FIRST EDITION · 2026

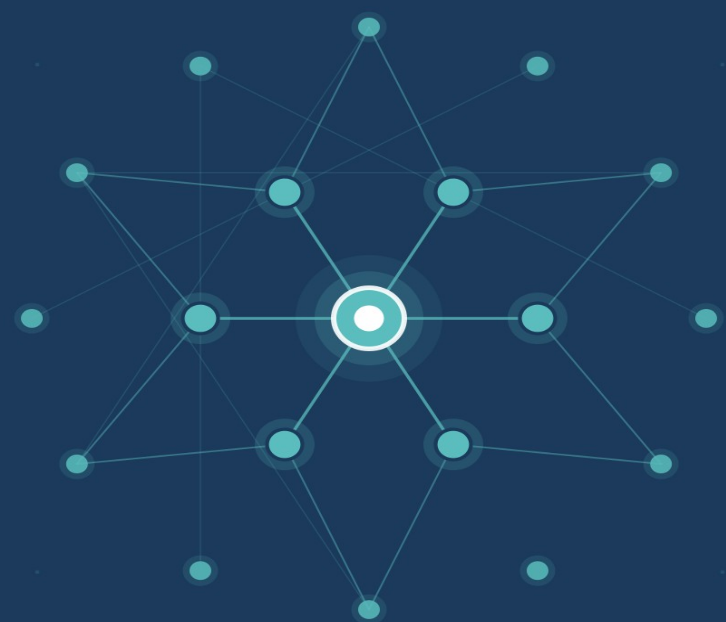

LEADING WITH ANALYTICS AND AI

A Global Boardroom Guide for the Emerging-Economy Era

Thanakit Ouanhlee

List of Case Studies

The book is structured around twenty-four substantive case studies — twelve opening cases that introduce each chapter through globally significant organisations, and twelve regional cases that ground the analytical content in emerging-economy practice, principally across Asia-Pacific. The running case study of Elastic Company, a 1,500-employee Thai textile manufacturer in Samut Sakhon Province, threads the cumulative analytical apparatus across all twelve chapters.

Opening Cases (Global)

Ch	Organisation	Origin	Subject
1	DHL	Germany / Global	The Leader's Role in the Age of Analytics and AI
2	Walmart	United States	The Data Foundation at Operational Scale
3	Capital One	United States	The Indicator Architecture
4	Microsoft	United States / Global	Strategy Translation Under Satya Nadella
5	IBM (Watson Health)	United States / Global	AI Governance Lessons from a High-Profile Setback
6	Zillow Offers	United States	The Cost of Inadequate Risk Management
7	Johns Hopkins University	United States	The COVID-19 Dashboard
8	Hans Rosling and Gapminder	Sweden / Global	Communicating Complex Data Visually
9	Schneider Electric	France / Global	ESG-AI Integration
10	BlackRock and Aladdin	United States / Global	Embedded Analytics at Global Scale
11	JPMorgan Chase	United States	The Sustained AI Narrative
12	NVIDIA	United States / Global	Three Decades of Sustained Capability Renewal

Regional Cases (Asia-Pacific Focus)

Ch	Organisation	Origin	Subject
1	DBS Singapore	Singapore	Digital Transformation in Banking
2	Grab	Singapore / SE Asia	The Regional Super-App's Analytical Foundation
3	Siam Commercial Bank (SCB)	Thailand	Indicator Architecture in Thai Banking
4	CP Group	Thailand / Asia	Strategy Translation in the Diversified Conglomerate
5	MAS Veritas Initiative	Singapore	Regulator-Led AI Governance
6	Petronas	Malaysia	Operational Risk in National Energy
7	Bank of Thailand	Thailand	Dashboards in Central Banking
8	Siam Cement Group (SCG)	Thailand / Asia	Multi-Audience Analytical Communication
9	Indorama Ventures	Thailand / Global	ESG-AI in Petrochemicals
10	Sea Group	Singapore / SE Asia	The Growth-to-Profitability Strategic Pivot
11	Ping An Insurance	China	Founder-Led Sustained AI Narrative
12	Tata Consultancy Services	India / Global	Multi-Decade Continuous Capability Renewal

The Running Case

Elastic Company, a 1,500-employee Thai textile manufacturer in Samut Sakhon Province, is the running case across all twelve chapters. The case begins in Chapter 1 with the strategic decision to deploy an AI-driven predictive turnover capability and develops cumulatively across the subsequent chapters into the integrated analytical apparatus operating across workforce, operations, commercial, and supply chain analytics. Chapter 12 returns to the operational record five years after the original deployment, validating the substantive direction of the original strategic decision.

Ch	Stage	Substantive Content
1	The Strategic Decision	Mode 3 deployment with override authority for the predictive turnover capability
2	The Technical Decisions	Gradient boosting model, AUC 0.78 at deployment, operational infrastructure
3	The Measurement Architecture	Three KPIs and three KRIs; Templates 2 and 3 instantiated
4	Strategy Translation	Premium customer pillar, THB 14M investment supporting THB 80M revenue target
5	The Governance Architecture	NIST AI RMF, Tier 2 classification, AI Risk Officer, AI Ethics Committee
6	Operational Risk	Twelve material risks, defence-in-depth, pre-positioned crisis response
7	The Dashboard Architecture	Three-tier dashboards: strategic, operational, analytical
8	Six-Audience Communication	Board, exec committee, workforce, line managers, regulator, public
9	Integrated Sustainability	ESG-AI integration, workforce pillar, ten-page integrated disclosure
10	Strategic Decisions	Capital allocation, divestiture, acquisition consideration, market entry, capability
11	The Boardroom Pitch	Analytical infrastructure expansion approved by the board
12	Five Years On	The cumulative apparatus and operational record validating direction

List of Frameworks, Templates, and Models

The book develops and applies several frameworks, templates, and models that the contemporary practitioner can use directly. The following list provides quick reference to the principal analytical artifacts the book introduces or applies.

Authoritative Templates

Template	Introduced	Fields
Template 2 — KPI/KRI Card (Ten Fields)	Ch 3	Metric Name, Definition, Unit, Time Window, Baseline, Target, Amber Threshold, Red Threshold, Owner, Action When Breached
Template 3 — Risk Register Entry (Six Fields)	Ch 3	Risk, Likelihood/Impact, Early Warning Indicator, Mitigation, Owner, Escalation Trigger + Action

Major Frameworks Applied

Framework	Chapter	Components
The Three Modes of Analytical Engagement	Ch 1	Mode 1 (informational), Mode 2 (decision-supporting), Mode 3 (autonomous with override)
NIST AI Risk Management Framework	Ch 5	Govern, Map, Measure, Manage functions and four-tier risk classification
Three-Lines-of-Defence Model (Applied to AI)	Ch 5	Line 1 (operational ownership), Line 2 (risk/compliance), Line 3 (internal audit)
AI Risk Taxonomy (Seven Categories)	Ch 6	Model, Data, Operational, Cybersecurity, Regulatory, Reputational, Strategic
Three-Tier Dashboard Architecture	Ch 7	Strategic (board), Operational (executive), Analytical (deep-dive)
Six-Audience Communication Map	Ch 8	Board, executive committee, workforce, line managers, regulator, public
ESG-AI Integration Framework	Ch 9	Materiality assessment, integrated disclosure, workforce pillar, just transition

Five Strategic Decision Categories	Ch 10	Capital allocation, portfolio, M&A, market entry, capability investment
The Five Pitch Structural Elements	Ch 11	Opening, substantive argument, supporting evidence, risk treatment, explicit ask
The AI Capability Renewal Cycle	Ch 12	Development, operational maturity, technology evolution pressure, substantive renewal
The Five Disciplines of the Learning Organisation (Senge, Adapted)	Ch 12	Personal mastery, mental models, shared vision, team learning, systems thinking

Supplementary Models Referenced

Model	Chapter	Substantive Content
Porter Five Forces	Ch 10	Competitive rivalry, threat of new entrants, threat of substitutes, buyer power, supplier power
BCG Growth-Share Matrix	Ch 10	Stars, cash cows, question marks, dogs
GE-McKinsey Nine-Box Matrix	Ch 10	Industry attractiveness × business unit strength
Ansoff Matrix	Ch 10	Existing/new product × existing/new market
Real Options Analysis	Ch 10	Application of options theory to capital decisions including management flexibility
Monte Carlo Simulation	Ch 10	Stochastic methods for parameter uncertainty in strategic decisions
The Mintzberg Critique of Strategic Planning	Ch 10	Planning fallacy, formalisation fallacy, predetermination fallacy
Resource-Based View	Ch 10	Internal resources and capabilities as principal source of competitive advantage
The Innovator's Dilemma (Christensen)	Ch 12	Pattern through which successful capabilities encounter transition difficulty
Double-Loop Learning (Argyris & Schön)	Ch 12	Learning that adjusts frameworks vs single-loop learning within existing frameworks

List of Callout Boxes

The book uses four distinctive callout types operating as visual and structural anchors for the substantive content. The callouts appear consistently across all twelve chapters using a shared colour scheme — navy as the principal accent, teal as the secondary accent — that supports the visual coherence of the cumulative book.

The Four Callout Types

Callout Type	Colour	Function
KEY CONCEPT	Teal	Substantive definition, framework articulation, or principle clarification
EXECUTIVE DECISION	Navy	Decision-oriented synthesis distilling the substantive choices that the analytical content supports
WORKED EXAMPLE	Gold	Detailed worked-through application of a discipline to a specific scenario
CASE STUDY	Grey	Synthesis of lessons from a substantive case discussion

Approximate Distribution Across Chapters

Each chapter typically contains four to six callouts distributed across the four types. The cumulative book contains approximately fifty to sixty callouts across the twelve chapters, with the cumulative distribution being approximately:

Type	Approx. Count	Notes
KEY CONCEPT	~18–20	Highest frequency; foundational framework articulations across all chapters
EXECUTIVE DECISION	~12–15	Decision-oriented synthesis at the close of major analytical developments
WORKED EXAMPLE	~8–10	Detailed applications including the Bezos memo (Ch 8) and objection sequences (Ch 11)
CASE STUDY	~12–14	Synthesis of lessons from opening, regional, and worked-

Preface

This book has been written from a position substantially shaped by my engagement with analytical and AI matters in the Thai manufacturing context across the past decade. The substantive engagement has extended across the academic dimensions of doctoral and post-doctoral research, the professional dimensions of work within a textile manufacturing operation employing approximately fifteen hundred personnel in Samut Sakhon Province, and the broader practitioner dimensions of engagement with the contemporary literature, the practitioner community, and the regulatory and institutional environment within which Thai industrial operations conduct their analytical activities.

The book is written for three principal audiences. The first audience is the MSc-level student of analytics, AI leadership, or related disciplines who is engaging substantively with the contemporary practice that the cumulative apparatus of the book addresses. The book is structured to support twelve weeks of substantive engagement, with each chapter providing approximately one week of substantive material including the substantive content, the case studies, the critical thinking questions, and the broader supplementary material. The second audience is the practitioner engaged with analytical and AI activities across substantively varying organisational contexts. The book provides the substantive frameworks, templates, and disciplines that the contemporary practice has developed, with the application to the practitioner's specific context being itself substantive work that the book supports rather than substitutes for. The third audience is the senior executive whose substantive responsibilities include analytical and AI matters. The book provides the substantive content that the executive engagement requires while recognising that the executive's substantive judgement and contextual engagement extend beyond what any text alone can provide.

The book operates from a particular perspective that warrants explicit articulation. The perspective is that of the emerging-economy practitioner — the analytical leader whose substantive engagement occurs principally outside the developed-economy contexts that the substantial majority of the contemporary practitioner literature has been written from and for. The emerging-economy perspective is not separable from the broader analytical leadership practice; the substantive disciplines that the book develops apply across both developed-economy and emerging-economy contexts. The perspective is nonetheless visible in the regional case selection, the running case construction, the broader contextual engagement, and the substantive engagement with regulatory and institutional environments that differ across operating jurisdictions. The perspective is offered not as opposition to developed-economy practice but as substantive complement to it; the contemporary analytical leadership practice operates across both contexts, and the cumulative practitioner community benefits from the integration of perspectives across them.

The book also operates from a particular methodological perspective that warrants explicit articulation. The perspective treats analytical and AI activities as substantive organisational practice rather than as principally technical matters. The substantive technical content is necessary but insufficient; the analytical and AI activities operate within the broader organisational dimensions of strategy, governance, risk management, communication, sustainability, strategic decision-making, and leadership that the book's twelve chapters address. The substantive integration of the technical and organisational dimensions is the principal content of contemporary analytical leadership practice. Books that operate principally through technical content without the organisational integration produce the partial preparation that the contemporary practitioner cannot deploy substantively; books that operate principally through organisational content without the technical integration produce the equally partial preparation that the cumulative practice requires both dimensions of.

The running case study of Elastic Company has been developed substantively across the twelve chapters as the principal illustrative thread the book operates through. The case is constructed; the company is not an operational entity in the form the case describes. The construction is nonetheless substantive rather than fictional; the case incorporates substantive content from real Thai textile manufacturing operations including labour dynamics, regulatory environment, customer relationships, technology adoption patterns, governance practices, and broader operational matters. The construction permits the substantive illustration of the cumulative apparatus across the twelve chapters in ways that any single operational entity would not have permitted given the substantive confidentiality considerations the substantive content engages. The construction is intended to support the substantive engagement of the reader with the cumulative apparatus rather than to provide proprietary operational detail of any specific entity.

The case studies of real organisations — the twelve opening cases and the twelve regional cases — have been developed from publicly available sources including annual reports, regulatory filings, financial disclosures, academic research, and journalistic reporting. The substantive interpretation of strategy and practice in each case is the author's own; the interpretation may not coincide with the organisations' own articulations of their strategy and practice, and the case studies should be read as analytical interpretations rather than as authoritative organisational accounts. Where specific facts have been articulated, the substantive accuracy has been verified against publicly available sources to the extent the verification has been possible. Errors of fact or interpretation that remain are the author's responsibility.

The book has been written across an extended period. The substantive content reflects the contemporary literature and practice as it has developed across the recent period, with the chapters operating principally from the position of contemporary practice as of approximately 2025–2026. The contemporary practice continues to develop substantively, with the consequence that specific elements of the book may require substantive revision across subsequent editions. The substantive framework, however, is intended to support the contemporary practitioner across the strategic horizon that the book's content addresses, with the framework being substantively durable across the specific developments that the broader practice will produce. The first edition is offered as the substantive starting point for the contemporary engagement; the substantive feedback from readers, the cumulative practitioner experience, and the broader developments of the contemporary moment will support the subsequent editions that the broader practice may warrant.

The substantive contribution that the book attempts is to support the substantive practice of analytical leadership in the contemporary moment. The book does not substitute for the substantive engagement of the practitioner with the specific organisational context the practitioner operates within. The book provides the substantive frameworks, templates, and disciplines that the contemporary practice has developed; the substantive application across the practitioner's context is itself substantive work that the practitioner must engage with deliberately. The substantive engagement of the reader with the cumulative content, the substantive application across the reader's specific context, and the substantive contribution of the reader to the broader practice are what the book is designed to support across the practitioner's career and across the broader development of the contemporary analytical leadership practice.

Thanakit Ouanhlee
Thailand 2026

Acknowledgements

This book has been written with substantive support from a community of colleagues, mentors, students, and practitioners whose engagement has shaped the substantive content across the twelve chapters. The acknowledgement of the community is itself substantive component of the book's broader positioning, with the cumulative practice that the book attempts to articulate being itself the product of many years of collective engagement across the broader community.

My doctoral and post-doctoral supervisors and colleagues across California Intercontinental University, the University of Salford, and the University of Lancashire have shaped my substantive engagement with the academic disciplines that the book operates within. The substantive engagement with quantitative research, qualitative research, mixed-methods research, and the broader academic practice has developed across the cumulative engagement with these institutions across approximately the past decade, with the cumulative learning being substantively visible in the book's analytical approach and methodological positioning.

My colleagues across the Thai manufacturing community — particularly within the textile and broader light manufacturing sectors of Samut Sakhon Province and the surrounding industrial regions — have provided the substantive operational context that the book's running case construction operates from. The substantive engagement with day-to-day operations, with workforce dynamics, with regulatory and institutional considerations, and with the broader practical content of Thai industrial operations has substantively shaped the book's content beyond what purely academic engagement could have produced.

The broader practitioner community across Thailand, Southeast Asia, and the broader emerging-economy contexts has provided the substantive perspective that distinguishes the book's positioning from the developed-economy practitioner literature alone. The substantive conversations, the broader professional engagement, and the cumulative perspective that the emerging-economy practitioner community operates with have substantively shaped the case selection, the analytical framing, and the broader content positioning that the book operates through.

The students of the MSc programmes I have engaged with across the recent period — particularly the cohorts of the BM4041 Leading with Analytics and AI course at the University of Lancashire — have provided substantive engagement with the analytical leadership content that has substantively shaped the book's pedagogical structure. The substantive questions, the substantive case discussions, the broader engagement with the contemporary material that the students have brought has supported the book's development beyond what purely individual reflection could have produced.

The broader academic community engaged with analytics, artificial intelligence, business administration, organisational behaviour, and related disciplines has provided the substantive literature that the book draws on across its twelve chapters. The substantive engagement with the contemporary practitioner literature, the academic research literature, and the broader analytical community has supported the integration that the book attempts; the analytical positions the book articulates have been developed in substantive engagement with the broader scholarly conversation that the cumulative community operates within.

My family has supported the substantive engagement that book writing across an extended period requires. The substantive engagement is not separable from the broader life context within which the engagement occurs, with the family support being itself substantive component of the broader work that the book represents. The substantive support across the period of writing has been substantive practical contribution that the cumulative work depends on, with the contribution being acknowledged here in substantive recognition of the broader content that the family relationships provide.

Errors of fact, interpretation, or analytical positioning that remain in the book are my own responsibility. The substantive support of the community has shaped the book's strengths; the limitations are independent of the support and reflect the substantive boundaries that any single author working across the breadth of contemporary analytical leadership practice necessarily encounters. The substantive feedback of readers across the practitioner and academic communities will be welcomed as substantive contribution to subsequent editions that the broader development of contemporary practice may warrant.

How to Use This Book

The book is structured to support multiple modes of engagement across the principal audiences. This section provides substantive guidance for each audience and use case.

For the MSc Student

The book is structured to support a twelve-week MSc course operating with one chapter per week. Each chapter provides approximately one week of substantive engagement at the level that a substantial postgraduate course warrants. The structural elements within each chapter support specific pedagogical functions that the student should engage with deliberately.

The Learning Objectives at the opening of each chapter articulate the substantive capabilities the chapter develops. The student should engage with the objectives at the beginning of the chapter as articulation of what the chapter will support and at the end of the chapter as substantive self-assessment of whether the substantive engagement has produced the articulated capabilities.

The Key Terms list at the opening of each chapter introduces the principal terminology the chapter operates through. The student should engage with the terms substantively as the chapter develops, with the cumulative terminology supporting the substantive engagement with the broader practice the book addresses. The consolidated glossary at the end of the book provides the integrated reference across all twelve chapters.

The Chapter Introduction substantively introduces the topic the chapter addresses through a seven-paragraph structure that the student should engage with as substantive content rather than as transitional material. The introduction articulates the substantive case for the chapter's content, the broader contextual positioning, and the substantive disciplines the chapter will develop.

The Opening Case Study, Running Case Study (Elastic Company), and Regional Case Study within each chapter provide substantive case material that the student should engage with through the discussion questions and the broader analytical engagement that case-based learning supports. The cases are not illustrations of the analytical content but substantive material that the analytical content engages with substantively.

The numbered sections within each chapter develop the substantive analytical content. The student should engage with the sections substantively rather than only at the summary level, with the substantive engagement supporting the cumulative learning that the chapter is designed to produce.

The Critical Thinking Questions at the end of each chapter — typically twelve questions in imperative form — are designed to test substantive judgement rather than memorisation. The student should engage with the questions substantively, with the substantive engagement frequently producing more substantive learning than the chapter content alone would have generated. Several questions have no single correct answer; the value is in the quality of the reasoning the student brings to them.

The Suggested Further Reading provides the substantive bibliography supporting deeper engagement with specific dimensions of the chapter content. The student engaged with a research project, dissertation, or other substantial analytical work should engage with the further reading substantively rather than treating it as supplementary matter.

For the Practitioner

The book is structured to support substantive practitioner engagement across multiple use cases. The substantive disciplines, frameworks, and templates that the book develops can be applied substantively across the practitioner's specific organisational context with appropriate adaptation.

The practitioner whose substantive engagement is with one specific area of analytical leadership can engage with the chapters specifically relevant to the area, with the cross-chapter references providing the substantive integration that the cumulative apparatus operates through. The cumulative apparatus benefits from engagement with all twelve chapters, but specific chapters can support specific substantive engagement.

The Templates introduced in Chapter 3 — Template 2 (the ten-field KPI/KRI card) and Template 3 (the six-field risk register entry) — are designed for direct application by the practitioner. The substantive completion of the templates against the practitioner's specific analytical and risk content produces the substantive measurement and risk architecture that the broader practice operates within.

The Frameworks across the chapters — the NIST AI RMF treatment in Chapter 5, the AI risk taxonomy in Chapter 6, the three-tier dashboard architecture in Chapter 7, the six-audience communication map in Chapter 8, the strategic decision categories in Chapter 10, the pitch structural elements in Chapter 11, the capability renewal cycle in Chapter 12 — are designed for substantive application with appropriate adaptation to the practitioner's specific context.

The Running Case Study of Elastic Company provides substantive material that the practitioner can engage with comparatively. The cumulative case across the twelve chapters illustrates the integrated apparatus that the broader practice operates within; the practitioner's own organisational context will differ substantively from the running case, with the comparative engagement producing substantive learning about the practitioner's specific context that the case alone would not have generated.

For the Senior Executive

The book is structured to support the senior executive's substantive engagement with analytical and AI matters at the level the executive role requires. The executive's substantive engagement typically operates at strategic, governance, and broader institutional levels rather than at operational levels, with the book's content supporting the substantive engagement across these levels.

The Executive Decision callouts within the chapters distil the substantive decisions that the analytical content supports. The executive should engage with the callouts substantively as articulation of the decisions the analytical content is designed to support.

The Chapter Summaries at the end of each chapter provide substantive synthesis that the executive can engage with as principal reading, with the deeper engagement with specific sections being substantive supplement rather than substantive requirement. The chapter summaries are designed to be substantively complete rather than abridged versions of the chapter content.

The Ethics and Governance Discussions within each chapter address substantive content that the executive role specifically requires. The discussions articulate the substantive ethical and governance considerations that the analytical work engages, with the considerations being substantively addressed rather than only documented.

The Opening Cases, Regional Cases, and Running Case together provide substantive material for executive engagement with comparable organisations. The substantive lessons from the cases can be applied to the executive's own organisational context with appropriate substantive judgement.

For the Course Instructor

The book is structured to support instructors operating one-semester or two-semester courses in analytics leadership, AI strategy, business analytics, or related disciplines. The twelve-chapter structure maps directly to a twelve-week course, with each chapter providing approximately one week of substantive course material.

Instructors operating shorter courses can select substantive subsets of the chapters supporting specific course objectives. A six-week course in analytical governance and risk might combine Chapters 1, 5, 6, 9, 10, and 11. An eight-week course in analytical operations might combine Chapters 1, 2, 3, 4, 6, 7, 8, and 12. The substantive integration across the chapters means that selected subsets benefit from explicit treatment of the integration across the selected material.

The Critical Thinking Questions at the end of each chapter can serve as discussion prompts, assignment prompts, and assessment prompts. The questions are designed to test judgement rather than memorisation; assessments using the questions should accordingly engage with the substantive reasoning the student brings rather than seeking specific correct answers.

The Case Studies — opening, regional, and running — provide substantive case material for course-based case discussion. The discussion questions following each opening case are designed for in-class engagement. The running case across the twelve chapters provides cumulative case material that can support a sustained case discussion thread across the course.

The supplementary instructor's manual (separate document) provides additional substantive material including discussion guidance for the case studies, suggested syllabus mappings for varying course lengths, assessment rubrics, and slide deck outlines per chapter. The substantive engagement of the instructor with the supplementary material supports the substantive course delivery that the cumulative book supports.

Chapter 1

The Leader's Role in the Age of Analytics and AI

Learning Objectives

After studying this chapter, the reader should be able to:

Trace the historical evolution of leadership decision-making from intuition-based authority to evidence-based and algorithm-augmented authority, and explain why each transition required new leadership capabilities.

Distinguish clearly between data, analytics, and artificial intelligence as leadership tools, and place each on the descriptive-diagnostic-predictive-prescriptive analytics continuum.

Identify and apply the four modes of decision authority that contemporary leaders exercise, recognising which mode is appropriate for which class of decision.

Articulate the five core responsibilities of the analytical leader — sponsorship, stewardship, sense-making, storytelling, and talent stewardship — and explain why each is non-delegable.

Apply the Davenport DELTA framework and the five-stage analytics maturity model to assess organisational readiness for analytics and AI adoption.

Recognise the four most common failure modes that derail leaders attempting analytical transformation, and design specific governance safeguards against each.

Evaluate the boardroom implications of AI deployment using comparative regulatory frameworks across the European Union, the United States, Singapore, Thailand, Brazil, and India.

Key Terms

Algorithmic accountability; analytical leadership; analytics maturity; artificial intelligence; augmented decision-making; black-box model; board-level data literacy; data-driven decision-making; DELTA framework; descriptive analytics; diagnostic analytics; evidence-based management; explainable AI (XAI); generative AI; human-in-the-loop; predictive analytics; prescriptive analytics; sponsorship; stewardship; supervised learning.

Chapter Introduction

Consider three leaders, each facing an analytics-driven decision on the same Monday morning in 2026. The first is a logistics chief executive in Bonn, deciding whether to act on an AI-generated forecast that a tropical cyclone in the Bay of Bengal will disrupt one-third of Asian container routes within seventy-two hours. The second is a banking chief executive in Singapore, weighing the launch of a generative AI customer-service platform that her engineers say will reduce call-centre headcount by forty percent — but whose ethical implications her board has not yet fully discussed. The third is a textile chief operating officer in Samut Sakhon, Thailand, opening a quarterly report that ranks two hundred and eighty-six of his workers as having an elevated probability of voluntary resignation within six months, each accompanied by an algorithmic recommendation for individualised retention intervention.

These three scenarios share a structural property that defines the contemporary leadership task. In each case, the decision is informed — perhaps even shaped — by analytical machinery that the leader did not personally build, cannot personally audit in detail, and yet is personally accountable for. The forecasts may be probabilistically sound; the recommendations may be statistically defensible; the algorithms may have been validated by competent technical teams. None of that reduces the leader's responsibility to decide whether to act, in what manner, with what safeguards, and toward what end. Analytics and artificial intelligence have not replaced leadership. They have raised the stakes of leadership.

The leader's role in the age of analytics and AI is a particular set of obligations, distinct from the roles that surround it. It is not the data scientist's role of constructing models, nor the technology director's role of running platforms, nor the classical executive's role of pure judgement informed only by experience. It is something more specific: the responsibility to govern, sponsor, interpret, and decide in an organisation where significant cognitive work — pattern recognition, prediction, recommendation, increasingly even generation — is performed by machines that the senior leader did not build and does not fully understand. The role is defined less by what the leader personally does at the keyboard than by what the leader is accountable for at the boundary between human judgement and machine output.

The novelty of the contemporary leader's situation is not the use of evidence in decision-making. Management has been an evidence-using discipline since Taylor, and the most penetrating writers on it — Drucker, Deming, Pfeffer, Davenport — have argued for a century that organisational performance improves when leaders subordinate intuition to systematic measurement. The novelty is twofold. First, machines now perform interpretive and predictive work that previously could only be done by skilled human professionals. Second, these machines increasingly act on personal data and produce consequential decisions about identifiable individuals, raising legitimacy and accountability questions that the older traditions of evidence-based management did not have to answer. The leader's role has therefore been redefined less by the existence of new computational tools than by the legitimacy questions those tools have brought into being.

The stakes of getting this role right are higher than they first appear. A leader who fails at analytical governance does not merely produce worse decisions in the technical sense; she produces decisions that lack moral and reputational legitimacy. A turnover prediction system that flags employees by algorithm without visible human accountability is not merely an analytical weakness when it fails — it is a governance failure that damages institutional trust. A pricing algorithm that produces unintended discriminatory outcomes is not merely a model defect — it is a board-level liability. A workforce planning model whose recommendations are accepted without review is not merely a process risk — it is the abdication of a fiduciary duty. The leader who has not personally engaged with the analytical systems operating under her authority cannot defend them when they are challenged, and they will be challenged.

Analytical leadership, understood properly, is the integrated practice of four competencies held simultaneously. The first is sufficient analytical literacy to ask substantive questions of technical teams and to recognise substantive answers — not the depth of a data scientist, but the fluency of an informed reader. The second is sufficient organisational fluency to govern how analytical capabilities are adopted, resisted, and embedded in workflow, because analytical systems that are technically excellent but organisationally rejected produce no value. The third is sufficient ethical clarity to draw the lines that algorithms cannot draw — what data may be used, what decisions warrant human review, what trade-offs the organisation will and will not make in the name of efficiency. The fourth is sufficient strategic discipline to direct analytical investment toward purposes that matter rather than toward technical novelty for its own sake. None of these competencies, on its own, is sufficient. Leaders who have only analytical sophistication produce technically excellent failures that the organisation cannot use. Leaders who have only governance instincts produce well-protected mediocrity. Leaders who have only strategic clarity produce visions that the organisation cannot implement. The combination is the discipline.

The leader's role in the age of analytics and AI is therefore not the passive one of receiving analytical outputs and approving them. It is the active construction of an organisation in which analytical capability, human judgement, and ethical commitment cohere — and in which the leader, personally, can be held accountable for what the integrated system produces. This is a more demanding conception of leadership than the one that prevailed in the data warehousing era of the 1990s or even in the big data decade that followed. It is also, for those who accept it, a more consequential one. The leaders whose contributions will be remembered from the present period are those who built organisations capable of operating intelligently at the boundary between human and machine — not those who delegated the boundary to specialists, and not those who pretended the boundary did not exist.

Opening Case Study

DHL Group: From Logistics Giant to Analytics-Driven Enterprise

In 2008, Frank Appel succeeded Klaus Zumwinkel as Chief Executive Officer of Deutsche Post DHL Group. He inherited a sprawling logistics enterprise — the result of the 1998 acquisition of DHL by the former German postal monopoly — that was operationally complex, geographically vast, and analytically primitive. The group operated in more than two hundred and twenty countries, employed close to five hundred thousand people, and moved more than one and a half billion shipments per year. Its data was abundant. Its analytical capability was not. Within fifteen years, by the time of Appel's retirement at the end of 2022, DHL Group had become one of the most consistently cited corporate exemplars of board-led, enterprise-wide analytics transformation in the management literature.

The transformation was not a single decision but a sustained sequence of decisions, each of which illustrates a different facet of analytical leadership. Five initiatives in particular merit close attention because they will reappear, in different forms, throughout this book.

Initiative 1 — Resilience360 and the Supply Chain Risk Platform

In 2014, DHL launched Resilience360, a supply chain risk management platform that aggregated data from weather services, geopolitical risk feeds, port congestion sensors, customs systems, and social media to provide real-time visibility into potential disruptions across global supply chains. The platform was significant for two reasons. First, it represented a strategic recognition that the most valuable analytical asset DHL possessed was not its operational data alone but its position as an aggregator of supply chain signals from across its customer base. Second, it was commercialised as a paid service to enterprise customers, transforming what could have been an internal cost centre into a revenue-generating analytical product. The leadership lesson is one the book returns to repeatedly: the most valuable analytical capabilities are often those that sit at the intersection of internal data and external context.

Initiative 2 — The Integrated Data and Enterprise Analytics Platform

Between 2016 and 2019, DHL invested heavily in unifying data assets that had previously been fragmented across business divisions — Express, Supply Chain, Global Forwarding, and Post & Parcel Germany. The Integrated Data and Enterprise Analytics platform was the technical embodiment of a strategic decision: that DHL would treat data as a corporate asset rather than a divisional one. This required not only technical work but governance work — agreement on shared data standards, common definitions of customer and shipment, and a central data office with authority across divisions. The leadership lesson here is structural: most analytical transformations stall at the boundary between divisions, and only a chief executive with sustained personal sponsorship can force the unification.

Initiative 3 — Predictive Network Management

From 2018 onward, DHL deployed machine learning models that predicted shipment volumes, transit times, and capacity bottlenecks across its global network with sufficient accuracy to allow proactive rerouting before disruption occurred. The Predictive Network Management capability did not replace human operations managers; it provided them with twelve- to forty-eight-hour forecasts that allowed them to make decisions before, rather than after, network stress materialised. The leadership lesson is about decision modes: the predictive system shifted operations managers from reactive to proactive decision authority, but the decisions themselves remained human, with the algorithm acting as a sense-making aid rather than as a decision-maker.

Initiative 4 — OnTrac and Last-Mile Route Optimisation

Last-mile delivery — the final segment from a distribution hub to the customer's address — accounts for a disproportionate share of logistics cost and customer-experience variance. DHL deployed AI-powered route optimisation systems that recalculated delivery sequences in near-real time based on traffic, customer availability, and parcel priority. In the most sophisticated deployments, the algorithm itself made route decisions within human-defined constraints, with the human role shifting to exception management. This is an example of what the chapter will later call the fourth mode of decision authority — algorithm-driven within bounded delegation.

Initiative 5 — Workforce Analytics and the People Dimension

From 2019, DHL extended its analytical capability to workforce decisions — recruitment, retention, engagement, and skills planning. This was the most ethically charged of the five initiatives and the one that required the most explicit governance. DHL's leadership made a public commitment that workforce analytics would inform decisions but never make them, that all employees would have access to the data held about them, and that retention recommendations would be reviewed by human managers with the authority to override. The leadership lesson is that the legitimacy of workforce analytics depends entirely on the visible accountability of the human leader, not on the sophistication of the algorithm.

Synthesis: Five Leadership Lessons from DHL

CASE STUDY — Recurring Lessons from the DHL Transformation

Lesson 1 — Sponsorship from the top is non-negotiable. The DHL transformation succeeded because Appel personally sponsored it for fifteen consecutive years. Analytical transformation is a multi-year project that outlasts most quarterly executive attention cycles; only sustained chief-executive attention produces sustained organisational follow-through.

Lesson 2 — Data assets must be unified across silos. Most analytical value is unlocked by combining data that previously lived in separate divisions. The chief executive is the only person with the authority to force this unification against the natural resistance of divisional leaders.

Lesson 3 — Investment in talent precedes investment in technology. DHL spent heavily on analytical talent — not only data scientists but data engineers, data product managers, and analytics-fluent business leaders — before scaling its technology investments. Without the human capability, the technology produces nothing.

Lesson 4 — Customer-facing analytical outcomes drive adoption faster than internal efficiency metrics. Initiatives that visibly improved customer experience created adoption momentum; initiatives framed only as internal cost-cutting struggled. The leader's framing of the purpose of analytics matters enormously for its acceptance.

Lesson 5 — Ethical frameworks must precede, not follow, scale deployment. The workforce analytics initiative succeeded because its governance was settled before deployment. When ethical frameworks are bolted on after the fact, they are typically weaker and slower to gain organisational legitimacy.

Discussion Questions on the DHL Case

DHL's transformation took approximately fifteen years from initial board commitment to fully integrated analytics culture. What constraints might prevent a mid-cap manufacturer — for example, a 1,500-employee textile company in Thailand — from following a similar timeline? What compressed approach might be feasible, and what would be lost in the compression?

Of the five DHL initiatives summarised above, classify each as primarily descriptive, predictive, prescriptive, or algorithm-driven, and explain how that classification affects the leadership skills required to govern each.

DHL chose to commercialise Resilience360 as a paid product. What does this decision reveal about the strategic value the board placed on its analytical capability, and what risks accompanied that decision?

The workforce analytics initiative was the most ethically charged. Imagining the role of the chief executive proposing it to the board, identify the three governance commitments that should appear in the opening paper, and justify each.

Chapter 2

From Data to Decisions: The Analytics and AI Toolkit

Learning Objectives

After studying this chapter, the reader should be able to:

Trace the eight-stage analytics lifecycle from problem framing through monitoring, and identify the leadership decision required at each stage.

Distinguish among the principal data infrastructure patterns — operational systems, data warehouses, data lakes, data lakehouses, and streaming architectures — and recognise when each is appropriate.

Identify the principal families of machine learning models — supervised classification and regression, unsupervised clustering and dimensionality reduction, time-series forecasting, reinforcement learning, and deep learning — and match each to the business problems it addresses.

Apply the principal model evaluation metrics — accuracy, precision, recall, F1 score, AUC, calibration, and fairness metrics — and recognise the situations in which each metric can mislead.

Frame the build-versus-buy-versus-partner decision for analytical capability, identifying the five factors that should determine the choice and the typical sequencing of capability development.

Recognise the leader-relevant elements of generative AI and agentic AI, including the distinctive governance challenges of hallucination, prompt injection, and autonomous action.

Engage substantively with technical teams using shared vocabulary, while maintaining appropriate humility about the limits of leadership-level technical literacy.

Key Terms

A/B testing; accuracy; AUC (area under the ROC curve); baseline model; build-versus-buy decision; calibration; classification; clustering; confusion matrix; cross-validation; data engineer; data lake; data lakehouse; data scientist; data warehouse; deployment; ETL and ELT; F1 score; feature engineering; foundation model; gradient boosting; hallucination; hyperparameter tuning; MLOps; model drift; precision; prompt engineering; recall; regression; retrieval-augmented generation (RAG); ROC curve; streaming analytics; supervised learning; training set; unsupervised learning; validation set.

Chapter Introduction

Three executives are reviewing the same analytical proposal. The first, the chief financial officer, scans the executive summary, sees the headline figure that the model achieves ninety-two percent accuracy, and approves the investment. The second, the chief operating officer, reads the same summary, pauses on the accuracy figure, and asks the data science team a different question: what is the false negative rate on the population that actually matters — the segment of employees flagged as high-performers? The answer turns out to be substantially less impressive than the headline accuracy suggests. The third, the chief executive, hears both colleagues and asks a question that neither has asked: how was the training set constructed, and does its construction reflect the population on which the model will operate? The answer reveals that the training data was drawn from a period in which the company's wage structure differed materially from its present structure, with implications for how the model will perform in the months ahead.

The three executives are not differently intelligent. They differ in their analytical literacy — the depth of technical understanding that informs the questions they ask. The first executive asked a question that any board paper could answer; he received a polished answer and learned nothing he did not already know. The second asked a question that exposed a substantive weakness; she received a more honest answer and approved the investment with a defined remediation plan. The third asked a question that exposed a deeper weakness; she sent the project back for redesign. The differences in outcome cost nothing to produce in technical terms. They are differences in the questions the leaders knew to ask.

The analytics and AI toolkit, properly understood, is the set of technical capabilities, methods, and concepts that an organisation uses to turn data into decisions. It consists of four interconnected layers. The first is the data layer — the systems, structures, and processes by which data is captured, stored, integrated, and made available for analysis. The second is the modelling layer — the families of statistical and machine-learning methods that produce predictions, classifications, recommendations, and increasingly the generative outputs of large language models. The third is the evaluation layer — the metrics, validation procedures, and tests by which analytical work is judged sound or unsound, useful or misleading. The fourth is the deployment layer — the mechanisms by which models move from analytical artefacts into operational use, and the ongoing monitoring that keeps them honest over time. The leader who has a working grasp of all four layers can engage substantively with each. The leader who has a working grasp of none operates with whatever the technical team chooses to present.

Three forces make this technical literacy more urgent in the present moment than it has been in earlier analytical eras. The first is the commodification of basic analytical capability. Twenty years ago, the construction of a serviceable predictive model required the work of a small team of specialists over several months. Today, an analyst with a laptop, a public dataset, and an open-source machine learning library can produce a model of comparable technical sophistication in a single afternoon. The barrier to producing models has collapsed; the barrier to using them well has not. The second is the sophistication of the failure modes. The models being deployed in contemporary organisations are more capable than their predecessors but also more opaque, more sensitive to subtle data shifts, and more capable of producing confidently wrong outputs in ways that cannot be detected by the casual reviewer. The third is the generative AI inflection. Foundation models — large language models trained on internet-scale text corpora and adapted for specific applications — present a category of analytical tool that did not exist in mature commercial form before 2022 and that introduces governance questions the older analytics literature did not have to address.

The stakes of inadequate technical literacy at the leadership level are operational, financial, and reputational. The operational cost is that the leader cannot distinguish a model that should be trusted from a model that should not, with the predictable consequence that bad models are trusted and good ones are over-questioned. The financial cost is that the organisation pays premium prices for capabilities it could build, builds capabilities it should buy, and partners where it should not — because the build-versus-buy-versus-partner conversation requires a level of technical understanding that the leader without literacy cannot bring to the room. The reputational cost is that the leader who cannot interrogate the analytical machinery operating under her authority cannot defend it when it is challenged in public, in court, or before a regulator. Each of these costs accumulates quietly across an analytical programme and becomes visible only when the programme is in difficulty, at which point remediation is expensive and slow.

Adequate technical literacy at the leadership level does not mean technical expertise. It does not require the leader to write code, derive equations, or construct neural network architectures. It requires something different: knowing what each tool in the analytical toolkit can and cannot do; knowing what questions to ask of a technical team in order to expose substantive weaknesses; recognising the warning signs that an analytical effort is in trouble; and possessing the vocabulary to participate as an informed reader rather than a passive recipient of conclusions. The line between informed reader and active practitioner is one the leader should not attempt to cross, for the same reason that a senior judge does not perform her own forensic accounting. But the line between informed reader and uninformed authoriser is one the leader must cross, because the alternative is to govern what cannot be understood.

Analytical leadership in the contemporary era requires, at minimum, the vocabulary of an informed reader of analytical work. The leader who is fluent in this vocabulary can engage with technical teams, can recognise the difference between a sound result and a polished one, can challenge the assumptions buried in any analytical artefact, and can hold the organisation to a standard that pure technical accountability would not produce. The leader who lacks this fluency is, in practice, dependent on the goodwill and self-discipline of her technical staff — a dependency that is often well-rewarded but is structurally precarious and that, when it fails, fails expensively.

Opening Case Study

Walmart: From Retail Link to Retail Intelligence at Planetary Scale

Walmart is among the most thoroughly documented examples of large-scale analytical capability development in modern business history. The company's relationship with data has been continuous and central since the 1980s, when its founder Sam Walton first invested in computerised inventory tracking at a scale that was unusual for general-merchandise retail at the time. What distinguishes Walmart from many of its peers is not that it adopted analytical technology — many retailers did — but that it sustained the investment, integrated it across the enterprise, and built a culture in which analytical capability was treated as a primary basis of competitive advantage rather than as an operational support function. The arc of that investment, traced across four decades, offers an unusually rich set of lessons about how an analytical toolkit is assembled, refreshed, and extended over the lifetime of an enterprise.

Phase 1 — Retail Link and the Information Symmetry Strategy (1990s)

In 1991, Walmart launched Retail Link, a supplier-facing data platform that gave the company's suppliers near-real-time access to sales and inventory data for their products across Walmart stores. The strategic logic was unusual. Most retailers of the period treated point-of-sale data as a competitive asset to be guarded from suppliers, who would otherwise use it to optimise their own commercial position. Walmart inverted the logic. The company recognised that the most valuable analytical work on its products would be done by the suppliers themselves, who possessed deeper category expertise than Walmart could replicate. By giving suppliers the data they needed to forecast demand, manage replenishment, and refine product mix, Walmart effectively externalised a substantial portion of its analytical work to its supplier base — and reduced its inventory carrying costs in the process. The platform was technologically unremarkable by modern standards but strategically pioneering. The toolkit lesson is foundational: the analytical infrastructure of a leading firm is not only the systems the firm operates for itself but also the systems it operates jointly with its partners.

Phase 2 — Data Warehousing at Unprecedented Scale (1990s through 2000s)

Walmart's data warehouse, operated for many years on the Teradata platform, was for an extended period the largest commercial data warehouse in the world. The investment was financially significant — at peak, the warehouse stored hundreds of terabytes of transactional and operational data, with continuous additions running into the gigabytes per day. The strategic logic was that scale of data was itself a competitive moat: any retailer could buy the same software and hire the same consultants, but only Walmart possessed the underlying data volume that made certain analytical questions answerable. Demand sensing at the level of an individual product in an individual store on an individual day required years of accumulated transactional history; competitors with smaller data assets could not perform the same analyses, regardless of how sophisticated their analytical methods. The toolkit lesson is that, for certain analytical applications, the durable advantage lies in the data itself rather than in the algorithms applied to it.

Phase 3 — Acquisition-Driven Capability Acceleration (2010s)

During the 2010s, Walmart accelerated its analytical capability development through a sequence of acquisitions whose primary value was not the acquired product but the acquired analytical talent and technology. Jet.com was acquired in 2016 for approximately three billion US dollars; while the e-commerce site itself was significant, the acquisition was widely understood to be principally about acquiring Marc Lore and his engineering organisation, who subsequently led Walmart's e-commerce transformation. Bonobos, Modcloth, and ShoeBuy followed; Flipkart was acquired in 2018 in a transaction valued at sixteen billion US dollars, bringing not only the leading Indian e-commerce platform but also one of the most sophisticated emerging-market data science organisations available for acquisition. The toolkit lesson is that, beyond a certain scale, analytical capability is built as efficiently by acquisition as by organic development — but only if the acquiring organisation has the integration discipline to absorb the acquired capability into its operations rather than allowing it to wither in isolation.

Phase 4 — The Machine Learning Platform Era (late 2010s onward)

By the late 2010s, Walmart had moved beyond traditional business intelligence and structured machine learning into a platform model in which a central machine learning infrastructure supported hundreds of distinct predictive applications across the enterprise — demand forecasting at the store-product-day level, dynamic pricing on selected product categories, supply chain optimisation, fraud detection, customer churn modelling, marketing personalisation, and others. The platform approach achieved economies of scale that project-by-project analytics could not match. A new use case could be developed against the existing platform in weeks rather than months; common challenges such as data versioning, model monitoring, and feature reuse were addressed once for the platform rather than repeatedly for each project. The toolkit lesson is that, at scale, the principal investment is not in individual models but in the platform infrastructure that allows models to be developed, deployed, and monitored efficiently.

Phase 5 — Generative AI and the Frontier (2023 onward)

From 2023 onward, Walmart has been among the most visible large enterprises to deploy generative AI in production. Applications have included generative AI assistance for store associates seeking product information, AI-augmented supplier negotiations using technology partners such as Pactum, generative search and discovery on the consumer-facing storefront, and code generation tools for the company's substantial internal engineering organisation. The deployments are notable not only for their scale but for the explicit governance frameworks that surround them — published responsible AI principles, named accountability for each application, structured human-review thresholds for consequential decisions. The toolkit lesson is that generative AI does not replace the earlier analytical capabilities but extends the toolkit into a new layer that interacts with all the layers beneath. Foundation models retrieve from data warehouses, are evaluated against bespoke metrics, and are deployed through platform infrastructure built for earlier model classes.

Synthesis: Five Toolkit Lessons from Walmart

CASE STUDY — Recurring Lessons from the Walmart Trajectory

Lesson 1 — Scale of data is a durable competitive moat. Walmart's leadership in data-driven retail rests on transactional data accumulated over decades. Competitors with less history cannot replicate certain analyses regardless of their algorithmic sophistication. The leader should think about the organisation's data assets as long-horizon investments whose value compounds.

Lesson 2 — Build, buy, and partner are not exclusive choices. Walmart built core infrastructure, partnered with suppliers on shared data platforms, acquired e-commerce capability where speed was decisive, and bought specialised tools where the build economics did not favour internal development. The toolkit decision is a portfolio decision, not a single choice.

Lesson 3 — Platform investment exceeds project investment in long-term value. Most of Walmart's contemporary analytical productivity comes from platform infrastructure that supports hundreds of use cases, not from individual celebrated projects. The leader should be sceptical of analytical organisations that show many projects but no platform underneath them.

Lesson 4 — Talent acquisition often proceeds through company acquisition. Walmart's most consequential talent additions came as parts of company acquisitions, not as individual hires. In tight talent markets — including most emerging-economy markets — this pattern is more efficient than open-market recruitment for senior analytical leadership.

Lesson 5 — New analytical tools extend the toolkit; they do not replace it. Generative AI sits on top of the data warehouses, machine learning platforms, and governance frameworks that came before. The leader who imagines that generative AI obsoletes the earlier toolkit will be unprepared for the operational realities of deploying generative AI in production.

Discussion Questions on the Walmart Case

Walmart's Retail Link platform externalised analytical work to its supplier base. Identify the conditions under which a similar strategy would be effective for a mid-cap manufacturer or distributor. Describe the conditions under which it would fail.

The platform-versus-project distinction is central to Walmart's contemporary analytical productivity. For an organisation at an earlier stage of analytical development, identify the threshold at which platform investment becomes economically justified, and the indicators that the threshold has been reached.

Walmart's acquisition of Jet.com is widely understood to have been a talent acquisition more than a business acquisition. Evaluate the governance and integration challenges this pattern presents, and identify the conditions under which acqui-hire is a more efficient path to capability than open-market recruitment.

Walmart deployed generative AI with explicit governance frameworks in advance. Compare this sequencing with the alternative pattern, in which generative AI is deployed first and governance is constructed retrospectively. Identify the trade-offs of each sequence.

END OF FREE SAMPLE

Thank you for reading this preview of

LEADING WITH ANALYTICS AND AI

A Global Boardroom Guide for the Emerging-Economy Era

This free sample has presented the cover, the front matter, and the opening material of Chapters 1 and 2, including two of the book's twenty-four substantive case studies. The complete book continues across twelve chapters organised into six parts.

The full edition develops the cumulative analytical apparatus across foundational positioning, the analytics and AI toolkit, KPI and KRI design, strategy translation, AI governance, risk management, dashboards and visualisation, communication, ESG and responsible AI, strategic decision-making, the boardroom conversation, and continuous learning. It includes the complete running case of Elastic Company across all twelve chapters, all twenty-four global and regional case studies, both authoritative templates, and 144 critical thinking questions.

THE COMPLETE EDITION

Twelve chapters · Six parts · Approximately 540 pages

Twenty-four case studies · One cumulative running case

144 critical thinking questions · Two authoritative templates

To continue reading, please purchase the complete edition.

Thanakit Ouanhlee

First Edition · 2026