.

It also has a structural failure mode with serious consequences for safety culture. Every layer of transmission introduces distortion, dilution and reinterpretation, and across several layers the cumulative effect can produce, at the frontline, a message only loosely resembling what senior leadership intended. The distortion is not malicious. It comes from the same dynamic behind every alignment gap: managers who have absorbed the message to different depths, interpreted it through their own values and experience, and

passed it on with whatever conviction, or lack of it, their own relationship to safety produces.

The particular irony of the cascade in safety is that it fails worst exactly where it is most needed: in the relationship between the supervisor and the frontline worker, where safety culture either becomes real in the daily experience of the workforce or stays an aspiration expressed in documents and meetings frontline workers never attend. This is where the engagement gap described in Chapter Eleven is widest, where safety culture failure is most often traced, and where organisations invest least in developing their people as safety communicators, far less than they spend on developing senior leaders. The cascade delivers the message least reliably where the message matters most.

The alternative to the cascade is not a single channel; nothing can substitute for the quality of the relationship between each worker and their direct supervisor. The alternative is a communication culture in which the meaning of safety is developed at every level through dialogue, rather than transmitted from the top through the hierarchy: one in which frontline workers are not the final recipients of a message designed and cascaded by leaders above them, but active participants in building a shared understanding of what safety means in this organisation and why it matters. That model is more demanding to build and sustain. It is also the only one that produces the alignment effective safety culture depends on.

Safety Communication That Works

The research on effective safety communication, together with the experience of organisations that have built strong safety cultures, points to three principles that distinguish communication that builds alignment from communication that produces only its appearance.

The first is that effective safety communication runs in both directions. It does not only transmit information and values from leadership to the workforce. It actively seeks, through both structured and informal channels, what the workforce knows about the safety reality of the work: the gap between how work is designed and how it is done, the conditions developing beneath the formal indicators, and the concerns that have not yet been formally raised. Two-way communication closes the alignment gap in the direction that matters most. It aligns workers with the organisation's safety

direction, and it aligns the organisation's understanding of safety conditions with what workers actually experience.

The second is that effective safety communication is specific rather than generic. Generic communication, the safety moment at the start of every meeting, the quarterly message from the chief executive, the annual safety day, produces a familiarity response. Workers who have heard the same message in the same format many times process it with the efficiency familiarity allows, which is to say with minimal attention. Specific communication, engaging with the actual safety challenges of a particular team in a particular week, a real conversation about a specific near-miss and what it reveals about the system, a leader's direct engagement with a concern a particular worker has raised, gets a different quality of attention precisely because it is specific.

The third is that effective safety communication is credible before it is inspirational. An inspiring safety message from a leader whose behaviour under pressure contradicts it is received as performance. The same message from a leader whose accumulated behaviour has earned credibility is received as conviction. Chapter Fifteen examines how that credibility is built and how quickly it is lost. For communication, the implication is simple: a message that cannot draw on credibility is not communication in the sense culture-building requires. It is noise.

The Engagement-Communication Loop

Engagement and communication are not independent. They form a reinforcing loop that either builds or erodes safety culture, and knowing which direction the loop is currently running is one of the most useful diagnostic questions a safety leader can ask.

The loop works like this. Workers who are engaged with safety talk about it: they raise concerns, share observations, take part in safety discussions, and supply the operational intelligence the formal safety management system cannot capture on its own. That communication gives leaders access to the real safety conditions of the organisation: what is developing in the space between what the system expects and what the work requires, concerns that have not reached the threshold of a formal report, and adaptations experienced workers have developed and newer workers are inheriting. Leaders who receive this communication and respond to it, acting on the concerns raised, using the intelligence to improve the conditions it reveals, and telling the

workers who contributed it what was done, reinforce the engagement that produced it. Workers whose communication is well received become more engaged. Workers who are more engaged communicate more. The loop builds.

It also runs the other way. Workers who raise concerns that are not acted on, provide intelligence that produces no visible response, or see that the leaders they talk to do not seem to use what they share, reduce their investment in safety communication. They learn from experience that it is a compliance ritual rather than a real channel of influence, and they adjust. The formal indicators, reporting rates, meeting participation and survey responses, stay at the compliance level, because compliance is what the organisation requires. The discretionary safety intelligence that real engagement produces dries up, because the loop has run in the negative direction long enough to establish that it is not worth the effort.

The practical implication is that the quality of safety communication does not depend on the quality of the safety communication programme. It depends on the quality of the loop: how far workers' safety communication produces visible, timely responses from leaders that reinforce their investment in speaking up. Building the loop means treating every piece of safety communication from workers, every concern raised, near-miss reported and observation volunteered, as a contribution that deserves a response. The response need not be complicated. It does need to be visible, personal and timely, and it needs to show, through the quality of the engagement, that what the worker shared was heard and valued.

Building the Communication Culture

The communication culture safety alignment requires is not built through a communication programme. It is built through the daily, relational, operationally specific communication of leaders at every level, accumulated over the long period real cultural change takes. The principles are not complicated, but applying them consistently takes a degree of self-discipline and cultural intelligence most leadership development programmes do not build.

Leaders who build effective safety communication cultures share three characteristics that show in their day-to-day engagement with the workforce.

The first is curiosity about the safety reality of the work: not the documented reality the management system describes, but the lived reality

The Communication Gap: Why Safety Messages Don't Land

workers experience in the operational moment. That curiosity shows in the questions leaders ask at the worksite: not questions that confirm what they already assume, but questions that try to understand the gap between how the work is designed and how it is actually done. Workers can tell the difference between a curious leader and one conducting a compliance inspection, and they answer curiosity with a candour inspections rarely get.

The second is visible follow-through. These leaders close the loop consistently, personally, and quickly enough to show that what they were told was taken seriously. Chapter Six described why closing the loop matters even when the answer is not the one the worker wanted. In communication terms, it is what turns a channel from a compliance mechanism into one workers believe in.

The third is the willingness to speak honestly about what is difficult. Safety communication cultures are damaged most not by a shortage of positive safety messages but by dishonest ones: leaders expressing a confidence they do not feel about conditions they know are not fully adequate, and the institutional habit of presenting safety performance in its best light at the expense of the more honest account workers know to be accurate. When workers see leaders talk about safety with a candour that acknowledges difficulty and uncertainty, rather than performing the confidence the role seems to demand, they respond with the kind of trust alignment is built on. The candid leader does not undermine confidence. They build it, because their people know that what they say can be believed.

CHAPTER FIFTEEN

Belief, Zero Harm, and the Psychology of Credibility

Every year, in organisations across every industry, large numbers of workers attend safety meetings at which the organisation's commitment to Zero Harm is reaffirmed. They hear the statistics: the improvement in the injury rate, the days since the last recordable incident, the progress toward the annual target. They are reminded of the organisation's commitment and asked to commit, or recommit, to their own safe behaviour. Some of them nod. Some of them write in the spaces provided on the form. A substantial proportion of them do not believe what they are being asked to believe.

This is not cynicism, and it is not disengagement. It is the rational product of experience: experience of a complexity that the Zero Harm framework, as most organisations use it, is not designed to accommodate. The workers who do not believe the Zero Harm aspiration are not saying they do not want to be safe. They are saying something more specific and more important: that the version of Zero Harm being presented to them does not match the operational reality they work in every day, and that the gap between the claim and the reality has grown wide enough to damage the credibility of the people making it.

When large numbers of workers report in confidential surveys that they do not believe zero harm is achievable, they are not providing evidence of a workforce attitude problem. They are providing evidence of a leadership credibility problem, and credibility is the foundational resource of safety culture. Without it, no amount of communication, no quality of systems and

no depth of leadership commitment can produce the engagement safety culture requires.

The Architecture of Credibility

Credibility is not simply the absence of dishonesty. It is a relational quality, built from experience, that determines how communication is received and whether it changes the beliefs and behaviour of the people who receive it. A leader with high credibility can communicate difficult truths and be believed. A leader with low credibility can communicate sincere commitments and be disbelieved. The difference lies not in what is said but in the history of experience through which the audience hears it.

In safety culture, credibility operates at two levels. The first is individual leadership credibility: how far each worker believes that their own leader means what they say about safety. This is the credibility of the Knight described in Chapter Ten, built through consistency between stated values and actual decisions and accumulated over repeated observation under real pressure. Individual leadership credibility is the main determinant of how safety communication is received at the worksite. Every other channel, the senior leadership message, the safety management system, the formal training programme, is filtered through each worker's assessment of their direct leader before it changes any belief or behaviour.

The second level is organisational credibility: how far workers believe that the organisation as a whole means what it says about safety. It is built from the same material, the accumulated experience of whether the organisation does what it says, but it is more complex, because it is shaped by the behaviour of the whole leadership system rather than one leader. An organisation can have individual leaders with high credibility whose impact is undermined by low organisational credibility: workers who trust their supervisor but not senior leadership, or who believe the commitment is real at the operational level but compromised by strategic decisions that trade safety for other priorities.

That relationship creates a particular challenge for safety culture transformation. The effort invested in developing individual safety leaders can be undone by organisational decisions that show a different set of values at the strategic level. A supervisor who has built credibility with their team through years of consistent, values-driven safety behaviour puts that credibility at risk every time a strategic decision, about budget, staffing or production targets,

shows that the organisation's stated commitment is conditional. Workers watch both levels, and their judgement of the organisation's credibility is a composite that neither level can determine alone.

Credibility is also asymmetric: it is built slowly and lost quickly. I am certain you may have heard someone say something like "it takes a lifetime to create trust and a moment to lose it". Paul Slovic's research⁸⁶ on trust found that trust-destroying events are more visible, carry more weight and are remembered longer than trust-building ones. Safety is no exception. A leader can spend years building credibility through consistent decisions, and lose much of it in a single afternoon: the one time a concern was brushed aside when the schedule was tight, or a stop was overruled because a client was waiting. Workers do not average a leader's behaviour. They remember the exceptions, because the exceptions tell them what the commitment is worth when it is tested. That is why credibility cannot be banked. It has to be earned again in every decision that matters.

The Trust Equation

David Maister, Charles Green and Robert Galford⁸⁷, writing about the trust clients place in their professional advisers, set out a model of trust that has since been applied far beyond the consulting room. They called it the Trust Equation. Trust, in their account, is the sum of three things, credibility, reliability and intimacy, divided by a fourth: self-orientation. It translates to safety leadership with very little adjustment, and it gives structure to much of what this book has been describing.

Credibility, in the equation, concerns only the words: whether what a leader says about safety is accurate, informed and honest. A



The diagram illustrates the Trust Equation. On the left, a blue circle contains a white 'T' with the word 'Trustworthiness' written below it. This is followed by an equals sign. To the right of the equals sign is a fraction. The numerator consists of three blue letters: 'C', 'R', and 'I', each with a label below it: 'Credibility', 'Reliability', and 'Intimacy' respectively. These three letters are separated by plus signs. The denominator is a single red letter 'S' with the label 'Self-Orientation' written below it. A horizontal line separates the numerator from the denominator.

leader who understands the hazards of the work, speaks about them truthfully and does not overstate what the organisation has achieved scores well here. A leader who asserts that Zero Harm is an achievable target, to a workforce that knows otherwise, scores badly, however sincere they are.

Reliability concerns the actions: whether a leader's behaviour matches their words, consistently and over time. It is what people usually mean when they call a leader credible, and it is built slowly and lost quickly. This is the Knight's territory, and it is the element this book has returned to most often: the leader who stops the job when stopping costs something, who follows up on the concern, and who behaves the same way under pressure as they do when nothing is at stake.

Intimacy concerns safety in the other sense: whether it feels safe to tell this person the truth. It is psychological safety at the level of a single relationship, and it is what the Carer builds. It is earned in small moments, such as how a leader reacts the first time someone admits a mistake. A leader can be accurate and reliable and still hear nothing important, because no one feels able to bring them bad news.

Self-orientation is the denominator, and it is the element safety leadership most often overlooks. It is the degree to which a leader appears to be acting in their own interest rather than the interest of the people they lead. Because it divides everything else, it can undo the rest. A leader who is knowledgeable, consistent and approachable will still not be trusted if workers conclude that what really matters to them is their own number, their own bonus or their own reputation. Workers rarely say so. They simply adjust what they tell that leader.

This is where the equation sharpens the argument about Zero Harm. When Zero Harm is a target on which leaders are assessed and rewarded, every leader's personal interest becomes bound up with the figure. Workers can see that. The leader who urges them to report every incident, while visibly anxious about what those reports will do to the monthly result, is sending two messages at once, and the one workers believe is the one about self-interest. The supervisor who handles an injury informally to protect the team's record may believe they are protecting the team. Everyone watching reads it as protecting the record.

The practical value of the equation is diagnostic. When trust is low, it asks leaders which element is failing. Is the problem what they say, what they do, whether people feel safe telling them the truth, or whose interests they appear to be serving? Each has a different remedy, and the last is the one leaders find hardest to see in themselves, because self-orientation is almost always visible to others before it is visible to the person displaying it.

The Zero Harm Credibility Problem

Zero Harm as a safety aspiration is morally and directionally right. The aspiration to eliminate harm at work entirely, to build organisations in which no one is injured or made ill by their work, is the only aspiration consistent with a real commitment to human life and dignity. This book does not argue against it. It argues against a specific and consequential error in how most organisations hold and communicate it: an error that turns a real moral commitment into a performance that destroys credibility.

The error is using Zero Harm as a metric rather than a direction. When Zero Harm is declared as the target, the measurable outcome against which safety performance is judged and safety leaders are held accountable, it produces a predictable set of consequences that undermine the safety culture it was meant to build.

The first consequence is what happens to reporting. A target is something people are rewarded for meeting and penalised for missing, and targets are managed. Under a Zero Harm target, every recorded injury is a failure against the number, and the organisation begins, without anyone deciding it, to manage the number. Injuries are reclassified to fall below the recordable threshold. Injured workers are kept on site on light duties so that the day does not count as lost time. Supervisors whose bonuses or reputations depend on the figure find ways to have incidents handled informally. Workers learn that reporting a minor injury or near-miss brings attention, paperwork and a sense of having let the team down, and they stop reporting the small events that are the organisation's best warning of the large ones⁸⁸. The dashboard improves. The information the organisation needs to prevent serious harm quietly disappears from it. When harm is concealed rather than surfaced, it cannot be contained or learned from; the gap between recorded performance and real conditions widens exactly as the organisation believes it is closing.

The second consequence is less visible and runs deeper: the damage Zero Harm as a metric does to credibility in the workforce expected to believe in it.

Workers have a direct and detailed understanding of the conditions under which harm occurs. They know the hazards in their work environment, the conditions under which those hazards are likely to cause harm, and how adequate the organisation's current risk controls are to those conditions. Many know from direct experience that the gap between the Zero Harm aspiration and the reality they work in is larger than the organisation's public

communication admits. When leaders stand in front of them and assert that Zero Harm is not only the aspiration but the achievable target of the safety programme, they are making a claim that a significant part of their audience knows, from experience, to be false⁸⁹.

The consequence of making a claim workers know to be false is not simply that the claim is disbelieved. The credibility of the person making it, and of the organisation's safety communication more broadly, is damaged. Workers who hear Zero Harm asserted as an achievable performance target, rather than a permanent direction, and who know from experience that it does not describe their reality, revise their assessment of the leader's credibility, and not upward. They carry that revised assessment into every later piece of safety communication, filtering it through an established pattern: this organisation's safety communication contains assertions that do not reflect the reality of the work.

The Psychology of Belief in Safety Culture

To understand why the credibility problem matters so much, it helps to understand the psychology of belief in organisations: how workers form their beliefs about safety, and what allows those beliefs to be held rather than performed.

Belief is not knowledge. A worker who knows the content of the safety policy, can recite the procedures, and can describe the organisation's safety values in the policy's own language may not believe any of it in the sense that matters for safety culture: the personal conviction that makes values actionable in the unobserved operational moment. Belief in a safety culture requires that the values the organisation espouses are experienced as consistent with what it is actually like to work there: that the commitment to safety is not simply asserted but shown, not only in formal communication but in the daily decisions workers observe and from which they draw their most reliable understanding of what the organisation values.

Safety belief forms along a consistent path. It begins with the experience of leadership behaviour, which provides evidence, for or against, that the organisation's stated safety values are real. It is reinforced or undermined by the experience of colleagues, whose beliefs provide the social setting in which

each individual's beliefs develop. And it settles into a stable state, held conviction or deeply sceptical performance, through the accumulated weight of that experience over time. Once reached, that state is not easily changed by communication. It changes only with new experience: specifically, sustained experience of safety leadership that is credibly different from what produced the existing belief.

This has a direct implication for Zero Harm. Workers who do not believe the organisation's Zero Harm aspiration are not expressing an attitude that better communication about Zero Harm can change. They are expressing the conclusion of an evidence-based assessment of whether the organisation's commitment to safety is real and adequate to the outcome it claims to pursue. Improving the communication of the aspiration without addressing the evidence behind the conclusion will not change the conclusion. It will add one more piece of aspirational communication to a body of evidence already producing scepticism.

The credibility needed to make Zero Harm believable cannot be built by declaring it more emphatically. It can only be built by showing, through the accumulated evidence workers observe, that the organisation's commitment to eliminating harm is real: that it shapes the operational decisions that most directly affect safety, governs the resource choices that decide whether risk controls are adequate, and holds in the moments of production pressure where most organisations have shown that other values compete more powerfully.

When the Culture Stops Believing

A safety culture that has lost belief in the organisation's safety commitment has not simply become sceptical. It has developed a specific adaptation, a dual culture of performed safety and real safety, that is both stable and dangerous. It is the gap between espoused and enacted values described in Chapter Four, hardened into a way of working. In the performed culture, workers and leaders treat safety management as a compliance exercise: they do what the system requires, present what audits assess, and say what the organisation's safety messaging expects them to say. In the real culture beneath it, decisions are made on the values and assumptions experience has established, including a far more complex and ambivalent view of the organisation's commitment to safety than the performed culture ever expresses.

The dual culture is stable because it lets the organisation function and because it is rational for the people within it. The performed culture satisfies the formal requirements of the safety management system and avoids the consequences of visible non-compliance. The real culture lets workers deal with the operational reality of their work using the values and judgement experience has given them. The two coexist without obvious friction because they operate in different registers: the performed culture is the public register of safety behaviour, and the real culture is the private register of safety decision-making.

It is dangerous for a reason familiar from earlier chapters: the conditions that actually decide whether harm occurs develop in the real register, not the performed one. A workforce whose safety decisions are governed by a belief in the organisation's commitment produces a different safety outcome from one that performs that commitment in the formal channels while deciding from a different set of values. The gap between the two registers is where the latent conditions for serious events accumulate, invisible to the formal indicators the organisation uses to judge its performance.

Rebuilding belief in a culture that has developed this dual response takes a different approach from building belief in the first place. It requires an explicit acknowledgement that the existing evidence has produced reasonable scepticism, and that the path back to belief runs through changed experience rather than better communication. That is an uncomfortable admission for organisations that have invested heavily in safety communication and expected cultural change in return. It is also the honest account of how belief works, and the account a leader with real credibility would be willing to give.

Reframing Zero Harm

The way beyond the Zero Harm credibility problem is not to abandon the aspiration. It is to reframe how the aspiration is held and communicated, aligning what the organisation says publicly about safety with an honest account of what safety management can and cannot achieve.

Zero Harm held as a direction communicates something different, and more credible, than Zero Harm held as a performance metric. It says that the organisation will never accept any level of harm as inevitable or acceptable; that every event in which someone is injured or made ill is a failure of the system that the organisation is committed to learning from; and that the

direction of its safety effort is permanently toward the elimination of harm, even if reaching it is a long-horizon aspiration rather than a near-term target. This framing is honest. It is aspirational in the true sense. It is credible because it does not ask workers to believe something their experience tells them is false.

The shift from metric to direction needs a matching shift in how safety performance is measured and safety leaders are assessed. An organisation that holds Zero Harm as a direction measures the state of the system, not only its outcomes: the quality of hazard identification, the completeness of near-miss reporting, the responsiveness of the corrective action process, the alignment between work as designed and work as done, and, above all, the quality of leadership engagement with operational reality. As Chapter One argued, these measures resist the dashboard compression lag indicators allow. That is their value. They describe the actual condition of the safety culture rather than the outcomes it happened to produce over a period that may or may not have been representative. They reward reporting rather than penalising it, they remove the pressure that makes a Zero Harm target conceal the very information it depends on.

Belief as a Leadership Achievement

Workers' belief in the organisation's safety commitment is not a communication outcome. It is a leadership achievement: the product of leaders who have built, through accumulated experience of consistent, values-driven behaviour, the credibility that lets their communication be received as sincere rather than as performance.

The leaders who achieve it share one characteristic that separates them from those who do not: they are honest about the gap between the aspiration and the current reality. They do not claim the organisation has already achieved what it is trying to build. They describe, with specificity, where the organisation is now, where it is trying to go, and what the journey between the two asks of everyone in it. That kind of communication is more demanding than the generic commitment message. It needs leaders to have an accurate understanding of their organisation's current safety culture: not the culture they would like it to be, or the one the formal indicators suggest, but the one the workforce actually lives in.

It also takes personal courage: to speak about safety in terms that acknowledge difficulty, uncertainty and real complexity, terms the performed

culture has learned to avoid because they seem at odds with the confidence leaders are supposed to project. Chapter Fourteen described what candour does for communication. For credibility, its effect is sharper still. Being honest about the distance between aspiration and reality is how a leader shows that the aspiration is real rather than cosmetic, and it is the behaviour that most clearly distinguishes the leader who believes in safety from the one who performs that belief.

Go back to the safety meeting at the start of this chapter: the statistics, the reaffirmed commitment, the form, the nods from people who do not believe what they are being asked to believe. Nothing said in that room will change what they believe, because what they believe was never formed in rooms like that. It was formed on the job, by watching what the organisation did when safety cost something. It will change the same way, or not at all.

CHAPTER SIXTEEN

The Transformational Safety® Integrated Model

In 1931, Herbert William Heinrich⁹⁰ published *Industrial Accident Prevention*, and with it a model of accident causation whose influence on the following century of safety management has been, on balance, deeply harmful^{91,92}. The model was simple: a sequence of five dominoes, each labelled with a causal factor. Ancestry and social environment produced a fault of person, which produced an unsafe act or physical condition, which produced an accident, which produced an injury. Its simplicity made it immediately useful to organisations not yet equipped for systemic thinking, and it spread accordingly, embedded in training programmes, investigation frameworks and the institutional assumptions of safety management in every industry.

The claim that did the most damage was never prominently labelled as such. Heinrich asserted, without credible empirical foundation, that around eighty-eight per cent of workplace accidents are caused by the unsafe acts of persons. That figure has been repeated, cited and built on for nearly a century. It has no reliable basis in data. Heinrich said his conclusions came from extensive accident records and interviews, but the records had been completed by supervisors, and the interviews were with supervisors, not with the injured workers themselves. He changed the ratios he proposed across the four editions of his book without explaining why. And when Dan Petersen was asked to edit a fifth edition and went looking for the source data, none could be found⁹³. No one has found it since.

What Heinrich's model has is an almost irresistible organisational usefulness — a seduction if you like: it locates the cause of accidents in the behaviour of workers, which is visible and manageable, and conveniently turns attention away from the systems, cultures and leadership decisions that are far harder to examine.

The figure was not science. It was a permission structure, and organisations accepted it gratefully.

Heinrich's dominoes have been revised, extended and critiqued many times since. The Swiss Cheese Model, bow-tie analysis, AcciMap and a range of other causal frameworks have added layers of complexity and analytical power that Heinrich's linear chain could not hold. What has proved far more resistant to revision than the models is the assumption the domino model was built to support: that the most important cause of workplace accidents is the behaviour of the person at the sharp end. That assumption has survived the critique of the linear model, and it continues to shape incident investigation, safety investment and safety culture development in ways this chapter will argue are inadequate to the accidents organisations are actually trying to prevent.

The Transformational Safety® Integrated Model does not revise Heinrich's framework. It replaces it, with a causal architecture grounded in systems theory and in James Reason's model of organisational accidents, which locates the determinants of safety outcomes not in the behaviour of individual workers but in the organisational, cultural and leadership conditions that produce that behaviour. It takes seriously the causal complexity the preceding chapters have described, and it gives leaders a practical way of seeing where, in the system that produces safety outcomes, the greatest leverage for improvement lies.

The Person-Centred Failure

The failure of the domino model, and of the person-centred thinking it institutionalised, is not that individual behaviour is irrelevant to safety outcomes. It is that individual behaviour is treated as the origin of the causal chain rather than its most visible expression. Placing personal factors at the start of the sequence implies that the characteristics that make individuals prone to accidents, their cognitive styles, risk tolerance, physical characteristics and social background, are the main determinants of whether accidents

happen. The model was simple, easy to draw on a whiteboard and easy to teach, and those qualities helped it spread well beyond the evidence that supported it. The practical consequence, built into decades of investigation practice, is a habit of naming the person whose behaviour most directly preceded an event as its main cause, attributing that behaviour to the individual rather than to the system, and responding with interventions aimed at the individual rather than the conditions. The toolbox talk on situational awareness, the formal warning, the refresher on a procedure the worker already knew: each is a familiar product of that habit, and each leaves the work itself exactly as it was.

That habit has proved so durable because it fits the way people naturally explain the behaviour of others. Psychologists have long described the tendency to explain someone else's actions by their character while overlooking the situation they were in, and to explain our own actions the other way round⁹⁴. Investigators are not immune to it. Hindsight makes it worse: once the outcome is known, the warning signs look obvious and the decision that preceded the event looks careless, when at the time it looked like any other decision on any other shift. The investigator knows how the story ends. The worker did not. The language of investigation reports reflects this, with its steady supply of "should have", "failed to" and "did not identify", each phrase measuring a person against information that only became available after the event. There is also a quieter, organisational reason. A person can be retrained, disciplined or moved on, and the matter closed. Conditions are harder to change, slower to fix and more expensive, and they often lead back to decisions made by the people commissioning the investigation. A finding that points at a worker points down the organisation. A finding that points at conditions tends to point up it, and investigations are rarely commissioned by the people at the bottom.

The Longford gas plant explosion in Victoria, Australia, in 1998 showed how far this can go. Two workers were killed, and the state was left without gas for around two weeks. The company's defence placed the blame on its operators. The Royal Commission that investigated rejected that account⁹⁵. It cleared the workers and found that the causes lay in the company's failure to train its people adequately and to give them up-to-date operating procedures: in the conditions the organisation had created, not in the decisions of the people working within them. The operators on shift that day were dealing with a situation they had not been trained to recognise, in equipment that had

become dangerously cold, and they had not been told what would happen if warm oil was reintroduced. Their actions made sense to them because nothing in their training or their procedures said otherwise. Andrew Hopkins' account of the disaster remains one of the clearest studies of what person-centred explanations hide⁹⁶.

The research does not support person-centred explanations. Individual characteristics do influence safety behaviour; the relationship between personality and accident involvement is real, if modest. Those individual factors are consistently and substantially shaped by the conditions of the work environment. The same person, working in a well-designed system with supportive leadership and a strong safety culture, produces safer outcomes than when working in a poorly designed system with absent leadership and a compliance-only culture. Fatigue, time pressure, unclear procedures, poorly designed equipment and the unspoken expectations of supervisors all act on the person long before the person acts on the hazard. An operator on the eleventh night of a twelve-night roster, working a procedure that does not match the plant, under a supervisor who has made it plain the job needs to be finished, is not freely choosing a risk. The individual is not the primary cause. The system is. An investigation framework that directs attention to the worker rather than the system produces findings that support individual interventions, retraining, discipline, procedural changes, without touching the conditions that made the worker's behaviour a predictable response to the pressures they faced. Replace the worker and the conditions remain, waiting for the next person to step into the same position. The event is investigated, the lessons are documented, and the conditions that produced it go on developing until the next event makes them visible again.

The cost does not end with the conditions left in place. Every worker watches what happens to a colleague after an event, and draws conclusions. When the finding is that one person got it wrong, the lesson the workforce learns is not about the hazard. It is about the risk of being the person standing nearest when something goes wrong. People become more careful about what they report, less willing to admit an error or describe a workaround honestly, and quicker to protect themselves. Near-miss reports thin out, not because there are fewer near misses but because each one now carries the question of who will be blamed for it. Workarounds do not disappear; they simply stop being discussed. Interviews after the next event become guarded, and the account investigators receive is the one that is safest to give rather than the

one that is most accurate. Person-centred investigation does not only miss the cause of the last event. It shuts down the information the organisation needs to prevent the next one.

A Different Starting Point

The Transformational Safety® Integrated Model begins where the domino model's analysis ends, with the organisational and cultural conditions, and traces how those conditions reach the event, rather than starting with the individual and working toward the consequence. The reversal is not cosmetic. It reflects a fundamentally different understanding of where causation lies, and therefore of where prevention has leverage. The domino model treats the organisation as the backdrop against which individuals make mistakes. The Integrated Model treats it as the source of the conditions those mistakes grow from. If the cause sits with the person, prevention means changing the person. If it sits in the conditions, prevention means changing the conditions, and that work belongs to the people who set them.

That reversal changes the question an investigation asks. Instead of asking what is wrong with this person, it asks what conditions produced this behaviour. An investigation that asks the first question will always find an answer, because any person, examined closely enough after an event, can be shown to have done something differently from how it was written down. The individual does not disappear from the analysis. The person at the sharp end still makes the decision, with their own experience, judgement and way of thinking. The key difference is that they are understood as someone responding to conditions the organisation created, not as the origin of the event. That shift changes everything that follows: what evidence is gathered, who is interviewed, what the findings say and what the organisation does next.

It starts with a simple principle: people do what makes sense to them at the time, given what they know, what they are trying to achieve and what they are paying attention to. Almost no one at work sets out to be harmed or to harm others. Seen from inside the moment, most unsafe decisions look reasonable. Seen from outside, after the outcome is known, they look reckless. The investigator's job is to recover the inside view. So the useful question about any unsafe decision is not why the person got it wrong, but why it made sense to them in that moment. What were they trying to get done? What did they know, and what could they not see? What pressure were they under, and

from whom? Was the procedure workable as written? Had the shortcut been used many times before without consequence? What had happened the last time someone in their position made the other choice? Answered honestly, these questions lead away from the individual and into the conditions around them.

A practical test helps keep an investigation honest. It is sometimes called the substitution test⁹⁷. Take the person out of the event and put in someone else with similar training and experience, facing the same conditions, the same pressures and the same information. Would they probably have done the same thing? If the answer is yes, the cause is not in the person. It is in the conditions, and replacing or retraining the person will change nothing except who is standing there next time. The test works because it takes hindsight and character out of the picture and puts the work back in. If the answer is no, the investigation has learned something about the individual, but it is not finished. It still needs to ask how that person came to be in that role, with that training, doing that task, without anyone noticing.


None of this removes accountability. There are still choices no reasonable person would make in any conditions, and they should be dealt with as such. But they are far rarer than investigation findings suggest, and treating them as the norm hides everything else. Holding people accountable for conditions they had no power to change is not accountability. It is blame, and the workforce knows the difference. Accountability, properly understood, applies at every level of the organisation, in proportion to the control each level has over the conditions of the work. The reversal also changes whose decisions come under examination. When an investigation starts with conditions, it follows them to where they were set, and that is usually well above the sharp end. The staffing level, the schedule, the deferred repair and the tolerated workaround were all somebody's decision. The model does not let those decisions stay out of view.

Those conditions are the organisational factors, local workplace factors and latent condition pathways that sit beneath every unsafe decision in the model. They are rarely visible on the day of the event. They build up over months and years, through budget rounds, restructures, schedule changes and small decisions that each made sense at the time they were made. Reason's organisational accident model began to map them. The Integrated Model extends it into the specific causal architecture of safety failure, making visible what person-centred investigation has spent a century ignoring.

The Architecture of the Integrated Model

The Integrated Model builds on the organisational accident model developed by James Reason and keeps its basic structure, shown in the figure below.

At the top of the model is the event itself: hazards, the defences the organisation has placed between those hazards and the people and assets they threaten, and the losses that follow when the defences fail. Every serious accident is, in the end, a trajectory from a hazard through the defences to a loss. The rest of the model explains how the defences come to fail. Some defences are physical: guards, interlocks, relief valves. Others are people: the operator who notices, the worker who stops the job. Whether the physical defences hold depends on leadership decisions about maintenance, resourcing and schedule. Whether the human defences hold depends on whether leaders stand behind them. The figure also shows the defences with arrows running through them, because they are not fixed. They move, and how much they move is one of the most important things the model has to say. The section on still defences below explains why.



At the foundation are organisational factors: the strategic decisions, cultural assumptions, resource choices and leadership behaviours that set the conditions for everything above them. They include the safety management system as designed (not as implemented), the production pressures the organisation's commercial position generates, the quality and consistency of senior safety leadership, the assumptions about production and safety the organisation has developed over time, and the formal and informal rewards that shape what workers and leaders are motivated to do. These are the conditions examined in Chapters Four to Thirteen. They are also the furthest from the event, the hardest to connect to it, and the most resistant to the event-driven investigation most safety management systems are designed to carry out.

Organisational factors shape local workplace factors: the immediate operational conditions in which work is done. These include the condition of equipment, the design of the work, the time and resources available for the task, the social dynamics of the team, how present and engaged supervision is, and the ambient production pressure of the shift. Workers cannot change these conditions. They inherit them when they arrive at work, and must navigate them with the values, habits of thought and decision frameworks the organisation has shaped. David Hofmann⁹⁸ and his colleagues described the same relationship in their 1995 study of high-reliability process industries: the attitudes and behaviours visible in a work group are the micro-elements of an organisation, shaped by the macro-elements of its management systems and practices. What happens at the sharp end is, to a large degree, what the organisation above it has made likely.

Local workplace factors, in turn, produce unsafe decisions: the choices made at the sharp end, and the actions that follow from them, that breach or bypass the defences. This is where most investigations stop, and where the domino model begins. In the Integrated Model, unsafe decisions sit near the top of the causal chain, not at its origin. They are the last link before the defences, not the first, and they are best understood as the predictable product of the conditions beneath them.

The model's most important feature is the one person-centred thinking cannot see: the latent condition pathways. These are the routes by which organisational factors and local workplace factors reach the defences directly, without passing through anyone's unsafe decision. Degraded controls, maintenance backlogs, design flaws, normalised deviations and silenced

concerns weaken the defences whether or not anyone at the sharp end does anything wrong. This is why an organisation can suffer a serious event in which no one acted unsafely in any ordinary sense. Chapter One introduced Reason's latent conditions. What the Integrated Model adds is the explicit inclusion of cultural and leadership conditions among them: the tolerated weaknesses in risk controls, the concerns no one felt able to raise, and the accumulated adaptations bridging procedure and reality, which cultural misalignment and absent leadership allow to develop.

Reason's structure has no place for the individual characteristics of the people in the system. The Integrated Model adds one: personality factors. They sit alongside the causal chain rather than at its start, and they act on it at three points. They influence unsafe decisions directly, through cognitive style, risk tolerance, the accuracy of risk perception, and the experience and judgement a person brings, along with the changing states examined in Chapter Thirteen: fatigue, stress, pain and emotional state. They influence local workplace factors, through how a team works together, how supervision is exercised, and how people respond to the conditions around them. They feed into organisational factors themselves, because an organisation is made up of people, and the characteristics of those who set its strategy, lead it and shape its culture become part of what those factors are. Placed this way, personality factors are part of the system rather than the origin of the event. They shape how people respond to conditions and how conditions come to be, but they do not determine the outcome on their own. The same characteristics produce different outcomes in different organisations, because the organisational and local conditions around them are different. Leaders' own characteristics are part of this too. They flow into the organisational factors, which is why self-awareness is the first leadership capability, and why the leadership assessment described in Chapter Seventeen matters.

Leadership appears in the model twice, because leaders sit in two places. Senior leadership is one of the organisational factors: what senior leaders reward, resource and tolerate becomes part of the foundation everything else stands on. Frontline supervision is one of the local workplace factors: the supervisor is part of the conditions a crew works in, not outside them. Between the two, middle managers carry organisational factors into local conditions. Leadership is therefore not a separate influence acting on the model from outside. It is inside it, at the two points where the conditions for everything above are set.

The model's causal arrows all point upward. They show how harm travels. The Integrated Model adds one arrow that points the other way. Information about what is happening at the sharp end, the near-miss, the concern, the weakened control someone has noticed, the workaround a crew has been living with, either travels back down to the supervisors and senior leaders who can change the conditions, or it does not. That return path is psychological safety, the condition examined in Chapter Six. It is how an organisation finds its own latent condition pathways before they reach the defences, and it completes leadership's place in the model: leaders shape the model through the conditions they set, and they learn about it through what comes back.

What the structure cannot show, on its own, is the quality of that leadership, and that is what the nine constructs of the Transformational Safety® Leadership framework in Chapter Ten describe. The same model behaves very differently depending on where its leaders sit on that range. Under passive and avoidant leadership, the Invisible Man and the Fireman, latent condition pathways develop unchecked, because no one is looking and the return path is close to silent, and leadership engages only once a pathway has produced a loss. Under transactional leadership, the Policeman and the Dealer, unsafe decisions are controlled while people are being watched and rewarded, and the return path carries what the system requires and little more, so the organisational factors and latent condition pathways are left untouched: the defences weaken from underneath while behaviour at the sharp end looks compliant. Under transformational leadership, the Motivator, the Knight, the Thinker, the Carer and the Missionary, the organisational and local factors themselves change, latent conditions are found and closed because the return path carries the truth, and compliance becomes the participation examined in Chapter Eleven. The model shows how accidents happen. The quality of leadership decides how often a trajectory makes it all the way from hazard to loss.

Some of the constructs also act on the defences directly. The Carer maintains the human defences: a leader who knows when someone is fatigued, stressed or not properly prepared can stop those conditions from opening holes in the people who are themselves part of the defences. The Thinker finds latent conditions before they line up, by asking how the organisation's own decisions might be weakening its protection. The Knight keeps stated values and actual resource decisions aligned, so that leadership itself does not create the latent conditions it later has to investigate. The Motivator reduces unsafe

decisions through engagement rather than enforcement, because people who care about the outcome do not need to be watched to protect it.

Still Defences: Vibration, Resonance and Containment

The Swiss Cheese Model is usually drawn as a row of fixed slices, each with fixed holes, and an accident as the rare moment when the holes happen to line up. Reason himself described something more dynamic: holes in the defences continually opening, closing and shifting position. That movement is easy to lose when the model becomes a diagram, and it is the part that matters most for leaders. I have developed it into what I call Cultural Stability Theory.


Picture each defence not as a fixed slice but as a plate that is always moving, and moving in two ways. It moves laterally, through a third plane across the path of the hazard, carrying its holes with it. And it moves forwards and backwards, closer to and further from the defences on either side, so that the way energy passes between them is always changing too. A defence that barely moves keeps its holes where they are. They are known, they can be managed, and they do not suddenly line up with the holes in the defences beside it. The faster and more erratically the defences move, the more likely it becomes that holes in several of them line up at the same moment, and that a trajectory runs from the hazard all the way to a loss. This is why two organisations with similar defences, and similar weaknesses in them, can have very different safety records. The defences do not have to be poor for an accident to happen. They only have to be moving. The question a leader needs to ask changes accordingly: not only where the holes are, but how fast the layers are moving, and what controls that movement.

The energy that sets the defences moving comes from the lower part of the Integrated Model: from organisational and local workplace factors. Production pressure, change, turnover, fatigue, deferred maintenance and the slow drift described in Chapter One all put energy into the defences. Some of that energy arrives in a rhythm: the annual shutdown, the end-of-quarter push, the budget cycle. When pressure arrives in the same rhythm, year after year, it strikes the same defences at the same point in the cycle, and the movement it produces builds on itself. That is resonance in its engineering sense: a system amplifying its own vibration because the pressures acting on it keep arriving in

step with it. Organisations that see the same kind of event recur at the same point in the calendar are usually looking at resonance.

POORLY CONTAINED

The defences move, and their holes line up



WELL CONTAINED

The defences stay still, and the hazard is stopped



Figure: Cultural Stability Theory. Defences that move, across the hazard's path and towards and away from each other, let their holes line up. A containing safety culture holds them still.

What keeps the defences still is the quality of the safety culture around them. Cultural Stability Theory treats safety culture as the stabilising mechanism of the whole system: the thing that decides how much, and how fast, the defences move. The best word for how it does that comes from psychology: containment. A container takes in pressure and anxiety and holds it, so that it is not passed on raw. Chapter Five described culture's role in

managing anxiety; containment is that role applied to the defences. A containing culture takes in the negative energy that would otherwise set the defences moving, and that, arriving in rhythm, builds into resonance. The pressure is still felt. It is simply held rather than passed down to the sharp end. This is where the leadership band in the Integrated Model does its work. Under passive and avoidant leadership there is almost no containment, and the defences move freely. Under transactional leadership containment holds only while someone is watching; when attention moves elsewhere, the vibration returns. Under transformational leadership containment works at the source, because the culture itself takes in the pressure rather than handing it on. Psychological safety, the return path, is how the organisation feels the vibration: the concern raised about a control that is starting to slip is the first sign that a defence has begun to move.

James Reason⁹⁹, whose model this chapter builds on, made a related point in 1998. He described safety culture as the engine of a safe system: the force that keeps driving it toward the greatest possible resistance to its hazards, even through long periods when nothing has gone wrong. Containment is the same idea seen from the side of the defences. The culture is not one more barrier in the row. It is what keeps the barriers where they should be.

One point needs to be clear, because it is easy to misread. Still defences are not a rigid organisation. Chapter Seven argued that strong but rigid cultures fail because they cannot adapt, and nothing here contradicts that. The aim is for the defences to be stable while the culture around them stays alert and adaptive. Containment does not lock a structure in place. It holds energy so the structure can keep working as intended. The same is true of a transformational safety culture: it notices, adjusts and learns, precisely so that the defences do not have to move.

For leaders, this changes the question. It is not only whether the organisation's defences are good enough. It is how much they are moving, what is putting energy into them, and whether the culture around them is containing that energy or passing it on.

What the Model Reveals

The Integrated Model reveals what the linear causal chain conceals: how far the behaviour that precedes a workplace accident is the predictable response of a reasonable person to the conditions the organisation has created.

It is not mainly a product of the individual's inadequacy, negligence or appetite for risk. It is mainly a product of the conditions, organisational, cultural, local and cognitive, in which that worker was operating when the decision was made.

This has an immediate consequence for incident investigation. An investigation that identifies the unsafe decision, or the immediate failure, and stops there has not found the cause. It has found the final expression of a causal chain that, in most serious events, began in the organisational factors months or years before the event. An investigation that actually serves to prevent recurrence has to work back down the model: from the unsafe decision to the local workplace factors that made it likely, and from there to the organisational factors, including the cultural and leadership conditions that made the unsafe decision the predictable, rational and adaptive response of the worker involved. And it has to trace the latent condition pathways, the ways the defences had been weakened directly, whatever anyone at the sharp end did. This is the reluctance to simplify described in Chapter Twelve, given a map.

The Integrated Model also shows why the most common responses to serious events, improved procedures and individual retraining, are inadequate to preventing similar ones. Procedure improvement works on local workplace factors: it changes the instructions available to workers in similar conditions. Retraining works on the individual: it changes the knowledge, and perhaps the attitudes, of the person whose unsafe decision preceded the event. Neither touches the organisational factors that produced the local conditions, created the production pressures, established the cultural assumptions and built the leadership environment in which the decision was the rational one. Neither closes the latent condition pathways that were weakening the defences all along. Those conditions stay unchanged, and the next event goes on developing.

Personality Factors and the Individual in the System

Including personality factors in the Integrated Model reflects a commitment to causal completeness that a purely systemic view of accident causation sometimes lacks. Individual differences in safety behaviour are real,

they are predictive, and they have legitimate implications for how organisations select and develop people. The question is not whether personality factors matter. It is how to weigh them within the fuller model.

The research on personality and safety consistently identifies a cluster of individual differences with the most robust relationships to safety outcomes: conscientiousness and agreeableness from the broad personality model; safety-specific variables such as safety motivation and safety knowledge; and domain-specific variables such as the accuracy of risk perception and cognitive style under stress¹⁰⁰. These variables have real predictive value. Workers with higher safety motivation, more accurate risk perception and greater conscientiousness in the safety-relevant sense do behave more safely across a range of conditions. But the effect sizes are modest, and they are consistently moderated by the quality of the work environment¹⁰¹. Personality predicts safety behaviour more strongly in poor safety cultures than in strong ones, because strong cultures, by establishing more consistent conditions for safe behaviour, reduce the variation personality can account for^{102,103}.

The practical implication is that personality factors matter most at the extremes: in identifying people at substantially elevated risk in safety-critical roles, and in understanding why particular individuals respond to common conditions in unusually risky ways. For the workforce as a whole, personality factors matter less than organisational and local workplace factors, and investment in those conditions will produce larger safety improvements than equivalent investment in individual selection and development.

Including personality factors is also meant to prevent the opposite error from the one the domino model promotes. The domino model over-weights individual factors by placing them at the origin of the causal chain. A purely systemic model that treats workers as interchangeable inputs to a system that entirely determines their behaviour under-weights them, ignoring real variation between people. The Integrated Model holds both truths: the system is the primary determinant of safety outcomes, and the people within it are not all the same. Safety leaders who want to build safe organisations need to attend to both.

Where Intervention Produces Leverage

The practical value of the Integrated Model lies less in its theoretical completeness than in the guidance it gives on where to direct safety leadership attention and investment. By making visible the full causal architecture of safety outcomes, from the organisational factors at its foundation to the unsafe decisions closest to the event, it shows where intervention produces the most leverage.

The evidence consistently points to organisational factors: the foundation on which the cultural assumptions, leadership behaviours and strategic decisions that shape everything above them are set. Changes made there flow up through the whole model, improving local workplace conditions, closing latent condition pathways, and changing the conditions in which people act. Changes made only close to the event, to procedures or to individuals, produce local improvements that do not flow back down, and that remain vulnerable to the organisational factors that have not changed.

The Integrated Model also shows why safety leadership has to be developed at every level. The organisational hierarchy maps onto the model from the foundation up. Senior leaders act mainly on organisational factors. Middle managers turn those factors into local conditions. Supervisors are part of the local workplace factors themselves. A leadership development effort confined to one level reaches one part of the causal chain. One that reaches every level reaches all of it. The research supports each part of that mapping. Studies of first-line supervisors, including early work by Marcel Simard and Alain Marchand¹⁰⁴, have found their leadership style, and in particular their visible concern for the wellbeing of their teams, to be among the strongest influences on safety at the frontline. Research on site and installation managers in the offshore oil and gas industry, by Rhona Flin and her colleagues¹⁰⁵, has linked their commitment and leadership style to the safety behaviour of the workers below them; and Longford is a reminder of where accountability for the organisational factors finally rests: with the senior leaders who set them.

Seen through the Integrated Model, the safety leader who spends most of their attention managing the safety management system, maintaining procedures, conducting audits and tracking corrective actions, is investing where the model shows the least leverage. The leader who spends it on organisational and cultural factors, the quality of the safety culture, the consistency of safety leadership behaviour, and the alignment between what

Beyond Compliance

the organisation says it values and what its operational decisions show, is investing where it shows the most.

Where a leader directs their attention is, in the end, what decides the safety culture their organisation builds. Choosing the foundation is a hard choice, because the work is less visible, harder to measure, and needs a sustained commitment the project-based approach to safety improvement does not accommodate. It is also the choice the evidence, the research and the history of industrial accidents consistently point to. The Transformational Safety Integrated Model shows, with precision, which parts of the system leadership is responsible for shaping, and which of them, shaped well, change everything above them.

PLEASE RESPECT COPYRIGHT

CHAPTER SEVENTEEN

The Maturity Journey: From Lethal to Transformational

The assessment results arrived on a Thursday afternoon as a radar chart: nine axes, each labelled with a Transformational Safety® Leadership construct, each plotted twice, once from the leader's own responses and once from the responses of twelve people who reported directly to him. He had been a safety manager for eleven years and a safety director for four. He had completed the assessment with the confidence his experience seemed to warrant, quickly and without much hesitation, because he believed he already knew what it would show. He had not expected to be uncomfortable with what he found.

The chart told a consistent story, though not the one he expected. The Invisible Man sat low, in the best zone that construct allows. Good. The transactional constructs, the Policeman and the Dealer, sat squarely in the green. Also broadly expected, and broadly comfortable. The Fireman was not. He had rated himself low on it. His people had placed it well into the red. Then the transformational constructs. Only the Knight reached amber in his team's ratings. The Motivator, the Thinker and the Carer all sat in the red. The Missionary was the lowest-rated construct of all.

Then he looked at his own line. On every transformational construct, his self-rating sat well above the ratings of the people he led, in places by three or four points. He had placed himself in the green as a Knight and a Carer. His people had placed him nowhere near it. The Fireman was the only construct where the gap ran the other way, and it was the one that explained the rest. He saw himself as a leader who was always present. They experienced a capable, fair and rather distant Policeman, who was seen most clearly when something had gone wrong: after the incident, at the investigation, in the meeting that followed. It was not that they thought badly of him. On the transactional

The Maturity Journey: From Lethal to Transformational

constructs their ratings and his were almost identical. They agreed he was fair, consistent and clear about the rules. What they did not see was the part of his leadership he thought mattered most. And twelve people, answering separately and without knowing one another's responses, had seen the same thing..




Figure: A Transformational Safety® Leadership profile, comparing a leader's self-rating with the ratings of his twelve direct reports. Green is the target range: you cannot see too much green. Amber can be lived with, provided development is under way to move it toward green. Red is where no leader wants to be found. For the passive constructs, the lowest band is amber, because any real presence of the Invisible Man or the Fireman is a concern.

He sat with the chart for a long time. The numbers surprised him, and the surprise was the first lesson. Fifteen years in safety, and he had not known how his own leadership landed on the people closest to him. The longer he looked, the less the numbers troubled him than what they implied about the safety culture he had spent eleven years building. A culture in the Developing range: systematically managed, procedurally disciplined, reliably compliant. A

culture whose leadership had invested heavily in the system and modestly in the people. A culture that was, in the language of the maturity model, exactly where its leadership had placed it. The system he had built was not the problem. By most measures it was a good one. The problem was that, somewhere along the way, he had come to believe the system was the culture.

That recognition, precise, undefended and uncomfortable, was the beginning of his development as a safety leader. Everything before it had been safety management.

The Maturity Landscape

The Transformational Safety® Culture Maturity Model describes not a destination but a landscape: a map of the ground safety culture occupies, from its most dangerous forms to its most effective. Each level is marked by a particular set of cultural assumptions, leadership behaviours and organisational conditions, which produce, with considerable reliability, the safety outcomes associated with that level. Knowing which level an organisation currently occupies, what produced it, and what the path to the next level requires is the foundational diagnostic capability of safety leadership.

A landscape can also be crossed in both directions. Organisations do not only climb; they slip. A change of chief executive, a merger, a sustained cost-cutting programme or the loss of a handful of trusted supervisors can move a culture down a level far faster than it took to climb it. The maturity an organisation has reached is not a possession. It is a position that has to be held, and holding it takes the same leadership that earned it.

The model builds on the safety culture ladder developed by Patrick Hudson¹⁰⁶. Working with Shell on its Hearts and Minds programme, Hudson extended Ron Westrum's typology¹⁰⁷ of organisational cultures into a ladder of five rungs: pathological, reactive, calculative, proactive and generative. Westrum's own insight was about information. He showed that cultures can be told apart by what they do with bad news. In a pathological culture, the messenger is punished. In a bureaucratic one, the messenger is tolerated and the message is filed. In a generative one, messengers are actively sought out and valued. The same distinction runs through every level of this model, and through the return path in Chapter Sixteen. Hudson's central insight was that safety cultures develop through recognisable stages, each with its own characteristic attitudes to information, blame and responsibility, and that an

The Maturity Journey: From Lethal to Transformational

organisation cannot skip a rung. The ladder has lasted in high-hazard industries largely because it lets organisations describe where they are in terms their own people recognise. The Transformational Safety® model keeps that insight and the five-stage structure. What it changes is the lens. Each level is defined by the leadership that produces it, and each transition is described by the leadership capability it requires. The colours follow the familiar convention, from the red of the Lethal level to the green of the Transformational.



The five levels are not arbitrary divisions of a continuous spectrum. They are qualitatively different states of safety culture, distinguished not mainly by the sophistication of the safety management system but by the depth at which safety values are embedded in the organisation's basic assumptions. Each transition requires something different from the one before, and the leadership capability needed at each level is not simply more of what was needed at the level below. It is a different kind of capability. The early transitions are largely achievable through management. The later ones are not.

Nor is the model a straight line that organisations move along in sequence. Organisations can sit at different levels in different parts of their operations at the same time: a Functional culture at corporate level and an Avoidant or Developing one at a remote operation, or a strong culture in one

business unit and a hollow compliance performance in the one next door. That variation is often as important as the average, because it shows which parts of the organisation have the leadership conditions that safety culture development needs, and which do not.

The Lethal Level

The Lethal level is, in the blunt language of the model, the level at which organisations are dangerous. There is no systematic safety management. Basic protections are absent or inadequate. The cultural assumptions governing the organisation's relationship to safety are essentially indifferent: harm is an expected part of the work, the organisation's responsibility is limited to the minimum the law requires, and the people most exposed to its hazards have the least access to protection and advocacy.

The Lethal culture is not mainly marked by malice. It is marked by an absence of understanding: of the knowledge, capability and organisational development that systematic safety management requires. Leaders at this level are rarely people who have decided to trade workers' safety for commercial advantage. More often they have inherited organisations whose relationship to safety has never been seriously examined, lack the knowledge to see what that examination would reveal, and work from the assumption that harm is a feature of their industry rather than a preventable outcome of their system.

The Lethal organisation is recognisable to anyone who has met one: safety rules that exist in name but not in practice, protective equipment required in policy but not enforced in operation, incident reporting that produces records without producing learning, and a workforce that expects harm to occur with a regularity the management system cannot prevent. The leadership profile is dominated by the Invisible Man and the Fireman: leaders functionally absent from the safety life of the organisation, who engage with safety only after events that cannot be ignored.

The move from Lethal to Avoidant requires the basic architecture of safety management: hazard identification, minimum risk controls, an incident reporting mechanism, and the beginnings of a leadership orientation toward safety as something that can be managed rather than an accepted hazard of the job. This is mainly a management challenge. It is achievable through investment of resources and expertise in the objective infrastructure of safety

management, and organisations that make it produce real improvements in safety outcomes. It is the beginning of the journey, not its end.

The Avoidant Level

The Avoidant culture is driven by avoidance. It manages safety mainly as the avoidance of harm: of incidents, injuries, and the regulatory and reputational consequences they bring. Its safety management system exists and is maintained, but its main function is managing events that have occurred rather than preventing conditions that are developing. Its leadership responds strongly to visible safety problems. The Fireman is prominent: leaders who engage actively after events and less consistently through the long ordinary periods between them.

The Avoidant culture differs from the Lethal one in a specific way: it knows safety is a problem and is trying to address it. What holds it at this level is the assumption that safety problems present themselves: that the organisation's task is to respond well when safety issues become apparent, rather than to create the conditions in which they can be detected before they do. That assumption produces safety intelligence that is mainly retrospective, built from events that have already happened, from patterns in the incident register, and from the findings of periodic external review.

The move from Avoidant to Developing requires a shift from retrospective to prospective safety management: systems that produce information about current conditions rather than simply recording past outcomes. That means investment in lead indicators, near-miss reporting, risk assessment that engages with operational reality, and audits that test whether risk controls work rather than whether they exist. It is still mainly a management challenge, but the leadership dimension begins to show. An Avoidant culture that cannot build the reporting culture lead indicators depend on has met a problem better systems alone cannot solve.

The Developing Level

The Developing culture is system-driven. It has a comprehensive safety management system and applies it with a consistency and discipline the Avoidant culture has not achieved. Its lead indicators are functioning: near-

miss reporting runs at reasonable rates, hazard identification is part of operational routine, and risk assessments are done at the required intervals. Its metrics are tracked and reviewed, its corrective actions managed, its audit programme operational. By the standards of most external assessments, it is a well-managed organisation.

It is also the level at which the difference between safety management and safety leadership matters most and is hardest to see. The system works well enough that the gaps in the culture, between what the system measures and what actually determines safety outcomes, and between administered safety performance and the lived safety experience of the workforce, do not appear in the formal data. The organisation measures what its system is designed to measure, and what it measures is improving. The conditions a deeper diagnosis would reveal, the underlying assumptions, the values enacted under pressure, the quality of leadership at the frontline, are not in the data, and so they are not in the conversation.

The assumptions of the Developing level carry a particular and stable overconfidence: the belief that if the system is working as designed, the safety performance it produces is adequate. This is the confident miscalibration described in Chapter One: a model of the organisation's own safety performance consistent enough with the available data to feel accurate, while being systematically wrong about the things the data does not capture. Leaders at this level are not complacent. They are engaged, diligent and committed to the system they have built. What they are not doing is asking whether the system's performance is an adequate measure of the culture's health.

The move from Developing to Functional is the most demanding transition in the model, because it requires something no earlier transition has: a shift in orientation from management to leadership. A Developing culture can go on improving its safety management system indefinitely without making this move. Making it requires leaders who will move from administering the system to diagnosing the assumptions beneath it, from tracking the metrics the system produces to understanding what those metrics are not telling them, and from securing compliance to building engagement.

The Functional Level

The Functional culture is driven by continuous improvement. It has moved beyond administering safety performance to engaging with it: an

ongoing inquiry into whether the current approach to safety is adequate to the conditions the operation is presenting, and deliberate improvement of that approach when it is not. Its workforce is involved in developing safety solutions, not as a communication strategy but because the operational knowledge of workers is an irreplaceable input to improving safety conditions.

A Functional culture looks and sounds different from a Developing one, in ways anyone who has worked in both will recognise. Workers raise safety concerns in the informal channels, in conversations before and after shifts, with a frequency and specificity that show engagement rather than performance. Near-miss reports describe the conditions that produced the near-miss rather than the behaviour that preceded it, because the culture looks for systemic conditions rather than the individual whose behaviour was closest to the event. Leaders ask about how work is actually done rather than whether procedures are being followed.

The leadership profile at the Functional level shows transformational constructs alongside the transactional foundation. The Motivator is visible and credible. The Knight's consistency between stated values and decisions under pressure is established and observed. The Thinker is creating the forums and the permission for real safety thinking. The Carer is known and trusted across the workforce, not as a management style but as a quality of relationship. Operating together and consistently, these constructs move the culture from system-managed safety, the mark of the Developing level, to the people-driven safety the Functional level depends on.

The move from Functional to Transformational is the longest and least predictable in the model. It cannot be managed the way earlier transitions can. It needs time, consistency of leadership behaviour, and the accumulation of shared experience that changes cultural assumptions: not just espoused values and artefacts, but the basic, taken-for-granted beliefs about what the organisation is and what it is for. Organisations that reach the Transformational level usually cannot point to a programme or intervention that got them there. They can point to a period of sustained, consistent, culturally intelligent leadership that added up to it.

The Transformational Level

The Transformational culture is value-driven. Safety is part of the heart and mind of the organisation: not as a policy, a commitment or a priority, but

as part of its identity, so fundamental that separating safety from its other core values would make no sense. The organisation does not have a safety culture as one component of its wider culture. It is a safety culture. The distinction between managing a safe organisation and being one has dissolved.

The difference from the Functional level is harder to describe than the differences lower down, and easier to recognise when you meet it. It lies in how natural safety is: in the absence of the deliberate, effortful quality safety commitment has at every lower level. In a Functional culture, workers are involved in safety, engaged with it and contributing to it. In a Transformational culture, the line between their work and their safety is hard to find, because safety is not something done in addition to the work. It is the way the work is done. Unsafe decisions feel wrong, not merely non-compliant.

The leadership profile at this level is marked by the Missionary. Not because the other constructs are absent, but because the Missionary's defining quality, the full integration of safety values into professional identity, has become the cultural norm rather than the exception. Leaders at this level do not talk about safety as a separate part of their leadership. They talk about operations, and safety is in everything they say about operations. They do not visit the worksite to check safety conditions. They visit because they are interested in the work, and that interest includes safety as part of what the work is.

..... A Transformational Safety[®] culture is not achieved by managing safety better. It is achieved by becoming the kind of organisation for which safety is not a management challenge but a way of being.

Diagnosing Your Current Position

Knowing where an organisation sits on this landscape is among the most important questions a safety leader can ask, and among the hardest to answer accurately, because every mechanism this book has described for producing false confidence also distorts the answer. A leader's own reading needs at least three sources: the formal safety data, read for patterns rather than headlines;

conversations with workers about how concerns are actually raised and received; and direct observation of work in its unmanaged moments, moving through the layers of culture described in Chapter Four. For the organisation as a whole, those readings need the structured assessment described in Chapter Eighteen, which hears from everyone, not only from the people a leader happens to meet.

What the Gap Reveals

There is a fourth source of evidence, and it is the one leaders find hardest to use: the gap between how they see their own leadership and how the people they lead experience it. The Transformational Safety® Leadership assessment is a 360-degree instrument for this reason. A leader's self-rating tells them what they intend. The ratings of their people tell them what actually lands. Where the two agree, the leader has an accurate picture of their own influence. Where they differ, the difference is the most useful thing in the report.

The direction of the gap matters. Research on self–other agreement in leadership, notably by Leanne Atwater and Francis Yammarino¹⁰⁸, has found that leaders who rate themselves substantially higher than their followers do tend to be less effective than those whose self-ratings are accurate or modest. The reason is not hard to see. A leader who overestimates their own consistency, approachability or visibility has no reason to change, and their people have learned that saying so is not worth the risk. In the language of the Integrated Model, the return path is quiet, and the leader has taken its silence for agreement.

The most common pattern, in organisations at the Developing level, is the one in this chapter's opening. The transactional constructs are rated much the same by leader and followers. The transformational constructs are rated markedly higher by the leader. Often, there is a passive construct that runs the other way: the leader rates themselves lower on the Fireman than their people do. This is the leader who believes they are present, and is experienced mainly in the aftermath of events. Leaders see their intent. Their people see their behaviour under pressure, and they notice when it arrives. A leader who believes they are operating as a Knight while their people experience a Fireman is not being dishonest. They are judging themselves by their values, and being judged by their decisions.

Beyond Compliance

What to do with the gap is simpler than facing it. First, treat it as information rather than criticism: it describes how leadership is received, not what the leader is worth. Second, choose one or two constructs where the gap is largest and the effect on the culture greatest, rather than trying to change everything at once. Third, ask the people who rated you what they would need to see differently, and make it clear that the question is real. And fourth, reassess, because the gap closes only when behaviour changes consistently enough, and for long enough, for people to notice. For a leader who acts on it, a large gap is not a verdict. It is a starting point, and the most accurate one they will ever be given.

The safety director with the radar chart did not change the culture he had built by redesigning its management system. He changed it by taking his own chart seriously. The Missionary showed the largest gap of all, but it is built from the constructs beneath it, so he started with the next two: the Carer and the Thinker. Then he went back to the twelve people who had rated him to ask what they needed to see. That was where the move from Developing to Functional began for his organisation: not in the system, but in one leader's willingness to be seen accurately.

PLEASE RESPECT COPYRIGHT

CHAPTER EIGHTEEN

Reading the Organisation: Assessing Safety Culture and Leadership

The senior management team were sure of one thing. Safety, they said, was among the very highest priorities in the business, and when their own survey responses came back, they said it more strongly than any other group in the organisation. They were not wrong to believe it. They had invested in systems, in training and in people. They meant what they said.

Three months later, small groups of workers drawn at random from across the business sat in a room with an outside facilitator and talked about what safety was actually like where they worked. The conversations covered a great deal, but one theme came back in group after group. Many of them believed that the main purpose of an incident investigation was to find someone to blame. Some said they had been told, by their own managers, that there was no such thing as an accident: if something went wrong, someone had made a mistake.

Both accounts were sincere. Both described the same organisation. The senior team were describing the culture they intended and could see. The workers were describing the culture they experienced. Neither on its own was the truth about that organisation's safety. Together, read carefully, they were.

This chapter is about reading an organisation in that way: not through the lens its leaders already hold, and not through a single survey score, but through a structured, honest account of the culture and the leadership as the whole workforce experiences them.

Why Organisations Misread Themselves

Earlier chapters have described the ways organisations come to believe things about their own safety that are not true: the illusion of safety, the virtual management system, the confident miscalibration of leaders whose information has been filtered on its way up to them. Each of these is, at heart, a failure of reading. The organisation looks at itself and sees the wrong things, or sees the right things through the wrong lens.

There are three reasons this happens so reliably. The first is that the information that reaches leaders is shaped by the system that produces it, and that system is usually built to show activity and compliance rather than conditions and belief. The second is that leaders read what they find through their own experience of the organisation, which is the experience of people at the top of it, with the authority, information and freedom that position brings. The third is that the people who know most about the organisation's real safety, the people doing the work, are rarely asked in a way that makes an honest answer feel safe.

Reading an organisation accurately means correcting for all three. It means gathering information that the safety management system does not produce. It means hearing from every part and level of the organisation, not just the part that leaders already see. And it means asking in a way that protects the people who answer.

Two Instruments, One Picture

The Transformational Safety® approach to organisational assessment uses two survey instruments, administered together, and then sets their results alongside what people say in their own words.

The first, the Transformational Safety® Culture Survey, asks people to describe the safety culture they work in across seventeen factors. Some of these concern the organisation's systems and practices: communication, competence, the management of change, compliance with safety systems, safety rules, and the way accidents and incidents are handled. Some concern the climate people work in: the supportiveness of the environment, the work environment itself, co-operation, involvement, and the predominant management style. Whilst some reach toward belief: the personal priority people give to safety, their appreciation of risk, the priority they see the

organisation giving it, management commitment, safe behaviours and shared values.

The second, the Transformational Safety® Leadership Survey, asks people to describe the safety leadership they experience across the nine constructs set out in Chapter Ten, from the Invisible Man through to the Missionary. It is the same model as the individual 360-degree assessment instrument described in Chapter Seventeen. The difference is the question being asked. The individual assessment asks how one leader is experienced by the people who work for them. The organisational survey asks what leadership feels like across the organisation as a whole: what kind of safety leadership the workforce lives under, wherever they happen to sit.

The two surveys are deliberately read together, because culture and leadership explain each other. A culture factor sitting in the red zone tells you where the problem is. The leadership profile in the same part of the organisation frequently tells you why.

Target Zones, Not Rankings.

Most organisational surveys compare their results with norms: the average of other organisations that have completed the same instrument. That comparison is comforting and truly almost useless. It tells an organisation how close it is to other people's problems. An organisation scoring above the industry average on management commitment may still have a workforce that does not believe its managers mean what they say; it simply has more company than it thought.

The Transformational Safety® surveys are read instead against best-practice target zones: ranges associated with strong safety performance, shown in the traffic-light colours used throughout this book. Green is the target range, and you cannot see too much green. Amber can be lived with, provided development is under way to move it toward green. Red is where no organisation wants to be found; the shorthand used with clients is that you do not want to be "dead in red". The zones are not the same for every measure. For the passive leadership constructs, any real presence is a concern. For the transformational constructs, more is always better. For the Policeman, the zone floats: some enforcement is necessary, and the right amount depends on the regulatory environment the organisation works in, but too much policing produces exactly the watched compliance described in Chapter Ten.

Reading against target zones changes the conversation leaders have about their results. The question stops being how the organisation compares with others and becomes how far it is from where it needs to be, and in which direction it is moving.

Slicing the Data

An organisation-wide average is the least useful number an assessment produces. Every average hides the variation inside it, and Chapter Seven argued that the variation is often the most valuable diagnostic information an organisation has. So, in order to maximise true value the results should always be sliced.

The first cut is usually by location: each site, facility or major operation read on its own. Within each location, the data is then sliced again by the characteristics most likely to separate one experience of the organisation from another: department, level in the hierarchy, length of service, age, and working arrangement, such as people on rotational rosters compared with those on regular hours. Each slice produces its own culture and leadership profile, read against the same target zones.

The slices show things no average can. They show where the stronger and weaker sub-cultures are, which tells leaders both where development is most needed and where existing strength can be built on. They show how the experience of leadership changes within the hierarchy, and at which level it breaks; Chapter Eleven showed that engagement falls between the boardroom and the worksite, and the slices show whether the quality of safety leadership falls with it. They show whether newer workers are experiencing a different organisation from the one longer-serving workers describe, which is often the first sign of cultural drift. And they show the organisational version of the gap Chapter Seventeen described in individual leaders: the distance between how the leadership team sees the organisation and how everyone else experiences it.

In the organisation that opened this chapter, the slices told a consistent story. Across a city office, an onshore facility and two offshore platforms, the Policeman and the Dealer sat in the green zone however the data was cut. Every transformational construct sat in amber. Most of the culture factors sat in amber too, with change management repeatedly drifting toward red. And the senior management team rated the priority of safety more highly than any

other group in the business. It was a clear Developing profile in the terms of Chapter Seventeen: a well-run compliance system, led by capable transactional leadership, without the transformational leadership that turns compliance into participation.

Hearing the Reasons

A survey tells you what people think. It rarely tells you why. And as Chapter Four showed, surveys mainly reach the espoused layer of culture: what people know they are supposed to believe, which is not always what governs their decisions. So a survey should never be relied upon alone.

Alongside it, smaller focus groups are drawn at random from across the workforce and facilitated by someone outside the organisation's management structure. The randomness matters, because it stops the groups being filled with the people management would choose to hear from. The outside facilitator matters, because it makes honesty safer. Also, the group format matters more than it might seem. In an individual interview, people describe their own views. In a group, they react to each other, use their own language, argue, agree and fill in each other's stories, and in doing so they show which issues carry real weight and why. The gap between what people say on a survey and what they actually do starts to become visible in the room.

Where the survey and the focus groups point in the same direction, the finding can be trusted. Where they diverge, the divergence is itself information. In the organisation that opened this chapter, the survey showed accidents and incidents being taken seriously at every level; people did not doubt that the organisation cared when something went wrong. It was the focus groups that explained why the same measure sat stubbornly in amber: the care was real, but the investigations were experienced as a search for someone to blame. No survey item would have surfaced that. It took people talking to each other.

What the Survey Cannot Tell You

A well-run survey is a powerful instrument, and it has limits that leaders need to understand before they read a single result.

It captures a moment, not unlike a photograph. An organisation surveyed in the weeks after a serious incident, in the middle of a restructure, or during

a bitter industrial dispute will describe itself differently from the same organisation in a calmer period. That does not make the results wrong. It means they have to be read with the moment in mind, and it is usually better to assess when the organisation is not in crisis, so that what is measured is the culture rather than the reaction.

It depends on who answers. In the organisation that opened this chapter, ninety plus percent of the workforce took part, which is high enough to draw conclusions with real confidence. When response rates are low, the people who do answer are often those with the strongest views, in either direction, and the silence of the rest is itself information about the return path. Contractors, who in many high-hazard operations do much of the most dangerous work, are too often left out of assessments altogether. A picture of the culture that does not include them is a picture of only part of the organisation.

It also depends on whether people can answer honestly. A workforce that does not trust the anonymity of the process, or that has learned from past surveys that nothing changes, will give the answers it thinks are expected. Literacy, language and the way questions are administered all matter too, particularly in diverse and multilingual workforces. None of this argues against surveying. It argues for surveying carefully, for reading the results alongside what people say in their own words, and for treating a survey as one voice in the account rather than the account itself.

Measuring the Return Path

Chapter Sixteen described psychological safety as the return path in the Integrated Model: the route by which what is happening at the sharp end travels back to the people who can change it. An organisational assessment can measure that path directly.

The approach I use distinguishes between two things that are often confused. The first is a person's own experience: whether they personally feel able to raise a concern, admit an error or challenge a decision. The second is their reading of the climate around them: whether they believe it is safe for people in general to do those things here. The two do not always agree. An organisation where most people feel personally safe but believe the climate is unsafe for others has a different problem from one where people believe the climate is supportive but would not themselves speak up. Measuring both, and

the gap between them, tells leaders not only whether the return path is open but where and why it is blocked.

Alongside this sit two further measures. The first is whether people can point to safety champions around them: individuals at any level who visibly take safety beyond what the system requires. The second is engagement with safety itself, the discretionary participation described in Chapter Eleven. Read together with the leadership profile, these show whether the organisation's leadership is producing compliance or participation, which is the question that matters most.

Reading One Slice

It helps to see how these pieces come together for a single part of an organisation. The example that follows is illustrative, drawn from patterns that recur across many assessments rather than from any one organisation.

Consider a maintenance crew working rotational night shifts on a large operating site. Across the organisation as a whole, the culture survey shows the management of change in amber. For this crew it sits in red. The leadership profile for the same slice shows the Fireman noticeably higher than elsewhere on the site, and the Carer noticeably lower. Few people in the crew can name a safety champion among their own leaders. On their own, these are numbers. Together, they already suggest a story: a crew that experiences change arriving without warning, and leadership that appears mainly when something goes wrong.

The focus group fills in the rest. The crew describe procedure changes and equipment modifications that are decided on day shift and handed to them by email at the start of a night, with no one available to explain the reasoning or answer questions. They describe concerns raised about those changes that seem to disappear. And they describe the only reliable way to get a senior leader's attention as an incident. Nobody in the group describes their supervisor as uncaring. They describe someone who is stretched across too many crews to know any of them well.

Read together, the slice tells the organisation something important and specific. The problem is not this crew's attitude to safety, and it is not a gap in the procedures. It is a set of conditions created by the way the organisation manages change and spreads its frontline leadership: exactly the kind of latent condition described in Chapter Sixteen, sitting in a part of the organisation

leaders rarely see at the hours the work is done. The return path from this crew is close to silent, and the defences in their area are moving more than anyone above them knows.

That reading points directly to what needs to happen. The management of change process needs a step that reaches night shift properly. The crew's supervisor needs a span of leadership that makes knowing people possible, and development in the Career. And senior leaders need to be seen on the night shift before the next incident, not after it. None of that would have come from the organisation-wide average, which in this example looked unremarkable.

Reading the Picture Together

The value of the assessment is not in any single number but in reading the pieces together into a coherent account of where the organisation is and why. That account is built around three questions.

The first is where the organisation sits on the maturity model of Chapter Seventeen. The leadership profile is often the clearest evidence. A culture whose leadership is dominated by the passive constructs, in which the Invisible Man and the Fireman are recognisable throughout the workforce's account, sits toward the Lethal and Avoidant end, where the most dangerous workplaces are found. A culture whose transactional constructs are strong and whose transformational constructs are not sits at the Developing level: compliance and expected effort, and average safety performance. A culture whose transformational constructs are established across the organisation, not just in a few exceptional teams, is reaching the Functional and Transformational levels, where participation replaces compliance.

The second is where the variation lies. Which parts of the organisation are already building the culture it needs, and which are not? What do the stronger parts have that the weaker ones lack? The answer is almost always found in the leadership profile of the stronger part, and it is the most useful practical finding an assessment produces, because it shows the organisation an example of its own future that already exists within it.

The third is what the workforce is telling the organisation that its leaders did not know. In most assessments there is at least one finding of this kind, often surfaced in the focus groups: a belief about investigations, a frustration with the management of change, a part of the system that everyone at the worksite knows does not work as written. These are the findings that matter

most, because they are the latent conditions of Chapter Sixteen, made visible before they have had the chance to line up.

Bringing Leaders Into the Findings

How leaders first meet the results matters almost as much as the results themselves. An assessment report delivered as a document, read alone and discussed at the next management meeting, tends to produce the defensive reactions Chapter Five described: findings explained away, the method questioned, the uncomfortable parts attributed to a few disgruntled respondents.

A better approach is to bring the leadership team into the findings together, in a facilitated session, before any report is circulated. One exercise is particularly useful. Before seeing the results, each leader predicts how the workforce has rated the culture and the leadership in their own part of the organisation. When the actual results are revealed, the gap between prediction and reality becomes the starting point for the conversation. It is the organisational equivalent of the self-rating gap in the individual 360-degree assessment, and it is just as revealing. Leaders who predicted accurately are usually the ones closest to their people. Leaders who were surprised have learned, in a way no report could teach them, how far their view of their own area has drifted from the experience of the people in it.

The aim of the session is ownership. A leadership team that has worked through its own results, argued about them, and accepted what they show is in a very different position from one that has received a consultant's report. The first has begun to change. The second has received a document.

Handling the Results With Care

An organisational assessment asks people to tell the truth about where they work. How the organisation handles what they say decides whether they will ever tell it again.

Anonymity has to be real, not merely promised. Responses are collected in a way that cannot be traced back to individuals, and no result is reported for a group small enough that its members could be identified. People can tell the

difference between an organisation that protects them and one that only says it will.

The results must never be used to rank departments, set targets for managers or decide rewards. The moment a culture or leadership score carries consequences, people begin to manage the score instead of the culture, for the same reasons Chapter Fifteen gave about Zero Harm targets and Chapter Nineteen gives about individual assessments. The assessment is a way of learning, and it stops working the moment it becomes a way of judging.

The findings belong first to the people who gave them. The most important thing an organisation does after an assessment is go back to its workforce and tell them, plainly, what was heard, what surprised the leadership, and what will be done as a result. An organisation that asks and then goes quiet has taught its people that asking was a performance. One that reports back honestly, including about the findings that were uncomfortable to hear, has shown them something about the return path that no policy statement could.

When to Assess, and How Often

The first assessment is best done before any significant development work begins, because it is the baseline against which everything that follows will be judged. Without it, an organisation has no way of knowing whether its investment in leadership and culture is changing anything, and it will tend to judge success by the measures that are easiest to collect, which are the ones this book has argued are least reliable.

After that, a full reassessment every eighteen months to two years is usually long enough for real change to show, and short enough to keep the work honest. Culture does not move quickly, and assessing too often produces fatigue and the impression that the organisation is more interested in measuring than in acting. Between full assessments, short pulse checks on a few key measures in particular parts of the organisation can show whether specific changes are landing.

What matters most over time is not any single score but direction. An organisation whose transformational constructs are moving from the lower to the upper end of amber, whose red zones are shrinking, and whose leadership gap is narrowing is changing, even if it has not yet reached green. One whose

numbers are static, however respectable, has stopped moving, and Chapter Seven explained what tends to happen to cultures that stop moving.

The First Step, Not the Last

An assessment is a beginning. On its own it changes nothing, and an organisation that commissions one, receives the report, and files it has done something worse than not asking at all.

Its real value is twofold. It gives the organisation an honest starting point, located on the maturity model and grounded in the experience of its own people, rather than the comfortable picture its systems and its leaders' own vantage point provide. Repeated over time, it shows whether the organisation as a whole is changing: not only whether the leaders who have been through development have changed, but whether the culture the workforce lives in has moved.

The organisation that opened this chapter was not short of commitment. Its leaders cared about safety and had invested in it. What it lacked was an accurate picture of what that investment had and had not produced, and in particular of what its own people believed about what happened when something went wrong. The assessment gave it that picture. What it did next is the subject of the work that follows.

PLEASE RESPECT COPYRIGHT

CHAPTER NINETEEN

The Transformational Safety[®] System

There is a moment in the development of every safety leader that is recognisable on reflection, though it rarely announces itself at the time. It is the moment when the ideas stop being about other organisations and other leaders, and become, with uncomfortable precision, a description of this organisation and this leader.

The moment is not always comfortable. The leader who has invested years in safety management, who has built competent systems and maintained them with real diligence, and who arrives at the honest assessment that what they have built is a Developing culture operating well below its potential, is feeling the particular discomfort that development requires. Discomfort in the service of understanding is not a problem to be managed. It is the beginning of transformation: it opens the gap between where the leader is and where they need to be, and it supplies the motivation to close it.

The Transformational Safety[®] System is the structure built to close that gap. It is not a training programme, a certification process or an audit framework. It is an integrated development system, grounded in the research, tested in practice, and designed to produce the kind of safety leadership capability that safety culture transformation requires. This chapter describes its structure, its components, and how they fit together.

The Architecture of Transformation

The system rests on a single insight: developing safety leadership capability is not mainly a knowledge challenge. It is a challenge of self-awareness and identity. A leader who needs to develop the Carer does not need to know more about individualised consideration. They need to see how their current behaviour is experienced by the people they lead, to build the capability to engage with workers as individuals rather than as positions in a structure, and to shift from being a leader who manages safety to a leader who cares about the people who do the work. Knowledge is a precondition of that development. It is not enough on its own.

The system's structure reflects this. It is organised around a personal leadership framework, ATLAS, that begins with the most fundamental questions of leadership identity and builds outward to the skills of putting safety leadership into practice. It uses a diagnostic assessment of each leader's construct profile, providing both the self-assessment and the follower-perception data needed to identify the most important development priorities. Alongside it, the organisational assessment described in Chapter Eighteen shows where in the organisation that development matters most, and whether, over time, it is changing the culture as a whole. It is delivered through a combination of structured development, practice in real work, and coaching that turns understanding into lasting change in behaviour. It is designed to be sustained over the long period real cultural change takes, rather than delivered as a single event.

The components are designed to work together, and each depends on the others. The ATLAS framework provides the foundation for authentic development; without it, new safety leadership behaviours are produced by training rather than by shifts in values and identity, and they do not hold under pressure. The construct assessment directs development effort where it will matter most, rather than into constructs that are already strong. The coaching provides the accountability and reflective support that make development continuous rather than an event. The organisational application connects individual development to the cultural conditions the organisation as a whole needs.

ATLAS: The Personal Leadership Framework

The ATLAS framework takes its name from the figure of Greek myth who, in the popular image, carries the world on his shoulders. The image is not meant as a metaphor for burden. It describes the quality of safety leadership the framework sets out to develop: the understanding that the safety of the people within a leader's sphere of responsibility is theirs to carry, personally. Not as a management obligation or a professional requirement, but as a personal accountability that no system, procedure or audit can take over or relieve.

The framework rests on a foundation of personal leadership: self-awareness, self-management and values. Self-awareness is the honest understanding of how one's own behaviour lands on other people. Self-management is the capacity to stay steady, consistent and clear-headed when pressure, fatigue and frustration push the other way; the same strain Chapter Thirteen described in the workforce affects leaders too, and a depleted leader does not lead safety well, whatever their knowledge or commitment. And values are what make the rest sincere rather than performed. Without that foundation, the seven pillars produce skilled performance without depth.

The Seven Pillars

The seven development pillars of the ATLAS framework are the domains in which safety leadership capability is built. They run from the most foundational, the questions of identity and value that underlie all effective leadership, to the most applied, the implementation skills that turn understanding into consistent practice. The sequence matters. Development that starts with the applied pillars without the foundational ones produces skilled performance without depth: safety leadership behaviours that are competent but not authentic, and that cannot sustain the credibility the Knight depends on.

The first pillar is Who am I? This is the domain of leadership identity: the self-understanding that determines how a leader interprets their role, what they regard as their core responsibilities, and what they will and will not do when operational pressure makes those choices hard. For safety leaders, the question becomes: is safety leadership a function I perform, a competency I have

developed, or an expression of who I am and what I stand for? The answer does more than anything else to decide whether the pillars that follow produce transformation or merely better performance. Leaders who develop their safety leadership from that identity build a different kind of culture from those who develop it as a professional skill.



The second pillar is What Drives Me? This is the domain of values and motivation. It examines what actually drives the leader's engagement with safety: compliance obligation, professional reputation, real commitment to the wellbeing of the people in their care, or some mix of the three. In my experience, and in the research on safety-specific leadership, leaders whose safety engagement is driven mainly by commitment to the wellbeing of others produce different safety cultures from those driven mainly by compliance or reputation, even when the surface behaviours look similar¹⁰⁹. Understanding one's own motivation is the precondition for aligning behaviour with values.

Figure: The ATLAS personal leadership framework. Seven pillars, from identity to application, built on a foundation of self-awareness, self-management and values.

Beyond Compliance

The third pillar is My Emotional Intelligence: the capabilities Chapter Eleven identified as the foundation of transformational safety leadership. They are not traits a leader either has or lacks. They can be developed through structured reflection, deliberate practice and the kind of feedback the construct assessment provides. In safety leadership, developing emotional intelligence means developing what the Carer, the Motivator and the Missionary need: the capacity to understand and engage with the emotional reality of the people whose safety the leader is responsible for.

The fourth pillar is My Leadership Strengths. It turns development away from the familiar deficit model, in which leaders are handed a list of weaknesses to fix, and toward identifying the strengths a leader already has and building on them deliberately, while making sure there is adequate competence everywhere else. It begins with an honest recognition that technical excellence is not leadership. Many safety leaders were promoted because they were outstanding engineers, operators or specialists, and the competencies that made them so are not the ones their leadership role now requires. The 360-degree construct assessment is the main diagnostic tool here: it shows which constructs a leader can already lead from, and which need building.

The fifth pillar is My Leadership Challenges. It helps leaders understand that people meet challenges and opportunities in very different ways, and that a team's safety depends on those differences being understood and used rather than flattened. The framework uses the image of a beach at sunrise, where people approach the same shoreline differently: one digs carefully to discover what lies beneath the surface, while others take quite different approaches to the same water. Each approach brings value, and each carries risk. For a safety leader, the challenge is to know their own default; to value the people whose approach differs from it, and to recognise which approach a particular risk calls for.

The sixth pillar is My Capacity to Absorb Change. Safety culture transformation is change, and change is where most organisational efforts fail, for reasons Chapter Twenty examines. This pillar asks leaders to examine their own relationship with change before they try to change anyone else: how they respond to uncertainty, how they carry their people through it, and how they keep safety steady while everything around it moves. In the language of Chapter Sixteen, it is the leader's own capacity for containment.

The seventh pillar is My Implementation Skills: the move from knowing to doing, from understanding safety leadership to practising it consistently. Its

measure is the leader of whom people say, they walked the talk: the one whose words and decisions, stated values and actions under pressure, line up so reliably that people will follow them through almost anything. That consistency is the Knight's defining quality, and it is where everything in the earlier pillars is finally tested.

The Transformational Safety® Leadership model and the ATLAS framework answer different questions. The model describes what effective safety leadership looks like, and the 360-degree assessment measures where a leader sits across its nine constructs. ATLAS is how a leader moves. Each pillar builds particular constructs, and the link matters, because the construct a leader most needs to develop tells you where the development has to happen.

ATLAS	BUILDS
1 Who am I?	Missionary · Knight
2 What drives me?	Knight · Motivator
3 Emotional intelligence	Carer · self-awareness
4 Leadership strengths	Building from strength
5 Leadership challenges	Thinker · transactional skills
6 Capacity to absorb change	Thinker · protects against the Invisible Man and the Fireman
7 Implementation skills	Knight · walking the talk
Foundation	The base every construct stands on

A weak transactional profile is usually a gap in skills and knowledge, and the work sits mainly in the fourth and fifth pillars: identifying the strengths a leader can build on and the challenges that need development, including the skills that let the Policeman and the Dealer do their jobs well. A weak transformational profile is rarely fixed by skills alone. The Knight's consistency, the Motivator's conviction, the Carer's attention and the Missionary's sense of purpose come from the first three pillars, identity, values

and emotional intelligence, and they are proven in the seventh, where knowledge has to become consistent action. That is why training on its own does not produce Knights or Carers. The Thinker draws on the fifth and sixth pillars: valuing approaches that differ from one's own, and being at ease with change. A passive profile often has a different source. Leaders who are exhausted, or overwhelmed by change, tend to withdraw, appearing as the Invisible Man, or to lead only in crisis, as the Fireman. The sixth pillar addresses this directly: a leader who cannot absorb change cannot contain it for anyone else.

Figure: How the ATLAS framework supports the Transformational Safety® Leadership constructs.

Development in the Work, Not Away From It

Most leadership development still happens away from the work: in a training room, over two or three days, with the leader's phone switched off and their operation running without them. Then the leader goes back, and within weeks most of what was learned has faded. This is not because the content was poor. Research on the transfer of training, going back to the influential work of Timothy Baldwin and Kevin Ford¹¹⁰, has shown for decades that what happens back at work decides whether learning is used. The most powerful factor is not the quality of the course. It is whether the leader's own manager expects, notices and supports the new behaviour, and whether the leader has real opportunities to practise it soon after learning it.

For safety leadership this matters even more than for most skills, because the behaviours the system is trying to build are exactly the ones most easily crowded out by operational pressure. A leader can leave a workshop resolved to spend more time at the worksite, to ask better questions, and to respond differently to bad news. The first production crisis will test all three, usually within a fortnight. If nothing in the leader's working life supports the new behaviour at that moment, the old one returns, and it returns with the added weight of a lesson learned: the new way was fine in theory but did not survive contact with the real job.

The system is therefore built around the work rather than around the classroom. Structured development sessions provide the ideas and the language, but the development itself happens in the leader's own operation, on

real problems, with real people. A leader working on the Thinker does not practise asking how work is actually done in a role-play. They practise it on their next shift walk, with their own crew, and bring what happened back to their coach. A leader working on the Carer does not rehearse a one-to-one conversation. They have one, with a member of their team they have not spoken to properly in months, and reflect afterwards on what they learned and what it cost them to have it. Development that happens in the work is harder to schedule and harder to measure. It is also the only kind that reliably changes what leaders do when the pressure comes.

The Role of Coaching

Coaching is what holds the system together between the moments of structured development. It is not advice-giving, and it is not performance management under another name. It is a regular, confidential conversation in which a leader examines their own safety leadership against the commitments they have made, in the light of what has actually happened since they last spoke.

Good safety leadership coaching has a few consistent features. It works on real situations, not hypothetical ones: the near-miss last Tuesday, the concern that was raised in the toolbox meeting and not followed up, the decision to let a job continue when something did not look right. It looks at the leader's decisions under pressure, because that is where their leadership is formed and where their people are watching. It holds the leader to what they said they would do, without judgement but without letting it slide. It keeps returning to the gap between intention and experience that the construct assessment revealed: not whether the leader meant well, but whether the people they lead would have noticed any difference.

Coaching also does something no other part of the system can. It gives the leader a place to be uncertain. Senior leaders in particular rarely have anyone they can say "I don't know how to handle this" to without some cost. The coaching relationship is one of the few places where they can, and the honesty it allows is the same honesty the system is asking them to create for their own people. A leader who has experienced a relationship in which admitting difficulty is safe is better placed to build one for others. Over time, as their own development matures, leaders begin to offer the same kind of

conversation to the leaders below them, and coaching stops being something done to the organisation and becomes something the organisation does.

The Behavioural Competency Assessment

The 360-degree construct assessment is the diagnostic engine of the system: the instrument that makes the gap between current and required safety leadership capability precise enough to act on. It can also be quite confronting. Leaders assess their own safety leadership across the nine constructs, and a nominated group of the people they lead, who observe their safety leadership in operation, assess it too. Chapter Seventeen showed what the result looks like and what the gap between the two ratings reveals. What matters here is how the instrument is built to reveal it.

The assessment is designed to capture the behavioural reality of safety leadership as the workforce experiences it: not the leader's intentions, stated values or self-image, but the observable behaviours that the people affected by that leadership can describe. I developed its items from long experience of applied safety management and direct observation of safety cultures, and built them to align with the leadership factors Bass and Avolio⁵⁷ identified, but with a specific focus on safety leadership behaviour. They are written to be answered by workers at any level of the organisation. They require no specialist knowledge, only direct observation of how the leader behaves in practice.

That design is what lets the assessment reach information no other mechanism in most organisations can. Within the system, the results are where development begins: they shape the work in the fourth and fifth pillars, and they give coaching its focus. A repeat assessment, usually a year later, then shows whether the changes the leader has made are visible to the people they were made for.

One Leader's Year

The safety director whose chart opened Chapter Seventeen shows how the parts of the system work together. He chose the Carer and the Thinker as his first priorities. But the finding that stayed with him was the Fireman: his people saw him most clearly when something had gone wrong. Whatever else

he changed, he had to become visible before events, not only after them. What happened next is typical of how development through the system unfolds.

He began, uncomfortably, with the first two pillars. Asked who he was as a safety leader, his honest first answer was that he was the person who made sure the system worked. Asked what drove him, he found that much of it was the fear of being the director on whose watch something went badly wrong. Neither answer was shameful. Both explained a great deal about the chart. A leader whose identity rests on the system, and whose motivation is avoiding failure, will build a well-run system, keep a certain distance from the people inside it, and arrive in force when the system fails. That was exactly what his people had reported.

The practical work started small. He committed to one unhurried conversation a week with a frontline worker, on an ordinary day, when nothing had happened. It was not about compliance, but about what their work was actually like and what got in the way of doing it safely. The timing was the point. For years, a visit from the safety director had meant that something had gone wrong. He wanted his presence to mean something else. In the first month, most of those conversations were stilted, because his people were not used to seeing him when nothing was wrong, and some assumed that something must be. By the third month, they had started to tell him things: a maintenance backlog that everyone on the floor knew about and nobody had formally reported, a procedure that could not be followed as written on night shift, a supervisor who made it clear that near-miss reports were unwelcome. Each time, his coach asked the same two questions. What did you do with it? And would the person who told you know that you did?

Not everything went well. Twice in the year he fell back under pressure, once overruling a stop during a critical shutdown and once responding to a reported error with visible irritation. Both times he heard about it, from his coach and from his own people, and both times he went back to the person involved and said so. Those two conversations, he later said, did more for his credibility than anything else he did that year, because they showed his people that the change was real enough to survive his own mistakes.

Twelve months later he took the assessment again. His self-ratings had barely moved. His team's ratings had. The Carer and the Thinker had risen, and the gap between how he saw himself and how they saw him had roughly halved. The Fireman had fallen most of all, because his people now saw him on ordinary days. Near-miss reporting in his area had also risen, which a year

earlier he would have read as a problem. He now read it for what it was: the return path starting to carry information that had been there all along.

How Development Goes Wrong

Safety leadership development fails in recognisable ways, and most of them have nothing to do with the quality of the content.

The first is choosing the comfortable construct. Faced with an assessment, leaders naturally gravitate toward developing what they already do reasonably well, because progress is easier and the feedback is kinder. Chapter Seventeen's advice, to start where the gap is largest, is simple to give and hard to follow, and it takes solid coaching to hold a leader to it.

The second is treating the assessment as a score to be improved rather than information to be understood. Once ratings are linked to performance reviews or bonuses, leaders start managing the number: choosing friendly raters, lobbying their teams before the survey, behaving differently in the weeks before it runs. The same logic that corrupts a Zero Harm target, described in Chapter Fifteen, corrupts a leadership assessment used the same way. The assessment works only when it is used to learn, not to judge.

The third, and the most damaging, is exempting the top. When frontline supervisors are assessed and developed while senior leaders are not, the workforce draws the obvious conclusion: that safety leadership is something asked of people further down. The most powerful signal an executive team can send is to go first: to take the assessment themselves, share what it showed, and be seen to be actively working on it.

Building the Living Culture

Turning individual leadership development into organisational cultural change takes more than adding up the development of individual leaders. It requires deliberately building the organisational conditions that let those behaviours take root, be modelled across levels, and accumulate into the shared assumptions of a transformational safety culture. Development programmes alone cannot produce those conditions. They are produced by aligning the organisation's formal systems, its recognition and accountability structures, communication processes, resource allocation and criteria for

leadership selection and succession, with the values the culture is trying to embed.

That alignment is among the most important and most neglected parts of safety culture development, and Chapter Twenty sets out what it involves in practice. The principle is simple. The culture will develop in the direction of what the organisation's formal accountability says it values, because that is where the consequences that shape behaviour are found.

The living culture, the one in which safety is part of the organisation's identity and present in the operational decisions of everyone in it, is built by three things working together: leaders developing their capability, an organisation aligning its formal systems with the values that capability expresses, and a workforce accumulating the shared experience of what that leadership looks like when it is tested. None of the three is enough on its own. Leaders who develop in an organisation that still rewards the old behaviour will be pulled back toward it. Systems aligned with values that no leader models will be read as another initiative. And a workforce that has not yet seen the new leadership hold under pressure will wait, reasonably, until it does.

CHAPTER TWENTY

Moving the Organisation: From Assessment to Transformation

The report had been read. The leadership team had sat together, predicted their own results, and watched the gap between prediction and reality open up in front of them. They had heard, in the workforce's own words, that incident investigations were experienced as a search for someone to blame, and that change arrived at the worksite without explanation. They had accepted, after some resistance, that their organisation was a capable, well-run compliance culture that had not yet learned to lead safety rather than manage it.

Then someone asked the question that every leadership team asks at this point. So what do we do now?

The answer most organisations reach for is a programme. A name, a logo, a launch event, a training schedule, a set of milestones and a project manager to track them. It feels like action, it can be reported to the board, and it can be finished. It is also, in most cases, the reason nothing changes. Chapter Thirteen explained why: a programme is a way of managing change, and culture is not changed by being managed. Chapter Twelve described what happens when the language of transformation is adopted without the practice. This chapter is about the alternative: how an organisation actually moves, from one level of the maturity model to the next, once it knows honestly where it stands.

Why the Programme Fails

It is worth being precise about why the programme approach fails so consistently, because the reasons point to what works instead.

Michael Hammer¹¹¹, whose work launched the re-engineering movement, was unusually candid that most re-engineering efforts did not achieve what they set out to. The reasons he identified apply with even more force to safety culture. Programmes are designed around the parts of an organisation that can be scheduled: systems, training and communication. They underestimate the parts that cannot: the beliefs and habits people have formed from years of experience, and the time it takes for new experience to change them. They are led by the people who designed them, rather than by the line leaders whose daily behaviour the culture actually responds to. And they end. The end date is the most damaging feature of all, because it tells everyone in the organisation, before the programme has even begun, that this too will pass.

Workers have usually seen several programmes come and go. They know the pattern: the launch, the posters, the training, the first few months of visible attention, and then the slow return to how things were once the attention has moved on. A workforce that has seen that pattern three times will, quite rationally, wait out the fourth. The single most important thing an organisation can do at the start of its journey is to behave in a way that breaks that pattern, and the way to break it is to change what leaders do, visibly, before asking anyone else to change anything.

A Team That Owns It

The first practical step is to decide who owns the work. In many organisations, responsibility for safety culture is spread between several functions that barely speak to each other: the safety department, human resources, operations, and sometimes a separate wellbeing or training team. Each has a piece of the picture, and none has the authority or the mandate to bring the pieces together.

A small safety leadership team, led directly by the chief executive or a senior executive with real authority, changes that. Its members are drawn from the functions that shape the culture, but the team itself sits outside the ordinary reporting lines, so that it can think about the whole organisation rather than defend any one part of it. It meets regularly, every fortnight or so in the early

stages. It becomes the place in the organisation where safety thinking happens: where the assessment findings are owned, where priorities are set, and where progress is honestly reviewed. And it has the authority to bring in anyone, from any level, who can add something to its work, which in practice often means frontline workers whose insight the assessment revealed.

The purpose of the team is not to run a programme. It is to lead the organisation's response, and to hold the organisation, including its own members, to account for doing what it said it would.

Choosing Where to Begin

No organisation can change everywhere at once, and the attempt to do so usually means changing nowhere very much. A better approach is to begin in one part of the organisation, learn what works there, and let that success become the evidence that persuades the rest.

Choosing that first area is a judgement, and it needs to balance two things that pull in opposite directions. The first is need. The assessment will have shown which parts of the organisation are most at risk: where the red zones cluster, where the return path is quietest, where the leadership profile is weakest. There is a strong argument for starting where the need is greatest. The second is the likelihood of success. A pilot that fails, or that is overwhelmed by conditions no leadership development could fix in its first year, will be used by sceptics across the organisation as proof that the whole idea does not work.

The best first area usually sits between the two. It has a real and visible need, so that improvement will mean something. It has a leader who is willing, even eager, to go first, because no amount of external support will compensate for a reluctant leader at the centre of the pilot. It has enough stability, in its people and its operations, for change to take hold. And it is visible enough to the rest of the organisation that what happens there will be noticed.

Three Phases

Organisations that move successfully tend to pass through three broad phases. They overlap, and they are not tidy, but naming them helps leaders understand where they are and what the current phase asks of them.

The first is the foundation phase. It begins with the assessment and its honest reporting back to the workforce, and it continues with the work of building a shared understanding of what the organisation is trying to do and why. Leaders at every level are introduced to the ideas in this book: the difference between managing and leading safety, the leadership constructs, the maturity model, and the principles of high reliability. Leaders take their own individual assessments, beginning at the top. And, crucially, the organisation does something visible about what its people told it.

The second is the capability phase. This is where most of the effort goes, and where most of the time is spent. Leaders develop the specific constructs their assessments showed they most need, through the combination of structured development, practice in real work and coaching described in Chapter Nineteen. At the same time, the organisation builds specific practices that put the high reliability principles to work: near-miss reviews that ask what made a decision reasonable rather than who was at fault, investigations that use the substitution test, shift walks designed to learn rather than inspect, and management of change processes that reach the people the changes affect.

The third is the integration phase. Here the new ways of working stop feeling like initiatives and start feeling like how things are done. The organisation's formal systems are aligned with the values its leaders are now showing: what it recognises, how it selects and promotes leaders, what it measures. Leaders who developed in the earlier phases begin coaching the leaders below them. And the culture starts to carry itself, needing less deliberate effort to sustain because it has become part of what people expect.

No phase has a fixed duration. The foundation phase may take months; the capability phase usually takes years; and integration is never complete, because a culture always needs tending. What matters is that leaders know which phase they are in and do not mistake the end of one for the end of the journey.

Earning Belief Early

The first few months matter more than their length would suggest, because they are when the workforce decides whether this time is different. And the workforce will decide on the basis of what happens, not what is said.

The most powerful early action is to respond, visibly and specifically, to what the assessment revealed. If people said that investigations felt like a

search for someone to blame, the way investigations are done should change, and people should be told how. If people said that change arrived without explanation, the next significant change should be handled differently, and it should be obvious that it is being handled differently. If people raised specific problems in the focus groups, a broken piece of equipment, a procedure that cannot be followed on night shift, a backlog everyone knows about, some of those problems should be fixed quickly and the people who raised them should hear that they have been.

These actions are small compared with the work ahead, and that is the point. They are evidence, delivered early, that the organisation heard what it was told. Chapter Fifteen explained that credibility is built slowly and lost quickly. The early months are when the first deposits are made. The one thing an organisation must not do in this period is promise more than it can deliver. A modest commitment kept is worth far more than an ambitious one that quietly lapses, because the workforce will remember the lapse long after it has forgotten the ambition.

Leaders in the Field

One of the simplest and most effective things senior leaders can do is spend real time at the worksite: not inspecting, not presenting, but talking to the people who do the work. One approach I have recommended to many organisations is for each senior executive to set aside one day a quarter, put on the boots and the high-visibility vest, and spend a full day in an operational area, talking with whoever they meet about their work and what gets in the way of doing it safely.

It works because it does several things at once. It shows the workforce that senior leaders are willing to give the one resource everyone knows they are short of, which is time. It gives those leaders a picture of the organisation that no report can provide. It puts something of the Carer and the Motivator into leaders who are often experienced by the workforce as distant or unreachable. And it gives people at the worksite a chance to say directly what they might never put in writing.

The way it is done decides whether it helps or harms. The executive should arrive to listen, not to check. There should be no clipboard, no checklist and no entourage. Whatever the executive is told should be followed up, and the person who raised it should hear what happened. The visits should

continue, quarter after quarter, long after the novelty has worn off, because a single visit followed by silence teaches the workforce exactly the wrong lesson.

Building Learning Into the Organisation

Most organisations' safety training is dominated by what the law requires: inductions, permits, confined space entry, first aid, the competencies a regulator will check. All of that is necessary. Very little of it builds the culture this book has described, because it is designed to produce compliance, and it is experienced by the workforce as something done to them.

A useful alternative is to bring the organisation's safety learning together under a single structure, sometimes called a safety college, with a clear identity and a curriculum shaped by three things: what regulation requires, what the assessment showed the organisation needs, and what the workforce itself asks for. The important design choice is the split between required and discretionary learning. In any given year, a defined share of the safety learning on offer should be discretionary: not compliance training, but development in leadership, in understanding risk, in the human factors described in Chapter Thirteen, and in the skills of speaking up and listening.

The structure matters as much as the content. Learning offered under a recognised name, with a proper curriculum and some form of recognised achievement, is valued differently from a training calendar. It tells people that the organisation regards safety learning as professional development, not as a box to tick before starting work. And it gives the leadership development described in Chapter Nineteen a home in the organisation, rather than leaving it as a series of events.

Recognition That Builds the Right Things

Every organisation recognises something, whether it means to or not. The question is whether what it recognises builds the culture it wants.

Recognition schemes built around outcomes, injury-free months, days without a lost-time injury, safety targets met, have the same problem as the Zero Harm targets described in Chapter Fifteen: they reward the absence of reported harm, and so they quietly reward not reporting it. Recognition that

builds a learning culture is aimed at behaviour instead: the worker who stopped a job when something did not look right, the crew that reported a near-miss in enough detail for others to learn from it, the supervisor who changed a procedure because their team said it could not be followed as written. As Chapter Twelve argued, the stop that turned out to be unnecessary deserves recognition as much as the one that turned out to be justified.

Recognition also has to be in the hands of the people closest to the work. A scheme run centrally, with nominations reviewed by a committee and awards presented once a year, misses the moment when recognition matters most, which is straight after the behaviour it is meant to reinforce. Supervisors need the means, the authority and the skill to recognise good safety behaviour on the spot, and they need to be developed in using it well. Done badly, recognition becomes the Dealer's exchange described in Chapter Ten: a transaction that buys compliance and teaches people to ask what they will get. Done well, it is one of the clearest signals leaders can send about what the organisation values.

Aligning the Systems

As the work moves into its later phases, the organisation's formal systems have to catch up with the culture its leaders are building. Systems that still reflect the old assumptions will pull people back toward them, however strong the leadership.

The investigation process is usually the first to need attention, because it is where the old assumptions are most visible. Investigations should be redesigned around the questions described in Chapter Sixteen: why the decision made sense at the time, and whether someone else in the same position would have done the same. Findings should lead back to conditions, not stop at a person.

The management of change process needs to reach the people the change affects, before the change arrives, with an explanation of why and a real chance to say what the change will mean in practice.

What the organisation measures needs to shift from outcomes alone to indicators of the culture's health: near-miss reporting and what happens to the reports, the closure of corrective actions against conditions rather than paperwork, the frequency and quality of leaders' time at the worksite, and the movement of the assessment results over time.

And the criteria for selecting, developing and promoting leaders need to include their safety leadership, as their people experience it, not only their technical and managerial competence. Chapter Eight described how the path to promotion reproduces the management default. Changing that path is one of the slowest changes an organisation makes, and one of the most lasting.

The Edges of the Organisation

Much of the most dangerous work in high-hazard industries is done by people who are not employees of the organisation that owns the risk. Contractors, labour-hire workers and specialist service companies often make up a large share of the workforce on any given day, and a far larger share during shutdowns and major projects, exactly when risk is highest. An organisation that develops its own leaders and leaves its contractors untouched has changed the culture of only part of the site.

The edges of the organisation need deliberate attention. Contractors should be included in the assessment, so that the organisation knows how its culture is experienced by the people who work inside it without belonging to it. The supervisors who lead contractor crews, whether employed by the organisation or by the contractor, are frontline leaders in every sense that matters, and they need access to the same development. And the commercial relationship itself sends signals. A contract that rewards speed above everything, penalises delays without asking why they happened, and treats a stopped job as a cost to be recovered will teach contractor crews exactly what the organisation values, whatever its safety policy says.

The same is true of the boundaries inside the organisation. Hazards do not respect the lines between departments, shifts or functions, and Chapter Ten described how a demarcated culture leaves holes at those boundaries. Some of the most useful work in the capability phase happens at the handovers: between day and night shift, between operations and maintenance, between the permit issuer and the crew doing the job. These are the places where information is most easily lost, where responsibility is most easily assumed to belong to someone else, and where the latent conditions of Chapter Sixteen most often hide.

Surviving the Leader Who Leaves

The greatest threat to an organisation's progress is rarely resistance. It is turnover at the top. A new chief executive, a new site manager, a new head of operations, arriving with different priorities and no personal investment in the work already done, can undo years of progress in months. It does not usually happen through a decision to stop. It happens through quiet neglect: the safety leadership team meets less often, the field days lapse, the investigation process drifts back toward blame, and the workforce, watching closely as always, concludes that the commitment belonged to the person who left rather than to the organisation.

Protecting the work against that risk has to start well before anyone leaves. The safety leadership team should be built so that it does not depend on a single person. The organisation's commitments, what it will do, how it will measure progress, how often it will reassess, should be written down and owned by the board as well as the executive. Safety leadership should be part of the criteria for appointing senior leaders, so that the people who arrive are chosen partly for their ability to continue the work. New leaders should be introduced to the assessment findings and the journey so far as part of their induction, not left to discover them. And the regular reassessment described in Chapter Eighteen gives the organisation an early warning: when a change of leader is followed by a slide in the results, the slide can be seen and addressed before it becomes the new normal.

Measuring the Journey

Leaders need to know whether the organisation is moving, and they need to know it in ways that do not repeat the mistakes this book has described. The incident rate is the measure most boards will ask about first, and it is the least reliable guide in the early years of the journey, for a reason that often alarms leaders who have not been warned about it.

As psychological safety improves, reporting rises. People begin to report the near-misses they would previously have kept to themselves, and sometimes the minor injuries too. For a period, the formal data can look worse even as the culture is getting better. An organisation that does not understand this will panic at precisely the moment its work is succeeding, and may abandon it. Leaders, and boards, need to be told in advance that rising reports early in the

journey are usually a sign that the return path is opening, not that the organisation is becoming less safe.

The better measures are the ones described in Chapter Eighteen: the direction of the culture and leadership results over successive assessments, the shrinking of red zones, the narrowing of the gap between leaders' views and the workforce's experience, and the spread of the stronger sub-cultures from the pilot area into the rest of the organisation. Alongside these sit practical indicators of whether the new practices are actually happening: how often leaders are at the worksite, how investigations are being conducted, how quickly reported concerns are closed and whether the people who raised them hear back.

The Long Middle

Every organisation that undertakes this work reaches a point, usually somewhere in the capability phase, where the early energy has gone and the end is nowhere in sight. The launch is a distant memory. The quick wins have been won. The leaders who were most enthusiastic have moved on to other things, or other organisations. Progress is real but slow, and some of it has gone backwards. This is the long middle, and it is where most organisations stop.

Getting through it is less about new initiatives than about steadiness. Leaders need to keep doing the things that matter, the field days, the follow-up, the honest investigations, the coaching, long after they have stopped being novel. The safety leadership team needs to keep meeting, keep reviewing progress honestly, and keep resisting the temptation to declare success or to replace the work with something newer. The leaders themselves need the capacity the sixth ATLAS pillar describes: the ability to absorb the frustrations, setbacks and slow progress of real change without either giving up or pushing harder in ways that undermine the culture they are trying to build.

It helps to remember what the organisation is actually trying to change. Not its systems, which can be changed quickly. Not its training, which can be redesigned in a few months. It is trying to change what thousands of people believe about what happens when they speak up, what their leaders really value, and what this organisation does when safety and something else compete. Those beliefs were formed over years, by experience. They will change the same way.

Beyond Compliance

The organisation described at the start of Chapter Eighteen had all the ingredients for that change: committed leaders, sound systems, and a workforce that was willing to tell the truth when asked properly. What it needed was the discipline to keep acting on what it had heard, long enough for its people to stop waiting for the programme to end, and to realise that this time, there was no programme. There was simply a different way of leading.

PLEASE RESPECT COPYRIGHT

CHAPTER TWENTY-ONE

Beyond Compliance: A New Philosophy of Safety

Return to the street. It is still early morning. The magpies are still working the nature strip. The condensation is still beading on the rooftops. The two men are still making their way along the footpath, almost at the front door, in the silence that no preparation fills.

Consider, now, a different question. Not what happened before that morning: the normalised deviations, the silent management system, the cultural drift, the production pressure, the Wednesday afternoon decision in the third week of a shutdown. That story has been told. Consider instead what would have had to be different for that morning never to arrive. Not different in the sense of the particular procedure that was not followed, the particular pressure check that was skipped, or the particular supervisor who chose not to say what he knew. Different at the level where it matters: the level of culture, leadership, and the assumptions from which every operational decision in that organisation was made.

The organisation would have had to be different. Its leaders would have had to understand safety not as a system to be administered but as a value to be lived. The gap between how the work was designed and how it was done would have had to be visible, not because the procedures were better written, but because the culture was built to surface it. The near-misses, the workarounds and the concerns no one had formally raised would have had to travel back down to the people who could act on them, because the people

who noticed them felt safe to speak. The pressure of a late shutdown would have had to be taken in and held by the culture around the defences, rather than passed down until it reached a man deciding whether forty minutes mattered. The operations manager's presence on site that week would have had to carry a different signal: not the ambient pressure of a schedule that was behind, but the presence of a leader whose consistency between stated values and actual decisions had established, over time, that safety and schedule were not in competition but were parts of the same obligation. And the technician would have had to be working in a culture that did not ask his judgement to bridge a gap leadership had created and left for the workforce to navigate alone.

That organisation was possible. It was not inevitable. It was a leadership decision. The decision was not made, not because no one cared, but because no one had understood, clearly enough and personally enough, what the decision required.

This book has been an attempt to make that requirement clear.

What the Book Has Argued

It is worth stating the book's conclusion plainly before drawing out what it means. Safety performance follows safety culture. Safety culture follows safety leadership. And safety leadership is a decision.

The argument began with the illusion of safety: the gap between what organisations believe about their own safety and the actual state of the conditions that determine it. It moved through the architecture of safety culture, the three layers of artefact, espoused value and basic assumption that explain how culture works and why it fails, and through the mechanisms of failure, anxiety management, identity defence and cultural closure, that explain why organisations with sincere safety commitments can still produce catastrophic outcomes. It set out psychological safety as the condition beneath the whole culture, and High Reliability Organising as a picture of what the organisations that fail least actually do. It developed the leadership framework, the Full Range Leadership Model®, the nine Transformational Safety® Leadership constructs and the ATLAS framework, to describe what safety leadership looks like in practice. And it brought these together in the Integrated Model and the maturity model: a causal account of how harm

travels, how the organisation learns what is travelling, and how the quality of leadership decides whether the defences stay still.

Throughout, a single thread has run: the distinction between compliance and safety, and the claim that the distance between them is not a technical gap that better systems can close but a leadership gap that only the behaviour of leaders can address. Compliance produces organisations that look safe. Leadership produces organisations that are safe. The difference between looking safe and being safe is the one that matters: invisible in the data, invisible in the audit, invisible in the corrective action register, and entirely apparent on the morning when the conditions that have been developing, undetected, for months or years finally produce the walk that two men make in the grey early light.

Safety as a Belief System

A belief system is not a set of propositions a person or an organisation endorses intellectually. It is the deepest layer of the framework through which experience is interpreted and decisions are made. Beliefs of this kind are not chosen consciously. They are formed through experience, through the accumulated evidence of what is true, what works and what the world is like, and they operate below deliberate thought in the moment of decision. They are the basic assumptions of Schein's model: things taken for granted as features of reality rather than recognised as beliefs that could be otherwise.

For safety to work as a belief system in an organisation, rather than as a compliance requirement, a stated priority or a valued aspiration, the organisation's basic assumptions about human life, human dignity and the responsibilities of leadership must rest on one foundational conviction: that every person who comes to work has the right to go home. Not as an aspiration. Not as a performance metric. As the assumption that governs every operational decision in which safety is at stake.

An organisation that holds this conviction at the level of basic assumption does not make the same decisions as one that holds safety as a priority. It does not make the same production decisions, the same decisions about deferring maintenance, or the same resourcing decisions. It does not send the same signals about what is valued and what can be traded. And it does not create the same environment for the frontline workers who make the choices no safety system can fully anticipate. Held at the level of assumption rather than

aspiration, the conviction that every person has the right to go home produces the conditions from which safe operations grow.

You cannot communicate your way to the level of basic assumption. Belief at that depth is built by leaders who act as though they hold it, and keep acting that way long enough for people to stop wondering whether they mean it. Not when it is easy. Not when it costs nothing. When it is hard. When it costs something. When the pressure is real, the schedule is tight, and the commercial consequence of holding to the conviction is visible and felt.

Those are the moments that build safety belief. They are the moments workers watch, and from which they take the most accurate account available of what the organisation values. And they are the moments from which safety leadership is made: not in the safety briefing, the annual safety day or the published commitment to Zero Harm, but in the daily, operational, pressured decisions most leaders make without ever thinking of them as safety culture decisions at all.

Safety as a Leadership Practice

A practice is a way of engaging with the world that has become habitual through deliberate cultivation: behaviours developed with intention and repeated with consistency until they no longer demand effort and are simply the way things are done. Building one takes what the ATLAS framework builds on: an identity that makes the practice meaningful, values that make it sincere, and sustained, self-aware application over time until deliberate behaviour becomes habit.

Safety as a leadership practice is not the set of things safety leaders do in their formal role: the safety walk, the toolbox meeting, the corrective action review, the monthly safety discussion with the executive team. Those are the artefacts of safety leadership. They are necessary, and they are not enough. Safety as a leadership practice is the quality of engagement leaders bring to every contact with the operational reality of their organisation: the curiosity, the honesty, the concern for the people they lead, and the willingness to act on what that engagement reveals, even when acting on it is inconvenient.

The practice shows in the questions leaders ask, and the ones they do not. A leader who visits the worksite and asks whether procedures are being followed is practising safety management. A leader who visits and asks what the crew is most worried about that is not in any procedure is practising safety

leadership. The difference is not sophistication or experience. It is orientation: the difference between assessing compliance and trying to understand the safety reality of the work, including the parts the formal system has not yet captured.

The practice also shows in how leaders respond when safety and something else compete. An organisation that has been working on its safety culture for several years knows, with some precision, what each of its leaders actually does when the schedule pressure is real, the safety concern is serious, and the decision cannot be avoided. The workers know. The leadership team knows. The culture that grows from those decisions is the safety culture the organisation is building, whatever its policy statements say. Where the two disagree, the practice is the policy, every time.

Developing safety as a leadership practice takes the same investment as any serious professional practice: structured development of the foundations, deliberate practice in real work, honest reflection on what the practice is producing, and coaching that provides both challenge and accountability over the years it takes. It is not achieved in a workshop, a module or a certification. It is achieved through a thousand decisions across a career, in which the leader chooses, consistently and with growing ease, to act from the values a safety culture needs.

Safety as a Cultural Identity

The final dimension of this philosophy is the most ambitious and the hardest to achieve: safety as a cultural identity. Not as a cultural value; organisations at the Functional level of the maturity model already hold real safety values. Not as a cultural priority; Chapter Three showed why that is not enough. Safety as cultural identity is the state in which the organisation's understanding of itself, its answer to the question of what kind of organisation this is and what it stands for, includes safety not as a separate commitment but as an inseparable part of its character.

Identity of this kind is recognisable where it exists. It shows in the way workers describe their organisation to outsiders: not in the language of safety metrics or certifications, but in the language of what this place values and how it treats its people. It shows in the way new people are brought into the culture: less through safety training and procedure familiarisation than through immersion in a culture whose relationship to safety is part of its character. And

it shows in the way the organisation responds to the events that test it: the near-miss investigated with the depth that learning needs, the concern addressed with the urgency that commitment demands, and the decision under pressure that reveals, more precisely than anything else, what the organisation actually is.

Building this identity is the work of years. It is built from the bottom of the maturity model upward, through the patient, sustained, culturally intelligent leadership that produces the shared experiences from which cultural assumptions are formed and reformed. It requires leaders who understand that they are not managing a safety programme. They are building an organisation, defining through every decision what it is and what it is for. The safety culture that results is not a product of the safety management system they maintain. It is the accumulated expression of the leadership they have provided, the culture they have shaped, and the identity they have built into the assumptions of everyone who works there.

The Call to Leaders

This book is addressed to every leader who has read it, not only to safety professionals, safety managers and directors, or the specialists to whom safety has traditionally been handed. It is addressed to everyone with authority over the safety of others: the production manager, the operations director, the general manager, the board member, the team leader. Safety is not a specialist domain whose management can be delegated to the safety function while the rest of the leadership attends to other things. Safety culture is a leadership decision, and it is made, actively or by default, by every leader whose choices shape the conditions in which other people work.

Leaders who read this book and recognise their own organisation in it, in the virtual safety management system of Chapter Two, the priority-versus-value distinction of Chapter Three, the cultural misalignment of Chapter Five, or the passive-avoidant or transactional profile of Chapter Ten, are not being indicted. They are being offered the most useful thing a serious book about safety leadership can offer: an honest account of where the current model falls short, and of the distance between where their culture is and where it needs to be.

The response to that recognition is a decision. The leader who recognises a Developing culture in their organisation and responds by investing more

heavily in the safety management system is deciding to keep doing what produced the culture they have. The leader who recognises it and begins the serious, sustained, personally demanding work of developing safety leadership, in themselves, in their direct reports and in the conditions of their organisation, is making a different decision. Its consequences will not show in the next safety performance review, the next audit or the next board report. They will show, gradually and then unmistakably, in the safety culture their organisation becomes.

That decision is not mainly an organisational commitment. It is a personal one. It asks each leader to sit with a hard question: is the safety culture I have built good enough to protect the people who work within it? Not good enough against the industry average, not good enough for the regulator, not good enough to produce acceptable metrics. Good enough in the most demanding and personal sense: that if the conditions now developing in the gap between Work-as-Imagined and Work-as-Done were to line up, as they eventually do, with everything else a serious event requires, the leader could say honestly that the culture they built had given every worker in it the protection they deserved.

Most leaders who ask that question honestly will find the answer is not yet. The culture is better than it was. The system is more complete than it was. The metrics are moving in the right direction. But the assumptions at the base of the culture, the values enacted under pressure, the quality of safety leadership at the frontline and the engagement of the workforce are not yet what they need to be to answer the question with confidence.

That recognition is not a failure. It is where the work begins. And the leaders who begin it, who choose the sustained, difficult, personally demanding work of safety leadership over the more comfortable path of continued safety management improvement, are making the decision that defines what they stand for.

The New Philosophy

A philosophy of safety equal to the challenge of transforming safety culture is not a more sophisticated version of the compliance model. It is a different kind of thing: a different way of understanding what safety is, what it requires, and what leaders owe the people who depend on their decisions.

In this philosophy, safety is not compliance. Compliance is the floor: the minimum below which no responsible organisation operates, and the

administrative foundation on which safety culture is built. The floor is not the destination. An organisation that achieves compliance and stops has achieved the minimum required by law, not the minimum required to protect the people it employs. Safety culture transformation begins where compliance ends.

In this philosophy, safety is not a priority. Priorities are comparative, contextual and adjustable under pressure. Safety held at the level of cultural identity and personal leadership commitment is not one priority among others. It is the lens through which operational decisions are made, the framework within which production choices are weighed, and the standard against which every risk control is finally judged. An organisation for which safety is a value does not trade it against production when the pressure is real. It finds a way to produce safely, or it accepts, honestly, that current conditions do not allow production on safe terms.

In this philosophy, safety is not a programme. Programmes have beginnings and ends. Safety culture transformation has a beginning and a direction. It does not end, because a culture is never finished: it is always being built or eroded, depending on the quality of the leadership shaping it. Leaders who understand this do not manage their organisation's transformation as a project with a completion date. They lead it as a continuing commitment, measured not by the final state of the culture but by the direction it is moving and the consistency of the leadership moving it.

In this philosophy, safety is not a system. Systems are necessary. They are not sufficient, and they are not the point. The point is the person who, on a Wednesday afternoon in the third week of a shutdown, makes a decision about a pressure check, and this time makes it differently. Not because a procedure has been revised, a corrective action closed or a safety message delivered. Because the culture in which he is deciding has been built, by years of safety leadership, into one in which the safer decision is not the harder one, or the slower one, or the one that takes courage to defend against the schedule. It is simply the obvious one: the one consistent with who he is, with what the organisation is, and with what every leader from his supervisor to the general manager has shown, through years of consistent and credible behaviour, that this organisation stands for.

* * *

The sensor light came on above the door, flat and white. The school shoes were still there on the mat, side by side, the way a child leaves them when

he has been told to. From inside came the kettle, the television, a cupboard door: a morning that had not yet heard.

Footsteps. Light, quick. Not hers.

The door opened. The boy, toast in hand, looked up at the two men with his father's colouring and an expression that had not yet learned to hide anything.

He turned back toward the hallway.

"Mum?"

* * *

The walk those two men made is not inevitable. It is the predictable outcome of organisations whose leaders have not yet understood that the culture they build, the decisions they make under pressure, and the signals they send about what is valued and what can be traded are continuously shaping the conditions in which people make decisions about pressure checks, and permit slips, and whether, this one time, it will probably be fine.

It is the outcome of organisations that believe they are safe because they are compliant, that their systems are working because the metrics are improving, and that their leadership is adequate because the audits are passing. Organisations that have mistaken the administration of safety for its substance, the performance of commitment for the thing itself, and the appearance of a good safety culture for the experience of working in one.

It is also the outcome that safety leadership can prevent. Not through better systems, though the systems need to be better. Not through more training, though the gaps in capability need to be closed. Through a continuous, active, daily decision by leaders at every level to understand the reality of the work being done under their authority, to build the conditions in which that reality can be honestly seen and effectively addressed, and to show, in every decision, in every relationship, and in every moment when the choice between safety and something else presents itself, that the people in their care are worth protecting.

That is the philosophy this book has been building toward. That is the decision it is asking you to make.

The system performs exactly as it is led.

Lead it well.

PLEASE RESPECT COPYRIGHT

Annotated Bibliography

- [1] Reason, J. (1997). **Managing the risks of organizational accidents.** Ashgate.

Relevance to the story: Reason's organisational-accident framework provides the book with its core distinction between active failures and latent conditions and its system-oriented approach to causation. It supports the argument that frontline actions are often the visible end of a longer causal chain created by design, management, organisational and cultural conditions.

- [2] UNC Kenan-Flagler Business School. (2010). **Out of sight, out of mind.** UNC Executive Development.

Relevance to the story: Provides the contemporary management context for the Deepwater Horizon safety-award episode used early in the book. It supports the book's central contrast between reassuring corporate safety signals and the actual condition of operational risk. The example illustrates how positive indicators and recognition can coexist with serious latent vulnerability.

- [3] Hopkins, A. (2011). **Management walk-arounds: Lessons from the Gulf of Mexico oil well blowout.** National Research Centre for OHS Regulation, Australian National University.

Relevance to the story: Hopkins uses the Macondo disaster to argue for management presence that is genuinely connected to operational reality rather than ceremonial or administrative visibility. This supports the book's emphasis on leaders going to the work, asking questions that expose risk, and distinguishing management walk-arounds from superficial demonstrations of safety commitment.

- [4] **U.S. Chemical Safety and Hazard Investigation Board. (2016). Drilling rig explosion and fire at the Macondo well (Report No. 2010-10-I-OS), Volume 3: Human factors, organizational learning, safety performance indicators, risk management practices, corporate governance, safety culture.**

Relevance to the story: The CSB's Macondo analysis documents human factors, organisational learning failures, weak safety indicators, risk-management shortcomings and governance problems. It supports the book's argument that catastrophic events emerge from interacting organisational conditions rather than a single frontline error, and that conventional indicators can fail to reveal deteriorating process risk.

- [5] **U.S. Chemical Safety and Hazard Investigation Board. (2007). Investigation report: Refinery explosion and fire, BP Texas City, Texas, March 23, 2005 (Report No. 2005-04-I-TX).**

Relevance to the story: The official Texas City investigation provides primary evidence of systemic process-safety failures behind the 2005 refinery explosion. It supports the book's use of Texas City as an example of how poor hazard control, weak organisational learning and misleading reliance on personal-safety performance can coexist before a major accident.

- [6] **Republic of the Marshall Islands Maritime Administrator. (2011). Deepwater Horizon marine casualty investigation report (Official No. 2213; IMO No. 8764597).**

Relevance to the story: This official marine casualty investigation documents the operational decisions, warnings and failures surrounding Deepwater Horizon. It supports the book's argument that weak signals can be present yet misinterpreted or discounted, and that frontline concerns are ineffective when organisational structures and leadership do not reliably hear, interpret and act on them.

- [7] **Baker, J. A., III, et al. (2007). The report of the BP U.S. refineries independent safety review panel. BP U.S. Refineries Independent Safety Review Panel.**

Relevance to the story: The Baker Panel is central to the book's distinction between personal safety metrics and process safety. Its findings support the argument that an organisation may report acceptable injury statistics while serious process hazards deteriorate underneath. It also strengthens the book's case that leadership and governance failures can sit well above the frontline.

- [8] **National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling. (2011). Deep water: The Gulf oil disaster and the future of offshore drilling. Report to the President. U.S. Government Printing Office.**

Relevance to the story: The National Commission's report places the Deepwater Horizon disaster in a wider system of technical, managerial, regulatory and organisational failures. It supports the book's rejection of simple operator-blame explanations and its argument that major accidents often reflect accumulated decisions and governance weaknesses across multiple organisational levels.

- [9] **Hollnagel, E. (2012). FRAM: The functional resonance analysis method: Modelling complex socio-technical systems. Ashgate.**

Relevance to the story: Hollnagel's work provides the conceptual foundation for the distinction between Work-as-Imagined and Work-as-Done used throughout the book. It supports the argument that adaptation is normal in complex work, and that safety leadership must understand the gap between formal systems and the variable reality of operations rather than treating deviation automatically as misconduct.

- [10] **Vaughan, D. (1996). The Challenger launch decision: Risky technology, culture, and deviance at NASA. University of Chicago Press.**

Relevance to the story: Vaughan's analysis of Challenger established the concept of normalisation of deviance as a social and organisational process. It supports the book's argument that repeated successful departures from formal expectations can gradually become accepted as normal, allowing risk to migrate without any single dramatic decision to abandon safety.

- [11] **Health and Safety Executive. (2006). Developing process safety indicators: A step-by-step guide for chemical and major hazard industries (HSG254). HSE Books.**

Relevance to the story: HSE's process-safety indicator guidance supports the book's distinction between lagging measures of past events and leading measures of the condition of safety-critical controls. It reinforces the argument that injury rates and completed activities are insufficient proxies for catastrophic risk, and that leaders need information about whether critical barriers are actually healthy.

- [12] **Boskeljon-Horst, L., de Boer, R. J., & Dekker, S. W. A. (2024). Reducing gaps between paper and practice requires more than a technical alignment. Safety Science, 170, 106291.**

Relevance to the story: This paper directly examines the persistent gap between formal procedures and actual practice. It supports the book's argument that closing the paper-practice gap requires more than rewriting procedures or technically aligning systems; cultural conditions, leadership, psychological safety and the way rules are enacted all shape whether formal controls become real work practices.

- [13] **Kamarova, S., Gagné, M., Holtrop, D., & Dunlop, P. D. (2025). Integrating behavior and organizational change literatures to uncover crucial psychological mechanisms underlying the adoption and maintenance of organizational change. *Journal of Organizational Behavior*, 46(2), 263–287.**

Relevance to the story: Kamarova and colleagues integrate behavioural and organisational-change research to explain why new practices are adopted, internalised and maintained - or why organisations revert. This supports the book’s discussion of organisational homeostasis and the need for sustained leadership, meaning, mastery and belonging if safety change is to survive beyond an intervention period.

- [14] **Schein, E. H. (2010). *Organizational culture and leadership* (4th ed.). Jossey-Bass.**

Relevance to the story: Schein’s three-level model of artefacts, espoused values and basic underlying assumptions provides the principal cultural architecture used in the book. It supports the argument that visible safety systems and stated values can be misleading when deeper assumptions produce different behaviour, and that effective diagnosis must reach beneath artefacts to the assumptions that organise everyday action.

- [15] **Blanc, F., Ottimofiore, G., & Myers, K. (2022). From OSH regulation to safety results: Using behavioral insights and a “supply chain” approach to improve outcomes — The experience of the Health and Safety Executive. *Safety Science*, 145, 105491.**

Relevance to the story: Blanc and colleagues show how regulatory requirements are translated through a chain of actors before they become real safety outcomes. The paper supports the book’s claim that compliance on paper does not automatically create safe practice, and that effective regulation depends on behaviour, implementation, communication and the conditions under which regulated organisations actually act.

- [16] **Kruglanski, A. W. (2004). The psychology of closed mindedness. Psychology Press.**

Relevance to the story: Kruglanski's work on need for closure explains why people and groups can prefer certainty, stable interpretation and rapid closure over continuing ambiguity. It supports the book's account of cultural closure: organisations under pressure can seize on reassuring explanations and then resist or filter information that would require reopening settled beliefs about safety.

- [17] **International Nuclear Safety Advisory Group. (1986). Summary report on the post-accident review meeting on the Chernobyl accident (Safety Series No. 75-INSAG-1). International Atomic Energy Agency.**

Relevance to the story: The INSAG report is an authoritative post-accident account of Chernobyl and the organisational conditions surrounding it. It supports the book's historical discussion of safety culture as a concept developed to explain why technical systems alone cannot account for catastrophic performance, and why organisational norms and management practices matter.

- [18] **Hashemian, S. M., & Triantis, K. (2023). Production pressure and its relationship to safety: A systematic review and future directions. Safety Science, 159, 106045.**

Relevance to the story: This systematic review synthesises evidence on production pressure and its relationship with safety outcomes. It supports the book's argument that schedule, cost and output pressures reshape attention, workload and decision-making, often encouraging shortcuts and competing priorities even when no leader explicitly instructs workers to compromise safety.

- [19] **Butler, T., & Brooks, R. (2024). Time for a paradigm change: Problems with the financial industry's approach to operational risk. *Risk Analysis*, 44(6), 1285–1304.**

Relevance to the story: Butler and Brooks critique formal operational-risk approaches that can become detached from how organisations actually understand and act on risk. Their analysis supports the book's distinction between documented risk systems and enacted culture, particularly the argument that assumptions, values and behaviour can undermine technically sophisticated controls.

- [20] **Petitta, L., & Martínez-Córcoles, M. (2023). A conceptual model of mindful organizing for effective safety and crisis management: The role of organizational culture. *Current Psychology*, 42, 25773–25792.**

Relevance to the story: Petitta and Martínez-Córcoles link organisational culture with mindful organising and reliable performance. Their model supports the book's argument that high reliability depends not only on techniques but on cultural conditions that sustain attentiveness, anticipation, learning and effective response to weak signals and unexpected events.

- [21] **Cantu, J., Gharehyakheh, A., Fritts, S., & Tolk, J. (2021). Assessing the HRO: Tools and techniques to determine the high-reliability state of an organization. *Safety Science*, 134, 105082.**

Relevance to the story: Cantu and colleagues review approaches for assessing whether organisations have actually achieved a high-reliability state. The paper supports the book's caution against 'HRO theatre': adopting the vocabulary of high reliability is not the same as demonstrating the practices, capabilities and organisational conditions associated with reliably managing hazardous work.

- [22] Hashemian, S. M., & Triantis, K. (2023). Production pressure and its relationship to safety: A systematic review and future directions. *Safety Science*, 159, 106045.

Relevance to the story: This systematic review provides the evidence base for the book's recurring treatment of production pressure as a system condition rather than a personal weakness. It shows how output and schedule demands interact with workload, attention, shortcuts, error and safety climate, supporting the claim that workers inherit pressures created by organisational priorities.

- [23] Flin, R., Mearns, K., O'Connor, P., & Bryden, R. (2000). Measuring safety climate: Identifying the common features. *Safety Science*, 34(1-3), 177-192.

Relevance to the story: Flin and colleagues identify recurring dimensions across safety-climate measures in high-hazard industries, including management commitment, communication, competence and procedures. This supports the book's claim that workers infer the real importance of safety from consistent organisational and leadership signals rather than from policy statements alone.

- [24] Cullen, W. D. (1990). The public inquiry into the Piper Alpha disaster. HMSO.

Relevance to the story: Cullen's inquiry into Piper Alpha is a foundational primary source on the systemic and managerial conditions preceding the disaster. It supports the book's argument that major accidents are rarely isolated technical events; they emerge from weaknesses in communication, control of work, organisational arrangements and management of competing operational demands.

- [25] **U.S. Chemical Safety and Hazard Investigation Board. (2007). Investigation report: Refinery explosion and fire, BP Texas City, Texas, March 23, 2005 (Report No. 2005-04-I-TX).**

Relevance to the story: The Texas City report provides detailed evidence of long-standing process-safety weaknesses, prior warnings and organisational failures before the explosion. It supports the book's themes of normalised risk, misleading performance indicators, and the danger of organisations believing they are safe because administrative systems and personal-safety measures appear satisfactory.

- [26] **Brielmaier, C., & Friesl, M. (2023). The attention-based view: Review and conceptual extension towards situated attention. International Journal of Management Reviews, 25(1), 99–129.**

Relevance to the story: The attention-based view explains how organisational behaviour is shaped by what leaders and systems direct attention toward. It supports the book's argument that leaders reveal priorities through the scarce resource of attention and time, and that what is repeatedly noticed, discussed and resourced becomes culturally significant while unattended risks can disappear from view.

- [27] **Deming, W. E. (1986). Out of the crisis. Massachusetts Institute of Technology, Center for Advanced Engineering Study.**

Relevance to the story: Deming's management philosophy distinguishes leadership from mere supervision and control. It supports the book's argument that safety cannot be created only through inspection, compliance checking and correction; leaders must improve the system, remove obstacles, understand variation and create conditions in which people can perform reliably.

- [28] **Argyris, C., & Schön, D. A. (1974). Theory in practice: Increasing professional effectiveness. Jossey-Bass.**

Relevance to the story: Argyris and Schön's distinction between espoused theory and theory-in-use directly supports the book's analysis of the gap between stated safety values and enacted behaviour. Their work also underpins the argument for double-loop learning: organisations must question the assumptions and governing beliefs behind practice, not merely correct deviations within an unchanged model.

- [29] **Silverman, J., & Aafjes-van Doorn, K. (2023). Coping and defense mechanisms: A scoping review. Clinical Psychology: Science and Practice, 30(4), 381–392.**

Relevance to the story: This scoping review clarifies the relationship between coping processes and defence mechanisms, including more adaptive and less adaptive responses to threat. It supports the book's explanation of how individuals and groups can protect themselves from anxiety-producing information through defensive processes that shape interpretation rather than through deliberate dishonesty.

- [30] **Harmon-Jones, E., & Harmon-Jones, C. (2023). Dissonance motivation from an action-based perspective: An updated review. Advances in Motivation Science, 10, 1–36.**

Relevance to the story: Harmon-Jones and Harmon-Jones review contemporary cognitive-dissonance theory from an action-based perspective. The work supports the book's claim that inconsistency between beliefs, identity and threatening evidence creates motivational pressure to reduce dissonance, which can contribute to rationalisation, reinterpretation or selective acceptance of safety information.

- [31] Ecker, U. K. H., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L. K., Brashier, N., Kendeou, P., Vraga, E. K., & Amazeen, M. A. (2022). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1, 13–29.

Relevance to the story: Ecker and colleagues review why false or preferred beliefs can persist even after corrective information is presented. This supports the book's discussion of organisational truth-avoidance and the difficulty of changing established safety narratives once they are integrated into identity, memory and decision-making routines.

- [32] Pearson, D., Watson, P., Albertella, L., & Le Pelley, M. E. (2022). Attentional economics links value-modulated attentional capture and decision-making. *Nature Reviews Psychology*, 1, 320–333.

Relevance to the story: Pearson and colleagues show that attention is not neutral: stimuli associated with value can capture attention and influence later decisions. This supports the book's argument that organisational priorities and reward structures shape what people notice, what they overlook and which signals receive cognitive weight under competing demands.

- [33] Fasolo, B., Heard, C., & Scopelliti, I. (2025). Mitigating cognitive bias to improve organizational decisions: An integrative review, framework, and research agenda. *Journal of Management*, 51(6), 2182–2211.

Relevance to the story: This integrative review examines cognitive bias in organisational decisions and approaches to debiasing. It supports the book's argument that leadership judgement is systematically vulnerable to confirmation, framing and other biases, and that safer decision environments must be deliberately designed rather than relying on individual awareness alone.

- [34] **Roberts, L. M., & Caza, B. B. (2025). Positive identity construction in diverse organizations. *Annual Review of Organizational Psychology and Organizational Behavior*, 12, 423–451.**

Relevance to the story: Roberts and Caza review how people construct and protect positive identities in organisations. Their work supports the book's argument that safety information can become identity-threatening when it challenges a group's self-concept as competent or safe, creating defensive responses that influence whether disconfirming evidence is accepted or reinterpreted.

- [35] **Harmon-Jones, E., & Harmon-Jones, C. (2023). Dissonance motivation from an action-based perspective: An updated review. *Advances in Motivation Science*, 10, 1–36.**

Relevance to the story: This updated review of cognitive dissonance supports the book's treatment of cultural failure as partly a motivational process. When evidence conflicts with established beliefs or identity, people are motivated to restore coherence; in organisations, that dynamic can help explain rationalisation and resistance to information that threatens a preferred safety narrative.

- [36] **Edmondson, A. C., & Bransby, D. P. (2023). Psychological safety comes of age: Observed themes in an established literature. *Annual Review of Organizational Psychology and Organizational Behavior*, 10, 55–78.**

Relevance to the story: Edmondson and Bransby synthesise the mature literature on psychological safety, showing its importance for speaking up, learning and interpersonal risk-taking. The review supports the book's claim that organisations cannot reliably surface weak signals, mistakes and concerns unless people believe they can contribute difficult information without unacceptable social or career consequences.

- [37] Sedlar, N., Irwin, A., Martin, D., & Roberts, R. (2023). A qualitative systematic review on the application of the normalization of deviance phenomenon within high-risk industries. *Journal of Safety Research*, 84, 290–305.

Relevance to the story: This systematic review identifies how normalisation of deviance develops across high-risk industries, including production pressure, organisational culture, risk normalisation and the absence of negative consequences. It supports the book's account of cultural drift in which repeated successful deviations gradually become accepted as ordinary practice.

- [38] Boskeljon-Horst, L., de Boer, R. J., & Dekker, S. W. A. (2024). Reducing gaps between paper and practice requires more than a technical alignment. *Safety Science*, 170, 106291.

Relevance to the story: The paper supports the book's argument that the gap between formal safety systems and actual work is not solved by procedure revision alone. Its emphasis on organisational and cultural alignment reinforces the need for leadership, psychological safety and just treatment when bringing Work-as-Imagined closer to Work-as-Done.

- [39] Schein, E. H., & Bennis, W. G. (1965). *Personal and organizational change through group methods*. Wiley.

Relevance to the story: Schein and Bennis are an early source for the concept of psychological safety in organisational change. Their work supports the book's historical account that psychological safety predates modern team research and was originally concerned with reducing the interpersonal threat involved in learning, unlearning and changing established behaviour.

- [40] Kahn, W. A. (1990). Psychological conditions of personal engagement and disengagement at work. *Academy of Management Journal*, 33(4), 692–724.

Relevance to the story: Kahn's engagement model identifies psychological safety as one of the conditions affecting whether people bring themselves fully into their work roles. It supports the book's argument that interpersonal risk, role conditions and the quality of the social environment influence whether workers engage, contribute and speak rather than withdraw or protect themselves.

- [41] Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383.

Relevance to the story: Edmondson's seminal team study operationalised psychological safety and linked it to learning behaviour. It provides the core empirical foundation for the book's treatment of psychological safety as a team-level condition that makes speaking up, admitting error, asking questions and learning from operational experience more likely.

- [42] Brimhall, K. C., Tsai, C.-Y., Eckardt, R., Dionne, S., Yang, B., & Sharp, A. (2023). The effects of leadership for self-worth, inclusion, trust, and psychological safety on medical error reporting. *Health Care Management Review*, 48(2), 120–129.

Relevance to the story: Brimhall and colleagues examine leadership, inclusion, trust, psychological safety and medical-error reporting. The study supports the book's argument that reporting is shaped by social conditions and leadership, not simply by formal reporting systems, and that trust and psychological safety influence how teams surface and learn from safety-relevant information.

- [43] Sert, Ö., & Mert, İ. S. (2026). Silent weak signals in port operations: Ethical leadership, psychological safety and reporting climate in safety-critical terminals. *Journal of Contingencies and Crisis Management*, 34(3), e70207.

Relevance to the story: This study of safety-critical port operations links ethical leadership, psychological safety and reporting climate with employee silence and weak-signal reporting. It supports the book's argument that leaders influence whether early warning information reaches decision-makers and that silence can persist even in technically sophisticated safety systems when interpersonal conditions are unsafe.

- [44] Dodoo, J. E., Surienty, L., & Al-Samarraie, H. (2023). The influence of learning-oriented leadership for promoting future-directed workplace safety in the mining industry. *Safety Science*, 159, 106010.

Relevance to the story: Dodoo and colleagues examine learning-oriented leadership in mining and its relationship with future-directed safety behaviours such as safety voice and stewardship. The study supports the book's argument that leadership can move safety beyond rule compliance by creating conditions in which workers actively contribute to learning and future risk prevention.

- [45] Noort, M. C., Reader, T. W., & Gillespie, A. (2021). Safety voice and safety listening during aviation accidents: Cockpit voice recordings reveal that speaking-up to power is not enough. *Safety Science*, 139, 105260.

Relevance to the story: Using cockpit voice recordings from aviation accidents, Noort and colleagues show that speaking up is not sufficient if those in authority do not listen effectively. The study supports the book's argument that psychological safety includes both willingness to challenge and leadership receptiveness to challenge, especially across steep authority gradients.

- [46] Perkins, K., Ghosh, S., Vera, J., Aragon, C., & Hyland, A. (2022). The persistence of safety silence: How flight deck microcultures influence the efficacy of crew resource management. *International Journal of Aviation, Aeronautics, and Aerospace*, 9(3), Article 6.

Relevance to the story: This aviation study examines how flight-deck microcultures can sustain safety silence despite formal Crew Resource Management training. It supports the book's argument that training cannot overcome a local culture that discourages challenge, and that enacted norms between crew members determine whether formal communication principles are actually used.

- [47] Transportation Safety Board of Canada. (2019). Air transportation safety investigation report A19Q0010. Transportation Safety Board of Canada.

Relevance to the story: The Transportation Safety Board report provides an applied accident-investigation example of communication, authority and cockpit decision-making under operational pressure. It supports the book's use of aviation to illustrate how warnings can be weakened or neutralised when challenge is not effectively heard and incorporated into decision-making.

- [48] Edmondson, A. C. (2018). *The fearless organization: Creating psychological safety in the workplace for learning, innovation, and growth*. Wiley.

Relevance to the story: Edmondson translates psychological-safety research into practical leadership principles for learning organisations. The book supports Beyond Compliance's distinction between psychological safety and comfort, and its argument that leaders must deliberately frame work, invite contribution and respond productively to bad news if learning and accountability are to coexist.

- [49] **National Research Council. (1982). Toward safer underground coal mines. National Academy Press.**

Relevance to the story: This National Research Council report examines the organisational, technical and regulatory conditions necessary for safer underground coal mining. It supports the book's longstanding argument that safety programmes, rules and recommendations have little value unless they are translated into effective action and operating conditions at the workplace.

- [50] **Ronald, L. A. (1998). Identifying the elements of successful safety programs: A literature review. Prevention Division, Workers' Compensation Board of British Columbia.**

Relevance to the story: Ronald's literature review synthesises characteristics of successful occupational-safety programmes and emphasises the importance of active management commitment and leadership. It supports the book's argument that the effectiveness of safety systems depends on how leaders enact and reinforce them, not simply on whether programme elements exist on paper.

- [51] **Dyregborg, J., Lipscomb, H. J., Nielsen, K., Törner, M., Rasmussen, K., Frydendall, K. B., Bay, H., Gensby, U., Bengtsen, E., Guldenmund, F., & Kines, P. (2022). Safety interventions for the prevention of accidents at work: A systematic review. Campbell Systematic Reviews, 18(2), e1234.**

Relevance to the story: This Campbell systematic review compares a broad range of workplace safety interventions across a very large evidence base. It supports the book's argument that interventions aimed at organisational, group and engineering conditions generally offer stronger leverage than relying primarily on individual behaviour change, training or administrative exhortation.

- [52] **Burke, M. J., Sarpy, S. A., Smith-Crowe, K., Chan-Serafin, S., Salvador, R. O., & Islam, G. (2006). Relative effectiveness of worker safety and health training methods. *American Journal of Public Health*, 96(2), 315–324.**

Relevance to the story: Burke and colleagues' meta-analysis shows that more engaging and participatory safety training methods produce stronger learning and safety outcomes than passive approaches. It supports the book's argument that training quality and method matter, while also reinforcing that training is only one component of a wider system influencing behaviour.

- [53] **Evans, A., & Parker, J. (2008). Beyond safety management systems. *AeroSafety World*, 3(5), 13–17.**

Relevance to the story: Evans and Parker argue that safety management systems are necessary but insufficient without leadership, describing leadership as the 'secret ingredient' beyond formal systems. The article directly supports the book's central distinction between managing safety administratively and leading the culture and behaviours through which those systems become effective.

- [54] **Bennis, W. G. (1999). Managing people is like herding cats. Executive Excellence Publishing.**

Relevance to the story: Bennis's work supports the book's distinction between management and leadership and the proposition that leadership can be developed rather than treated as an innate trait. It underpins the book's emphasis on building leaders who create direction, meaning and commitment rather than relying only on managerial control.

- [55] **Collins, J. (2001). Good to great: Why some companies make the leap... and others don't. HarperBusiness.**

Relevance to the story: Collins' research programme emphasises the disproportionate role of leadership in organisational performance while cautioning against simplistic charisma-based explanations. It supports the book's argument that sustained organisational change depends on leadership discipline, humility and consistency, not merely on programmes or systems introduced from above.

- [56] Li, P., Sun, J.-M., Taris, T. W., Xing, L., & Peeters, M. C. W. (2021). Country differences in the relationship between leadership and employee engagement: A meta-analysis. *The Leadership Quarterly*, 32(1), 101458.

Relevance to the story: This meta-analysis examines how leadership relates to employee engagement across national contexts. It supports the book's claim that leadership is a major antecedent of engagement while also showing that context matters; leadership effects are shaped by cultural and national conditions rather than operating identically everywhere.

- [57] Bass, B. M., & Avolio, B. J. (Eds.). (1994). *Improving organizational effectiveness through transformational leadership*. Sage.

Relevance to the story: Bass and Avolio provide the foundational Full Range Leadership Model used in the book's leadership chapters. The framework distinguishes passive-avoidant, transactional and transformational behaviours and supports the book's argument that leadership is a behavioural range with different consequences for motivation, performance and safety culture.

- [58] Barling, J., Loughlin, C., & Kelloway, E. K. (2002). Development and test of a model linking safety-specific transformational leadership and occupational safety. *Journal of Applied Psychology*, 87(3), 488–496.

Relevance to the story: Barling, Loughlin and Kelloway provide direct empirical evidence that safety-specific transformational leadership is associated with better occupational-safety outcomes through safety climate and safety consciousness. The study supports the book's argument that general leadership principles become especially powerful when leaders make safety an explicit part of how they influence people.

- [59] Zohar, D. (2000). A group-level model of safety climate: Testing the effect of group climate on microaccidents in manufacturing jobs. *Journal of Applied Psychology*, 85(4), 587–596.

Relevance to the story: Zohar's group-level safety-climate model shows how supervisory practices create local safety climates that predict safety outcomes. It supports the book's emphasis on the frontline supervisor as a powerful cultural signal: workers learn the real priority of safety from day-to-day supervisory behaviour rather than corporate statements alone.

- [60] Flin, R., & Yule, S. (2004). Leadership for safety: Industrial experience. *Quality and Safety in Health Care*, 13(Suppl. 2), ii45–ii51.

Relevance to the story: Flin and Yule review leadership lessons from high-risk industries and their relevance to healthcare. The paper supports the book's argument that safety leadership involves visible commitment, communication, involvement and the active shaping of conditions for safe work, rather than leaving safety solely to specialist functions.

- [61] Lyubych, Z., Turner, N., Herscovis, M. S., & Deng, C. (2022). A meta-analysis of leadership and workplace safety: Examining relative importance, contextual contingencies, and methodological moderators. *Journal of Applied Psychology*, 107(12), 2149–2175.

Relevance to the story: This large meta-analysis establishes a broad evidence base linking leadership with workplace safety and compares different leadership orientations. It supports the book's claim that leadership effects on safety are robust but context-dependent, and that task, relationship and safety-specific leadership behaviours each contribute in different ways.

- [62] Judge, T. A., & Piccolo, R. F. (2004). Transformational and transactional leadership: A meta-analytic test of their relative validity. *Journal of Applied Psychology*, 89(5), 755–768.

Relevance to the story: Judge and Piccolo’s meta-analysis compares transformational, contingent-reward, management-by-exception and laissez-faire leadership. It supports the book’s Full Range Leadership discussion by showing that active leadership forms are generally associated with better outcomes, while passive and avoidant leadership is consistently weaker.

- [63] Kapp, E. A. (2012). The influence of supervisor leadership practices and perceived group safety climate on employee safety performance. *Safety Science*, 50(4), 1119–1124.

Relevance to the story: Kapp links supervisor leadership practices and perceived group safety climate with employee safety performance. The study supports the book’s argument that supervisors translate organisational expectations into local behaviour and that the climate workers experience is partly constructed through visible day-to-day leadership practices.

- [64] Brown, W., Birnstihl, E. A., & Wheeler, D. W. (1996). Leading without authority: An examination of the impact of transformational leadership Cooperative Extension work groups and teams. *Journal of Extension*, 34(5), Article 4.

Relevance to the story: Brown, Birnstihl and Wheeler provide an exploratory application of transformational leadership in work groups, showing very strong associations between the transformational dimensions and perceived extra effort and effectiveness. The study supports the book’s explanation of how inspirational, developmental and intellectually stimulating leadership can mobilise discretionary effort.

- [65] **Clarke, S. (2013). Safety leadership: A meta-analytic review of transformational and transactional leadership styles as antecedents of safety behaviours. *Journal of Occupational and Organizational Psychology*, 86(1), 22–49.**

Relevance to the story: Clarke's meta-analysis directly examines transformational and transactional leadership as antecedents of safety behaviour. It supports the book's balanced argument that transformational leadership strengthens safety climate and participation while active transactional leadership also contributes, particularly to compliance, and that effective safety leadership can draw on both.

- [66] **Zhao, L., Yang, D., Liu, S., & Nkrumah, E. N. K. (2022). The effect of safety leadership on safety participation of employee: A meta-analysis. *Frontiers in Psychology*, 13, 827694.**

Relevance to the story: This meta-analysis shows a positive relationship between safety leadership and employee safety participation. It supports the book's emphasis on discretionary safety behaviour - speaking up, helping, improving and contributing beyond minimum compliance - and shows that transformational safety leadership is particularly relevant to participation.

- [67] **Towers Perrin. (2006). Winning strategies for a global workforce: Attracting, retaining and engaging employees for competitive advantage (Towers Perrin Global Workforce Study, Executive Report). Towers Perrin.**

Relevance to the story: The Towers Perrin global workforce study provides the engagement data used in the book to illustrate the limited proportion of employees who are highly engaged. It supports the argument that discretionary effort cannot be assumed and that leadership, meaningful work, recognition and organisational connection influence whether people contribute beyond formal requirements.

- [68] **Quansah, P. E., Zhu, Y., & Guo, M. (2023). Assessing the effects of safety leadership, employee engagement, and psychological safety on safety performance. *Journal of Safety Research*, 86, 226–244.**

Relevance to the story: Quansah and colleagues link safety leadership with safety compliance and participation and examine employee engagement and psychological safety as mechanisms. The study supports the book's integrated argument that leadership shapes safety partly by influencing whether workers are engaged enough, and feel safe enough, to contribute actively.

- [69] **Wachter, J. K., & Yorio, P. L. (2014). A system of safety management practices and worker engagement for reducing and preventing accidents: An empirical and theoretical investigation. *Accident Analysis & Prevention*, 68, 117–130.**

Relevance to the story: Wachter and Yorio examine safety-management practices together with worker engagement and accident outcomes. Their findings support the book's proposition that systems work best when workers are actively engaged in them, and that engagement can mediate the relationship between formal safety practices and actual safety performance.

- [70] **Chênevert, D., Hill, K., & Kilroy, S. (2022). Employees' perceptions of non-monetary recognition practices and turnover: Does recognition source alignment and contrast matter? *Human Resource Management Journal*, 32(1), 40-57.**

Relevance to the story: Chênevert, Hill and Kilroy examine how non-monetary recognition from organisational and supervisory sources relates to turnover. The study supports the book's wider engagement argument that recognition is a social signal of value, and that alignment between organisational messages and the behaviour of direct supervisors affects whether employees remain connected and committed.

- [71] **Griffin, M. A., & Neal, A. (2000). Perceptions of safety at work: A framework for linking safety climate to safety performance, knowledge, and motivation. *Journal of Occupational Health Psychology*, 5(3), 347–358.**

Relevance to the story: Griffin and Neal provide the influential distinction between safety compliance and safety participation and link safety climate with safety knowledge and motivation. This framework is central to the book's argument that doing the required minimum is different from discretionary contribution, and that climate influences both capability and motivation for safe behaviour.

- [72] **O'Dea, A., & Flin, R. (2000). Safety leadership in the offshore oil and gas industry. In *SPE International Conference on Health, Safety and Environment in Oil and Gas Exploration and Production*. Society of Petroleum Engineers.**

Relevance to the story: O'Dea and Flin examine safety leadership in offshore oil and gas, linking management and supervisory leadership with workforce safety behaviour. The study supports the book's argument that leaders at different organisational levels influence different parts of the safety system, including initiative, participation and everyday compliance.

- [73] **Goleman, D., Boyatzis, R., & McKee, A. (2002). The emotional reality of teams. *Journal of Organizational Excellence*, 21(2), 55–65.**

Relevance to the story: Goleman, Boyatzis and McKee examine the emotional dynamics of teams and the role leaders play in shaping collective emotional climate. The article supports the book's argument that effective safety leadership requires emotional awareness and relationship management because leader affect and interpersonal behaviour influence openness, trust and engagement.

- [74] Cherniss, C., & Goleman, D. (Eds.). (2001). **The emotionally intelligent workplace: How to select for, measure, and improve emotional intelligence in individuals, groups, and organizations.** Jossey-Bass.

Relevance to the story: This edited volume provides the broader evidence base for emotional intelligence in workplace performance and leadership. It supports the book's treatment of self-awareness, self-management, social awareness and relationship management as foundational capabilities for leaders who must build trust, influence climate and respond effectively under pressure.

- [75] Rochlin, G. I., La Porte, T. R., & Roberts, K. H. (1987). **The self-designing high-reliability organization: Aircraft carrier flight operations at sea.** *Naval War College Review*, 40(4), 76–92.

Relevance to the story: Rochlin, La Porte and Roberts' study of aircraft-carrier flight operations is a foundational account of high-reliability organising. It supports the book's description of organisations that achieve sustained reliability in hazardous, tightly coupled environments through disciplined coordination, expertise, learning and continuous sensitivity to operations.

- [76] Weick, K. E., & Sutcliffe, K. M. (2015). **Managing the unexpected: Sustained performance in a complex world (3rd ed.).** Wiley.

Relevance to the story: Weick and Sutcliffe synthesise the core processes of high-reliability organising: preoccupation with failure, reluctance to simplify, sensitivity to operations, commitment to resilience and deference to expertise. Their framework provides the principal conceptual structure for the book's HRO chapter and its leadership implications.

- [77] Edmondson, A. C. (2003). Speaking up in the operating room: How team leaders promote learning in interdisciplinary action teams. *Journal of Management Studies*, 40(6), 1419–1452.

Relevance to the story: Edmondson's operating-room research shows how team leaders can promote speaking up and learning in interdisciplinary action teams. It supports the book's argument that leadership behaviour creates the interpersonal conditions in which expertise can cross status boundaries and operational concerns can be voiced before they become harm.

- [78] Chassin, M. R., & Loeb, J. M. (2013). High-reliability health care: Getting there from here. *The Milbank Quarterly*, 91(3), 459–490.

Relevance to the story: Chassin and Loeb apply high-reliability concepts to healthcare and describe the organisational changes required to move beyond isolated safety initiatives. Their work supports the book's argument that HRO is not a label or programme but a sustained organisational capability requiring leadership, culture, process improvement and disciplined learning.

- [79] Perrow, C. (1984). Normal accidents: Living with high-risk technologies. Basic Books.

Relevance to the story: Perrow's Normal Accidents theory provides the principal counterpoint to optimistic interpretations of high reliability. It supports the book's acknowledgement that interactive complexity and tight coupling can make some accidents inherently difficult to prevent, strengthening the case for humility, resilience and scepticism about claims of complete control.

- [80] Hopkins, A. (2009). The problem of defining high reliability organisations. In A. Hopkins (Ed.), *Learning from high reliability organisations*. CCH Australia.

Relevance to the story: Hopkins critiques the difficulty of defining High Reliability Organisations and the risk of treating HRO as a vague aspirational label. This supports the book's 'HRO theatre' argument and its insistence that high reliability should be judged by concrete organisational practices and capabilities rather than terminology or self-description.

- [81] **Williamson, A. M., & Feyer, A.-M. (2000). Moderate sleep deprivation produces impairments in cognitive and motor performance equivalent to legally prescribed levels of alcohol intoxication. *Occupational and Environmental Medicine*, 57(10), 649–655.**

Relevance to the story: Williamson and Feyer provide experimental evidence that extended wakefulness can impair cognitive and motor performance to levels comparable with legally significant alcohol intoxication. The study supports the book's argument that fatigue is an operational risk condition affecting judgement, attention and performance, not merely an individual wellness issue.

- [82] **Rogers, T. (2023). Human information processing under stress. *Journal of Emergency Management*, 21(2), 137–146.**

Relevance to the story: Rogers reviews how stress affects human information processing, including attention, perception and decision-making. It supports the book's argument that people under pressure do not process the same information in the same way, making workload, time pressure and stress important contextual influences on safety-critical judgement.

- [83] **Gay, C., Mahony, M., & Graves, J. (2005). Best practices in employee communication: A study of global challenges and approaches. IABC Research Foundation.**

Relevance to the story: The IABC study examines communication challenges across hundreds of organisations and identifies factors associated with employee alignment. It supports the book's argument that communication is not simply message transmission: employees need to understand strategy, trust leadership and see how their own work connects to organisational goals.

- [84] **Boswell, W. R., Bingham, J. B., & Colvin, A. J. S. (2006). Aligning employees through “line of sight.” *Business Horizons*, 49(6), 499–509.**

Relevance to the story: Boswell, Bingham and Colvin develop the concept of employee ‘line of sight’ - understanding organisational objectives and how one’s own work contributes to them. It supports the book’s argument that direct managers are critical translators of strategy and safety expectations into meaningful, actionable local priorities.

- [85] **Peker, M., Doğru, O. C., & Meşe, G. (2022). Role of supervisor behavioral integrity for safety in the relationship between top-management safety climate, safety motivation, and safety performance. *Safety and Health at Work*, 13(2), 192–200.**

Relevance to the story: Peker and colleagues show that supervisor behavioural integrity - consistency between safety words and actions - strengthens the pathway from top-management safety climate through motivation to safety performance. This directly supports the book’s argument that workers believe what leaders repeatedly do more than what organisations formally say.

- [86] **Slovic, P. (1993). Perceived risk, trust, and democracy. *Risk Analysis*, 13(6), 675–682.**

Relevance to the story: Slovic’s classic paper explains the asymmetry of trust: negative events that destroy trust tend to carry more weight than positive events that build it. This supports the book’s discussion of safety credibility, where a single visible contradiction between stated values and leadership behaviour can undermine a much larger volume of safety communication.

- [87] **Maister, D. H., Green, C. H., & Galford, R. M. (2000). *The trusted advisor*. Free Press.**

Relevance to the story: The Trusted Advisor introduces the trust equation linking credibility, reliability, intimacy and self-orientation. The framework supports the book’s practical discussion of leadership trust by showing that technical credibility alone is insufficient; consistent behaviour, relational safety and low self-orientation shape whether people believe and confide in leaders.

- [88] **Karanikas, N., & Zerguine, H. (2025). Contextualising new safety paradigms: A study in a large Australian construction company. *Safety Science*, 185, 106784.**

Relevance to the story: Karanikas and Zerguine examine how newer safety paradigms are understood and implemented inside a large Australian construction company. The study supports the book's argument that safety ideas must be translated into local organisational context and that participation, leadership and implementation conditions determine whether new paradigms become lived practice.

- [89] **Sherratt, F., Harch, D., & Perez, A. (2024). Making zero work for construction safety in a post-zero world. *Journal of Safety Research*, 91, 193–200.**

Relevance to the story: Sherratt, Harch and Perez examine the evidence and limitations surrounding 'zero' approaches in construction safety. The paper supports the book's critique of Zero Harm as a credibility problem when aspirational language is disconnected from evidence, operational reality or the way workers experience risk.

- [90] **Heinrich, H. W. (1931). *Industrial accident prevention: A scientific approach*. McGraw-Hill.**

Relevance to the story: Heinrich's early accident-prevention text is historically important because it shaped decades of thinking about unsafe acts, accident ratios and prevention. The book uses it as a reference point for explaining how influential safety models can become institutionalised even when their assumptions and evidentiary foundations later require critical reassessment.

- [91] **Manuele, F. A. (2011). Reviewing Heinrich: Dislodging two myths from the practice of safety. *Professional Safety*, 56(10), 52–61.**

Relevance to the story: Manuele critically re-examines two enduring ideas attributed to Heinrich, including the dominance of unsafe acts and simple frequency-severity relationships. The article supports the book's challenge to person-centred and triangle-based safety thinking and its argument that historical models should not be treated as unquestioned empirical laws.

- [92] Li, Z., Yao, M., Luo, Z., Wang, X., Liu, T., Huang, Q., & Su, C. (2024). A chemical accident cause text mining method based on improved accident triangle. *BMC Public Health*, 24, 39.

Relevance to the story: Li and colleagues apply text mining and an updated accident-triangle approach to chemical accident causes. The paper supports the book's discussion of contemporary attempts to refine older accident-ratio thinking, while also illustrating that modern analysis increasingly needs to incorporate organisational and causal information rather than simple numerical ratios.

- [93] Dekker, S. W. A. (2024). *Preventing industrial accidents: Reappraising H. W. Heinrich—More than triangles and dominoes*. Routledge.

Relevance to the story: Dekker's reappraisal places Heinrich in historical context and revisits what his work did and did not establish. It supports the book's more nuanced critique: Heinrich's legacy cannot be reduced to slogans, yet the data and assumptions behind widely repeated safety triangles and person-focused interpretations warrant careful scrutiny.

- [94] Han, Y., Liu, Y., & Loewenstein, G. (2023). Confusing context with character: Correspondence bias in economic interactions. *Management Science*, 69(2), 1070–1091.

Relevance to the story: Han, Liu and Loewenstein experimentally examine correspondence bias - the tendency to infer character from behaviour while underweighting context. This directly supports the book's argument against reflexive worker blame: behaviour that looks like a personal failing may be a predictable response to incentives, constraints and conditions created by the system.

- [95] **Longford Royal Commission. (1999). The Esso Longford gas plant accident: Report of the Longford Royal Commission (Parliamentary Paper No. 61). Government Printer for the State of Victoria.**

Relevance to the story: The Longford Royal Commission is the primary authoritative account of the 1998 Esso gas-plant disaster. It supports the book's argument that attributing failure to operators can obscure organisational causes; the Commission identified deeper deficiencies in training, procedures, knowledge and management of a complex process system.

- [96] **Hopkins, A. (2000). Lessons from Longford: The Esso gas plant explosion. CCH Australia.**

Relevance to the story: Hopkins' analysis of Longford translates the Royal Commission findings into a broader organisational-safety lesson. It supports the book's rejection of frontline blame and shows how management systems, organisational knowledge and leadership conditions can create the circumstances in which operator actions become the final visible step in a systemic failure.

- [97] **Meadows, S., Baker, K., & Butler, J. (2005). The incident decision tree: Guidelines for action following patient safety incidents. In K. Henriksen, J. B. Battles, E. S. Marks, & D. I. Lewin (Eds.), Advances in patient safety: From research to implementation (Vol. 4: Programs, tools, and products). Agency for Healthcare Research and Quality.**

Relevance to the story: The Incident Decision Tree operationalises James Reason's culpability approach and includes the substitution test: asking whether another similarly qualified person might have acted similarly in the same circumstances. It supports the book's practical method for distinguishing individual culpability from behaviour shaped by system conditions.

- [98] Hofmann, D. A., Jacobs, R., & Landy, F. (1995). High reliability process industries: Individual, micro, and macro organizational influences on safety performance. *Journal of Safety Research*, 26(3), 131–149.

Relevance to the story: Hofmann, Jacobs and Landy integrate individual, work-group and organisational influences on safety performance in high-reliability process industries. The paper supports the book's multilevel model by showing that frontline attitudes and behaviour are embedded within, and shaped by, broader management systems and organisational practices.

- [99] Reason, J. (1998). Achieving a safe culture: Theory and practice. *Work & Stress*, 12(3), 293–306.

Relevance to the story: Reason's paper describes safety culture as the 'engine' that drives a system toward sustained resistance to operational hazards. It supports the book's argument that safe performance requires an active cultural force that maintains vigilance and learning even during long periods without accidents, when complacency is otherwise likely to grow.

- [100] Beus, J. M., Dhanani, L. Y., & McCord, M. A. (2015). A meta-analysis of personality and workplace safety: Addressing unanswered questions. *Journal of Applied Psychology*, 100(2), 481–498.

Relevance to the story: Beus, Dhanani and McCord meta-analyse personality and workplace safety, showing that traits such as conscientiousness and agreeableness have real but limited relationships with unsafe behaviour and safety outcomes. The findings support the book's balanced position that individual differences matter, but cannot adequately explain safety performance on their own.

- [101] Priolo, G., Vignoli, M., & Nielsen, K. (2025). Risk perception and safety behaviors in high-risk workers: A systematic literature review. *Safety Science*, 186, 106811.

Relevance to the story: This systematic review examines the relationship between risk perception and safety behaviour in high-risk workers. It supports the book's argument that how accurately workers perceive and interpret risk influences behaviour, while also showing that the relationship is context-dependent and not strong enough to justify person-only explanations.

- [102] Lee, S., & Dalal, R. S. (2016). Climate as situational strength: Safety climate strength as a cross-level moderator of the relationship between conscientiousness and safety behaviour. *European Journal of Work and Organizational Psychology*, 25(1), 120–132.

Relevance to the story: Lee and Dalal provide direct evidence for safety climate as 'situational strength'. They found that conscientiousness predicted safety behaviour less strongly when safety climate was strong and consistent. This directly supports the book's argument that good organisational conditions reduce the amount of safety performance left to individual personality differences.

- [103] Doerr, A. J. (2020). When and how personality predicts workplace safety: Evaluating a moderated mediation model. *Journal of Safety Research*, 75, 275–283.

Relevance to the story: Doerr tests when and how personality predicts workplace safety, including the mediating role of safety motivation and the moderating role of safety climate. The study supports the book's argument that personality effects are conditional: stronger safety environments can attenuate the behavioural influence of individual conscientiousness.

- [104] Simard, M., & Marchand, A. (1995). A multilevel analysis of organisational factors related to the taking of safety initiatives by work groups. *Safety Science*, 21(2), 113–129.

Relevance to the story: Simard and Marchand's multilevel analysis shows that work-group safety initiative is strongly influenced by organisational and supervisory conditions, particularly participative safety management. It supports the book's argument that first-line supervisors are part of the causal environment of safety, shaping whether workers merely comply or actively contribute to improvement.

- [105] O'Dea, A., & Flin, R. (2001). Site managers and safety leadership in the offshore oil and gas industry. *Safety Science*, 37(1), 39–57.

Relevance to the story: O'Dea and Flin examine offshore installation managers and the relationship between leadership, safety attitudes and workforce behaviour. The study supports the book's claim that site-level leadership commitment and style influence how safety is enacted below them, making leadership behaviour a meaningful part of the operational safety system.

- [106] Hudson, P. (2003). Applying the lessons of high risk industries to health care. *Quality and Safety in Health Care*, 12(Suppl. 1), i7–i12.

Relevance to the story: Hudson's evolutionary safety-culture model provides the five-stage ladder from pathological through reactive, calculative and proactive to generative. It supports the book's maturity model by showing that cultures can be distinguished by characteristic ways of thinking about safety, information, responsibility and learning as they develop.

- [107] Westrum, R. (2004). A typology of organisational cultures. *Quality and Safety in Health Care*, 13(Suppl. 2), ii22–ii27.

Relevance to the story: Westrum's typology distinguishes pathological, bureaucratic and generative cultures according to how organisations process information, especially bad news. It supports the book's emphasis on information flow as a diagnostic of culture: whether messengers are punished, tolerated or actively sought reveals how an organisation is likely to respond to risk signals.

- [108] Yammarino, F. J., & Atwater, L. E. (1993). Understanding self-perception accuracy: Implications for human resource management. *Human Resource Management*, 32(2–3), 231–247.

Relevance to the story: Yammarino and Atwater develop the self-other agreement framework for leadership perceptions. Their work supports the book's argument that leaders who substantially overestimate their own effectiveness can be less capable of change because the feedback path is broken; accurate self-perception is therefore a necessary condition for leadership development.

- [109] Fruhen, L. S., Andrei, D. M., & Griffin, M. A. (2022). Leaders as motivators and meaning makers: How perceived leader behaviors and leader safety commitment attributions shape employees' safety behaviors. *Safety Science*, 152, 105775.

Relevance to the story: Fruhen, Andrei and Griffin show that workers respond not only to safety-leadership behaviour but also to the motives they attribute to it. The study supports the book's argument that affective or intrinsic safety commitment produces different responses from calculative commitment, even where outward leadership behaviours may appear similar.

- [110] Baldwin, T. T., & Ford, J. K. (1988). Transfer of training: A review and directions for future research. *Personnel Psychology*, 41(1), 63–105.

Relevance to the story: Baldwin and Ford provide the foundational model of training transfer, showing that lasting application depends on trainee characteristics, training design and the post-training work environment. This supports the book's argument that leadership development fails when the workplace does not provide expectation, reinforcement, opportunity and support for new behaviour.

- [111] **Hammer, M., & Champy, J. (1993). Reengineering the corporation: A manifesto for business revolution. HarperBusiness.**

Relevance to the story: Hammer and Champy's re-engineering work is used as a caution about large-scale organisational change. Their acknowledgement of high failure rates supports the book's argument that programmes often underestimate culture, habits and line leadership, and that change initiatives fail when they are treated as temporary projects rather than sustained changes in everyday work.