AI · METAVERSE · QUANTUM

THE WORLD THAT'S BEGINNING

Navigating the AI – Metaverse – Quantum Convergence

If everything we built was for the world that's ending,
what do we build for the world that's beginning?

ANIL PANDEY

Founder, Plexus Digitals · AV Academy

Author of Reimagining Reality



BlueRose ONE .com

Stories Matter

New Delhi • London

BLUEROSE PUBLISHERS
India | U.K.

Copyright © Anil Pandey 2026

All rights reserved by author. 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 or otherwise, without the prior permission of the author. Although every precaution has been taken to verify the accuracy of the information contained herein, the publisher assumes no responsibility for any errors or omissions. No liability is assumed for damages that may result from the use of information contained within.

BlueRose Publishers takes no responsibility for any damages, losses, or liabilities that may arise from the use or misuse of the information, products, or services provided in this publication.



For permissions requests or inquiries regarding this publication, please contact:

BLUEROSE PUBLISHERS
www.BlueRoseONE.com
info@bluerosepublishers.com
+91 8882 898 898
+4407342408967

ISBN: 978-81-69822-12-1

Cover design: Shariq Ansari
Typesetting: Simran Maurya

First Edition: October 2026

DEDICATION

For my parents, who taught me that the world is always becoming.

For my wife and our children, who have been patient with the hours this book demanded, and who are the reason any of it matters.

For my students at AV Academy, whose questions are the reason it exists.

For my colleagues at Plexus Digitals, who built the practice from which these arguments are drawn.

And for the next generation of Indian builders — who will not have to ask permission to lead this decade.

EPIGRAPH

"We shape our tools, and thereafter our tools shape us."
– Marshall McLuhan

"The future is already here – it's just not very evenly distributed."
– William Gibson

"You cannot step into the same river twice,
for other waters are continually flowing on."
– Heraclitus

"Yad bhavam, tad bhavati." ("As the disposition, so the becoming.")
– traditional Sanskrit aphorism

PREFACE

What this Book is, and What it is Not

This is the book that the second chapter of my previous book, *Reimagining Reality*, pointed toward but could not contain. *Reimagining Reality* was an attempt to map the convergence of artificial intelligence, the metaverse, and quantum computing at the level of an entire technology landscape. This book takes one frame from that landscape — the claim that the unit of value in the next decade is no longer the product but the experience — and follows it as far as it will go.

The argument arrived at me, as most useful arguments do, through customers. Over the past four years, the leaders I have worked with at Plexus Digitals have stopped asking me to build them better products. They have started asking me to build them better experiences — for their customers, for their employees, for their citizens, for their patients. The wording was not coordinated. It came from a real-estate developer in Lucknow, a state utility commissioner in Andhra Pradesh, a hospital chairman in Delhi, a faith institution in Vrindavan, a defence body in the capital. Each of them, independently, in their own language, was reaching for the same thing.

I began to listen for the pattern, and the pattern was unmistakable. The technologies they were buying — AI, spatial computing,

sensors, agents, immersive interfaces — were not, in their minds, technologies. They were ingredients in something larger that they did not yet have a clean word for. The closest word was "experience," but they meant something more specific than the way the marketing world uses the term. They meant: the integrated, sensed, remembered, adaptive thing that the customer actually moves through.

This book is my attempt to give that thing the language, the structure, and the playbook it deserves. It is written for builders, leaders, investors, and policymakers who have crossed the line from interest to commitment. It is written, particularly, from India and for India, because the experience economy that is now arriving will be built disproportionately from the civilisations that have thought longest about experience as a first-class category — and India has thought about little else for three thousand years.

A note on structure. The book has thirteen parts and sixty-four chapters, organised in three movements. The first movement (Parts I through VII) sets the strategic frame: the world that is ending, the world that is beginning, the technology forces at work, the new economics of value, a Tuesday in 2030 examined hour by hour, the new product categories that do not yet exist, the post-screen design principles, and the leader's playbook for the next six months, two years, and five years. The second movement (Parts VIII through XI) takes the frame into the field: five industry re-imaginings (banking, retail, healthcare, government, defence), the investor's lens on capital allocation in the convergence era, the policy and public-interest questions only governments can answer, and four extended case studies from my own practice. The third movement (Parts XII and XIII) takes the longer view: what 2040 looks like, the civilisational question, the deeper quantum-and-human-flourishing arc that asks what abundance is for and what humans are for in an era of unbottlenecked intelligence.

A note on voice. I write in the first person throughout the body of the book, because the things I have to say are most honestly said as a builder rather than a theorist. Each chapter opens with a grounded scene — almost always drawn from my own work or my own life — and closes with three decisions I believe any serious reader should make in the next ninety days. The decisions are concrete by intent. If the reader of this book executes even ten of the one hundred and sixty-five decisions it contains, the book will have been worth its shelf.

A note on India. Most of the technology literature of the past two decades has been written from American coasts for American institutions, and much of it has been written well. But the world whose future this decade will decide is not predominantly American. The Indian vantage — and, more broadly, the vantage of the civilisations where most of humanity actually lives — has been under-represented in the corpus. *Reimagining Reality* was one small attempt to correct that imbalance. *The World That's Beginning* is another. There will need to be many more.

Finally, a note on the title. The phrase came from a question I have asked, in some form, at every leadership conversation I have led for the past two years: *if everything we built was for the world that's ending, what do we build for the world that's beginning?* The question is the spine of this book. The chapters that follow are my attempt at an answer, offered in good faith, in the company of leaders who I expect will improve on it considerably.

— Anil Pandey
Lucknow, 2026

Advance Praise

The World That's Beginning

Navigating the AI – Metaverse – Quantum Convergence

Anil Pandey

"Real estate in India is at the inflection Anil Pandey describes — from the sale of square feet to the delivery of a living experience — and most developers have not yet felt it in their balance sheets, but they will. The chapter on the home as platform is the most precise articulation I have encountered of where the category is going and what it will demand of its participants: different capital structure, different metrics, different relationship with the customer, and above all a different understanding of what the product actually is. He writes with the authority of someone who has built the technology and sat with the developer across the table. The industry would do well to listen."

— S. K. Jaggi

Chief Operating Officer, Eldeco Housing & Industries Ltd.

"Real estate has always been about more than property — it has been about aspiration, belonging, and the story a family tells about itself through where it chooses to live. What Anil has done in this book is give that instinct a rigorous commercial language for the decade ahead. The framework of the home as platform, the argument that the developer's relationship with the household should extend thirty years beyond the registration deed, and the case that experience compounds while product depreciates — these ideas will reshape how serious developers think about what they are actually building. The convergence technologies he describes are not future speculation; they are present-tense tools that the smartest operators in our industry are already deploying. This book is the clearest map I have seen of where the category is going."

— Mukesh Sayani

Real Estate Business Leader & Thought Leader

"In life sciences, the gap between what the science makes possible and what the patient actually receives has always been a design problem — a failure of experience architecture, not a failure of knowledge. Pandey's framework of the experience stack, and his argument that the unit of value has migrated from the product to the continuous relationship, maps precisely onto what the most forward-thinking healthcare and pharma organisations are now grappling with. His treatment of molecular services, AI-augmented clinical workflows, and the shift from episodic to continuous care is both technically grounded and strategically urgent. This is the rare technology book that a life sciences leader can read not as background context but as an operational guide. I recommend it to every practitioner navigating the convergence of AI, biology, and patient experience."

— Dr. Seema Nair Sayani

Author | NA Practice Lead | SME Life Sciences | CSPO®

Dallas, Texas, USA

"Anil Pandey writes the way a builder thinks — from the ground up, from real engagements, from decisions that actually had to be made. The World That's Beginning is not a forecast document; it is a field manual for leaders who have already crossed the line from interest to commitment. The Experience Stack framework cuts through years of consulting jargon in a single diagram, and the chapter on India's civilisational advantage in the experience economy is the most honest account I have read of what this country actually brings to the convergence decade. Every serious leader I know should read this book before they read another technology headline."

— Dr. Ashok Kumar Tiwari

Chief Scientist (Retd.), CSIR–Indian Institute of Chemical Technology, Hyderabad

Professor, Biological Sciences, Academy of Scientific & Innovative Research, Ghaziabad

"A long career in public administration teaches you that the world changes not in press releases but in the quiet re-platforming of the institutions that run it. Anil Pandey has written the clearest account I have seen of how the convergence of AI, spatial computing, and quantum will re-platform every institution — government and private, large and small — over the next decade. His treatment of digital public infrastructure, AI governance, and the labour transition is rigorous and humane in equal measure. This is a book that should sit on the desk of every senior civil servant and every district collector in India."

— R. C. Mishra

Retired Civil Servant, Indian Administrative Service (IAS)

"Everyone building intelligent systems eventually confronts the same uncomfortable truth: the world that is beginning does not reward better products — it rewards better experiences, better memories, better meaning. Anil Pandey has named this migration with the precision of someone who has watched it happen in real boardrooms, real clinics, real classrooms. The chapters on agentic AI and the six new product categories are the most operationally useful framing of where the technology is actually going that I have encountered in print. Required reading for anyone allocating capital or talent in this decade."

— J. Koganti

Founder, Stealth AI | Former Vice President, Estée Lauder, USA

"The deepest frameworks of our civilisation have always understood that value resides not in the object but in the encounter — in the quality of presence, of alignment, of the space between the giver and the receiver. What Anil Pandey calls the experience economy is, from a Vedic perspective, simply the rediscovery of what was always true: that it is the rasa, the felt essence, that the human being is actually seeking. What is new is that the technologies of this decade finally allow this ancient understanding to become a business model at industrial scale. This book is the bridge between that ancient knowing and the present opportunity, and India needs exactly this bridge."

— Dr. Vivek Tangri

Founder, Vedic Gurukul | Trustee, Lete Hue Hanuman Ji Mandir, Lucknow

"In education we are in the business of building people's relationship with the future — which means we had better understand what the future actually holds. The World That's Beginning is the most grounded account I have read of what the convergence decade will ask of schools, of students, and of the institutions that serve them. Anil Pandey's chapter on learning unbundled is, frankly, a more useful guide to educational transformation than most of the policy literature I have encountered. His argument that India can leapfrog the factory model of schooling, the way it leapfrogged landlines for mobile, is both correct and urgent. Every educator and school leader in this country should read it."

— Dr. Sarvesh Goyal

Entrepreneur and Business Leader

Chairman, G.D. Goenka Public School, Lucknow | The Centrum

ACKNOWLEDGEMENTS

Heartful thanks, in roughly the order they would arrive at the door if this book were a house.

This book belongs to many people more than it belongs to me. I have learned, over fifteen years of trying to build useful things in technology, that the moments when ideas finally clarify into language are almost never solitary moments. They are conversations that happened, often years earlier, with people who had no idea they were contributing. The acknowledgements below are partial. The list could fill another chapter; the omissions are mine and not theirs.

To my parents, who in a small town in north India taught me — by example more than instruction — that learning is not separable from living, that a person's word matters more than her possessions, and that the world is always becoming what its inhabitants attend to. They could not have predicted the technologies I would end up writing about. They prepared me, all the same, to think clearly when these technologies arrived.

To my wife, Anupama, who has been the quiet co-author of every consequential thing I have ever made. The hours this book required were hours subtracted from a household she was holding alone, sometimes with grace and sometimes with the kind of honest impatience that is, in hindsight, more useful than the grace. The book exists because she did. The dedication above is not a courtesy.

To my children, Isha & Aryan, who have endured a father who has, on too many evenings of the past two years, has been somewhere else in his own world. They will, I hope, find in these pages, when they are old enough, an account of why their father believed the

work mattered enough to be sometimes absent. I owe them more than the book repays.

To my brothers and sisters and the wider extended family in this world — the fabric of the small dense web of relationships that makes Indian family life what it is — who have offered, through the years, the steady ground without which no one writes anything at all. The texture of attention I describe in Chapter Sixty-Two is, in the end, modeled on the texture of attention I have received from this fabric for as long as I have been alive.

To my colleagues at Plexus Digitals — the engineers, designers, strategists, account managers, and operators who have, over the past decade, built the practice from which most of the arguments in this book are drawn. The clients in the case studies became clients because someone on this team earned their trust. The frameworks in this book were stress-tested in real engagements where their failure would have cost real money. There is no version of this book that does not depend on the work the team has done, and no version of the team's work that I would have known how to write about without the team itself. They are not named here individually only because the list would be too long; they will, I hope, recognise their fingerprints throughout the text.

To the founders, partners, and associates I have built alongside in the broader Indian technology ecosystem — at conferences, on advisory boards, in early-morning calls and late-night chats, in the patient peer relationships that the start-up world rarely names but every founder depends on. The book reflects a great deal of what I have learned from these peers. Where it advances arguments they would disagree with, the disagreement is the result of conversations they helped me have.

To the team at AV Academy — the instructors, the operations team, the curriculum designers, and the thousands of students who have, in cohort after cohort, asked the questions that have forced me to

think more carefully than I would otherwise have done. AV Academy was founded on the belief that technology education in India deserves to be at least as rigorous and as humane as anything taught anywhere in the world. The book is a continuation of that belief. The students are, in the most direct sense, its first audience.

To the clients who appear, in composite and altered form, in the case studies of Part XI and throughout the book — the chairman of the bank, the chairman of the hospital chain, the commissioner of the public utility, the leader of the faith institution, the developer who replatformed his real estate firm, and the many others whose direct words I cannot quote but whose substance I have drawn on. They have entrusted me, through the years, with conversations and decisions and reflections that no consultant earns lightly. This book is, in part, an attempt to honour that trust by extracting general lessons from the specific work without compromising the specifics. Where I have got it right, the credit is largely theirs; where I have got it wrong, the responsibility is entirely mine.

To the dozens of leaders in government, in defence, and in regulatory institutions who have, in off-the-record settings, walked me through the workings of their domains with patience and candour. The chapters on policy, on data sovereignty, on digital public infrastructure, and on defence draw on these conversations more than the citations indicate. The Indian state has done, over the past fifteen years, work in digital public infrastructure that I think is among the most consequential public-sector accomplishments of any country in this period. The leaders of that work do not always get the recognition they deserve, including in this book; I will try to do better in subsequent writing. · · ·

To the small group of friends who read drafts of this manuscript at various stages and gave me the kind of feedback that only friends can give — frank, specific, and free of professional courtesy. Several chapters were substantially restructured because of their

readings. A few chapters were saved from publication only because they pointed out, gently, that I should not have written them. The final manuscript is meaningfully better for their generosity.

To the intellectual influences whose books and papers have shaped my thinking — many of whom are listed in the bibliography in Appendix D, and many more of whom are not because I have absorbed their thinking so thoroughly that I no longer remember where my reading ends and my own thinking begins. The list of authors whose work I have leaned on includes Peter Drucker, Clayton Christensen, Joseph Pine and James Gilmore, Daniel Kahneman, Cal Newport, Donald Norman, Stuart Russell, Toby Ord, Lynda Gratton, Nandan Nilekani, Gurcharan Das, Amartya Sen, Pavan Varma, Vaclav Smil, and many others. I do not pretend to have improved on any of them. I have tried to weave their insights into a specifically Indian and specifically practitioner-oriented frame; if I have succeeded at all, the success is owed largely to them.

To the long line of teachers — formal and informal — who taught me how to read carefully, write honestly, and think for myself. The capacities these teachers gave me are the capacities the convergence decade is most likely to require. I am acutely aware that an entire generation is now being educated in conditions where these capacities are harder to acquire than they were when I was being educated. The book is, partly, a plea on behalf of these capacities, written in gratitude to the teachers who instilled them in me.

◆ ◆ ◆

To the city of Lucknow, where most of this book was written — in homes, in cafés, in long walks along the Gomti, in early-morning sessions before the household woke, in late-evening sessions after the household slept. Lucknow is, in my estimation, an underrated

city for thinking. It has the depth of an ancient civilisation, the practicality of a modern commercial centre, and the calm of a place that does not have to prove itself to anyone. Several of the chapters were, in real fact, drafted at the same desk I have used for fifteen years. The city is, in this sense, a co-author whose contribution is felt rather than seen.

To India itself — the civilisation, the country, the impossible plural noun that it is — for being the kind of place that produces the questions this book is trying to answer. I have written, throughout, from a vantage that takes Indian experience seriously as a source of insight rather than only as a market for products built elsewhere. The country deserves much more of this kind of writing than it has received, and I hope this book contributes a small amount to a larger correction that many Indian writers, builders, and scholars are now undertaking together.

◆ ◆ ◆

And finally, to the reader.

The truth about a book like this is that it is incomplete without its readers. Every chapter is an invitation for the reader to argue, to disagree, to apply, to refuse, to extend. The decisions at the end of each chapter are not commands; they are starting points. The reader who reads this book carefully, takes it seriously, and then proceeds to build something better than I have described will have done the book a greater honour than any review could.

If you have read this far, you have given the book your attention — which, by my own argument, is the most precious resource you possess in this decade. I am grateful. I hope it has been worth the time. I hope what comes after is, if not better, at least different in ways the book could not have predicted.

— Anil Pandey
Lucknow, 2026

AUTHOR'S NOTE



A few practical notes for the reader.

On voice. I have written this book in the first person throughout the body chapters. Each chapter opens with a scene, almost always drawn from my own work or my own life, and closes with a small set of decisions I believe the serious reader should consider. Where the reader disagrees with my framing, the decisions are still useful — they are the prompts a thoughtful leader should be working through regardless of whether she shares my conclusions.

On composite cases. The case studies in Part XI and the smaller examples scattered throughout the book are composites. They are drawn from real engagements but altered to protect the confidentiality of the firms and individuals involved. The substance of each case is real; the specifics — names, numbers, geographic detail — have been changed enough that no single firm or person should be identifiable. Where I have used a real public-domain example, I have done so explicitly.

On forecasts. I have made specific forecasts throughout the book — about technology trajectories, about industry structure, about timing of inflections. The forecasts represent my honest best judgement, anchored where possible in the work of researchers and forecasters whose work I trust. They are, however, forecasts. They will be wrong in specifics. The reader who treats them as precise predictions is misusing them; the reader who treats them as directional indicators of what to plan for is using them as I intend.

On the chapter numbering. The book has thirteen parts and sixty-four chapters, plus an epilogue and six appendices. The chapter

numbering is stable; the index in Appendix E is keyed to chapter numbers rather than page numbers, so the index should remain useful across editions and printings.

On the decisions. The book contains 165 decisions distributed across the chapters and gathered in Appendix B. They are intended to be acted upon. Supplement Five in Appendix F describes a practical pacing — five or six decisions per quarter, sustained over twelve quarters — that I have seen work in real leadership teams. The reader who treats the decisions as commitments rather than as suggestions will get more from the book than the reader who treats them as interesting reading.

On revisions. The convergence decade is a moving target. I expect to revise this book substantially within three years. Readers who notice errors, who disagree productively, or who can offer evidence that improves the arguments are warmly invited to write to me. The contact details are at the back of the book in the "About the Author" section. The next edition will be better for the readers who help me make it so.

TABLE OF CONTENTS

Part I: The Frame 1

1. The World That's Ending, The World That's Beginning 3
2. The Five Forces Crossing the Threshold 11
3. The Great Value Migration 20

Part II: The Forces Beneath 29

4. Where AI Actually is in 2026 31
5. From Generative to Agentic 38
6. The Metaverse, Graduated 45
7. Quantum, the Silent Revolution 52

Part III: The New Unit of Value 59

8. Why the Unit of Value has Shifted 61
9. The Three New Currencies 70
10. The Experience Stack 77
11. Memory is the Real Metric 84
12. Bharat as Experience Laboratory 91

Part IV: A Tuesday in 2030 99

13. Aarav, Meera, Two Kids, Lucknow 101
14. Work Re-Composed 108
15. Learning Unbundled 114
16. Healthcare, Continuous and Personalised 120
17. The Home That Lives 126
18. Money, Identity, Ownership 132

Part V: The New Product Categories 139

19. Six Categories That Don't Yet Exist 141

20. Sentient Spaces 146
21. Always-On Companions 152
22. Generative Worlds 158
23. Molecular Services 164
24. The Provable Truth Layer 169
25. Adaptive Everything 174

Part VI: Designing for What's Beginning 181
26. From Screens to Spatial to Neural 183
27. The Five Post-Screen Principles 190
28. The Irreducibly Human 196
29. The Devices of Tomorrow 202

Part VII: The Leader's Playbook 207
30. The Honest Conversation About Work 209
31. The First Six Months 217
32. The Twelve to Twenty-Four Months 224
33. The Three to Five Year Bets 232
34. The Risks We Must Design For 239

Part VIII: Industry Re-Imaginings 247
35. Banking and Financial Services Re-imagined 249
36. Retail and Consumer Re-imagined 257
37. Healthcare and Pharma Re-imagined 264
38. Government and Public Service Re-imagined 271
39. Defence and Strategic Sectors Re-imagined 278

Part IX: The Investor's Lens 283
40. How Capital Allocation Changes 285
41. Valuation in the Convergence Era 291
42. What to Fund, What to Avoid 297
43. India's Capital Architecture 303

Part X: Policy and the Public Interest 309
44. AI Governance — From Principles to Practice 311
45. Data Sovereignty and Digital Public Infrastructure 317
46. Labour, Reskilling, and the Social Contract 323
47. Education Policy for the Convergence Decade 329

Part XI: Case Studies from the Field 335
48. A Real Estate Developer, Re-platformed 337
49. A Hospital Chain Goes Continuous 343
50. A Public Utility Becomes Predictive 348
51. A Faith Institution Goes Digital 353
52. Patterns Across the Cases 359

Part XII: The Longer View 365
53. Beyond 2030 — What 2040 Looks Like 367
54. The Civilisational Question 374
55. A Builder's Generational Responsibility 379

Part XIII: The Quantum Horizon 385
56. Beyond Classical Limits 387
57. The Throughput Ceiling and How We Pass It 394
58. Intelligence Unbottlenecked 400
59. The Abundance Frontier 407
60. Climate, Repair, and the Quantum Toolkit 414
61. Cryptography, Trust, and the Privacy Renaissance 420
62. The Stress-Free Life 426
63. Meaning, Meditation, and the Examined Life 434
64. The Quantum-Human Synthesis 441

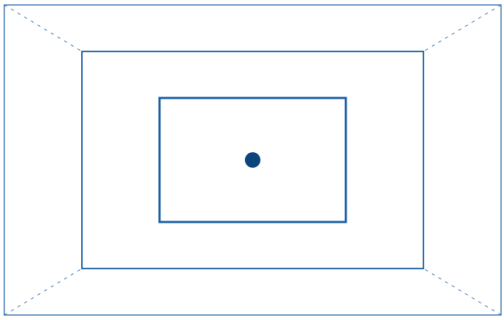
Epilogue: A Letter to the Builder Reading this in 2032 447
Appendix A: Glossary of Terms 450
Appendix B: Index of Decisions 454

Appendix C: Further Reading 470
About the Author 474
Colophon 475
Appendix D: Bibliography 476
Appendix E: General Index 485
Appendix F: Supplements 493

PART ONE

I

The Frame



Why the unit of value is shifting beneath every business in the room — and what it asks of the people who build them.

THE WORLD THAT'S BEGINNING

Why the unit of value is shifting beneath every business in the room — and what it asks of the people who build them.

CHAPTER ONE

1

The World That's Ending, The World That's Beginning

If everything we built was for the world that's ending —
what do we build for the world that's beginning?

—

"We do not see things as they are; we see them as we are."
— The Talmud

◆ ◆ ◆

In the autumn of 2025 I was invited to speak to a leadership forum in Lucknow that the leadership of one of north India's most respected real-estate companies had assembled. Three hours, no agenda, every slide debatable. The chairman who hosted me had asked, with the politeness that long careers earn, that I bring my "honest view of the decade ahead" rather than a vendor pitch. I had spent four months preparing for the conversation. I had built the slides three times. I had thrown two of them away.

What I had not thrown away was a single sentence. It sat at the top of the deck, in white type on a black field, and I had returned to it every time the rest of the deck began to lose its argument. The sentence asked: *if everything we built was for the world that's ending, what do we build for the world that's beginning?*

The sentence is the spine of this book. It is the question I now ask before I sign any client engagement, before I commit to any new product line at Plexus, before I admit a new cohort to AV Academy, before I commit to any policy submission. It is the question I am asking in writing now, at length, of any reader who has had the patience to open these pages. It is, I have come to believe, the only question that really matters this decade — because the answer to it, taken seriously, ends up reorganising almost everything that comes after it.

The trouble with the question is that it does not feel urgent at first. The world that is ending does not announce its ending. Most worlds do not. The worlds that ended within living memory — the world before the personal computer, the world before the mobile phone, the world before the broadband internet — did not formally retire. They did not issue press releases. They simply ceased to be the world inside which it was sensible to make plans. People who had not noticed continued to make plans inside them, until one day the plans no longer cashed.

This chapter is about the seam.

◆ ◆ ◆

I. The shape of an inflection

Every technology era has a shape. There is a long beginning, in which the new thing is curious and unreliable and serves no large purpose; a sharp middle, in which the new thing collapses unit costs by an order of magnitude and reaches the people who were not the early adopters; and a long plateau, in which the new thing becomes infrastructure that no one notices because it no longer needs to be noticed. The personal computer entered its sharp middle in the early 1990s. The mobile phone in the late 2000s. The broadband internet in the early 2010s. Each of these middles re-shaped the assumptions of the businesses around it, which means each ended a world.

What is unusual about 2026 is that we are not living through a single sharp middle. We are living through five at once. Artificial intelligence has crossed from curiosity into substrate; spatial and immersive computing has crossed from gimmick into infrastructure; quantum computing has crossed the long-elusive threshold of error-corrected operation; programmable biology has crossed from research-grade into clinical practice; energy and storage technology has crossed solar parity in two-thirds of the world. Each of these crossings is a once-a-generation event. The fact that all five have arrived in the same five-year window is, as far as I can tell, unprecedented in the history of industrial civilisation.

I have spent some time looking for a comparable moment. The closest analogue I can find is the period between roughly 1880 and 1910, when electrification, the internal combustion engine, the telephone, the modern factory, modern chemistry, and the modern corporation all crossed their own thresholds in the same generation. That period produced the Ford Motor Company and General Electric and AT&T and Bayer and Dow. It also produced the assembly line, the public school system, the modern city, and the mass middle class. It ended a world — the world of agricultural and craft economies — and began a world that lasted nearly a century. The mistake of every leader who lived through it was the same: they imagined the new technologies as additions to the world they already knew, rather than as the substrate of a world that was about to replace it.

We are at a comparable seam now. And the same mistake is, I observe, available to be made.

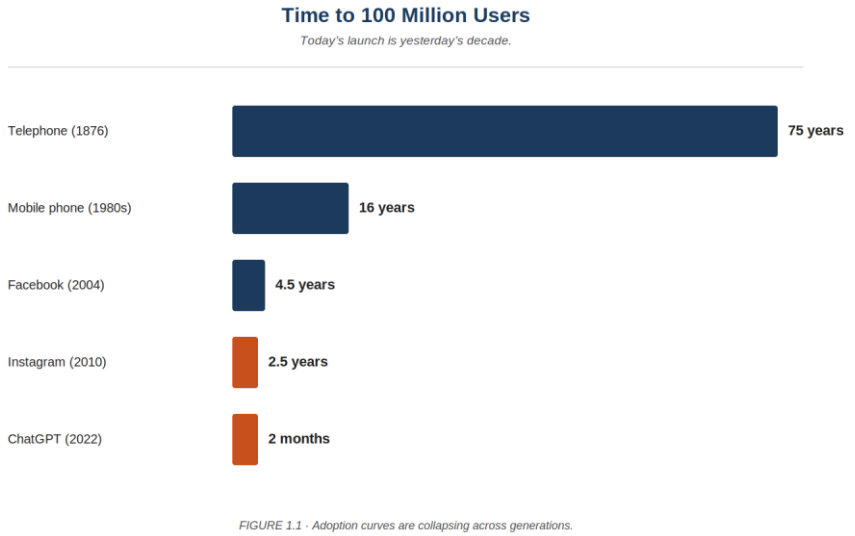


Figure 1.1

◆ ◆ ◆

II. What ending looks like from the inside

The thing about the world that is ending is that it works. That is what makes the seam so easy to miss. The car still drives, the bank still issues loans, the school still graduates the cohort, the hospital still discharges the patient, the corporation still files the quarterly report. None of this fails dramatically. It fails quietly, in the slow loss of margin, the slow drift of the customer to a competitor whose offer is somehow more right, the slow puzzlement of the product team that cannot quite say why the new entrant is winning when its features are no better.

I have watched this happen, up close, in three industries in the past two years. A retail bank that lost twelve percent of its under-thirty customers to an Indian fintech without losing a single product feature comparison. A real-estate developer that watched its premium project sell more slowly than its commodity project, because the commodity project bundled a three-year integrated wellness service that the premium project did not. A manufacturer

of high-quality industrial equipment that lost a tender to a competitor whose equipment was technically inferior but whose digital twin allowed the customer to simulate two decades of operation before signing. In each case the incumbent was selling a product. In each case the customer was buying an experience.

The seam is real. It is not a theory.

The world that is ending has two structural properties that are worth naming carefully, because they are the properties that no longer apply. The first property is that *the unit of value was the product itself*. You sold a thing. The customer took the thing home. The transaction completed. The relationship, if there was one, was warranty and recall. The second property is that *competition was on features and price*. Every product had a feature comparison sheet. The team that won was the team with the best sheet at the lowest cost. Both of these properties produced a world in which engineering, manufacturing, and supply chain were the value-creating functions of the enterprise. They were where the margin came from.

The world that is beginning has two different structural properties. The first is that *the unit of value is the experience*, by which I mean the integrated, sensed, remembered, adaptive thing the customer moves through over time. The product is one component of that experience, not the experience itself. The second is that *competition is on memory and meaning*. Memory, because what the customer recalls thirty days after the encounter, not what they evaluate at point of sale, predicts the second purchase. Meaning, because in a world of saturated choice the question the customer is asking is not "is this functional?" but "is this me?" Both of these properties produce a world in which design, intelligence, narrative, and trust are the value-creating functions. That is where the margin will come from.

◆ ◆ ◆

III. Two worlds, side by side

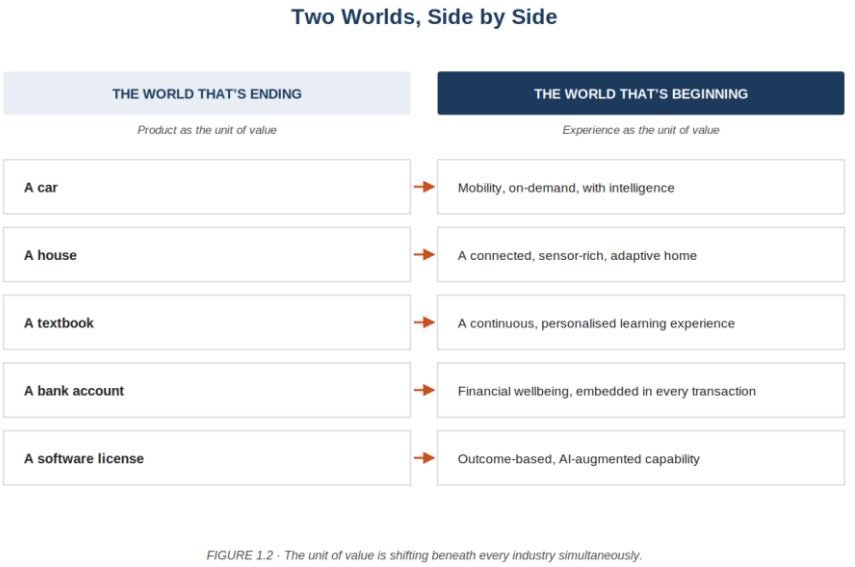


Figure 1.2

Take the figure above seriously, even though it is just a table. Each row in it is a several-billion-dollar industry that is now in the middle of a re-platforming. The car company that still believes it is selling cars in 2030 will lose the industry to the company that understands it is selling mobility, intelligence, and the continuous improvement of both over the fifteen years a customer drives. The textbook publisher that still believes it is selling textbooks in 2030 will lose the category to the company that understands it is selling a learner's growth across decades. The bank that still believes it is selling accounts will lose to the company that understands it is selling financial well-being embedded inside every other interaction in the customer's life. None of this is hypothetical. All of it is already underway, in pilot, somewhere in the world. Most of the leaders I work with are aware of one or two of these examples. Few of them are aware that the pattern is universal.

This is what makes the next decade different from any decade most of us have lived through professionally. It is not that one industry

is about to be disrupted while the others remain stable. It is that the *unit of value* is shifting beneath every industry simultaneously. The leaders who understand this as a unified pattern, rather than as a series of unrelated industry stories, will be in a position to lead. The ones who do not will spend the decade reacting to headlines about pieces of the pattern that they should have read as one.

◆ ◆ ◆

IV. The honest question

There is a question I have asked every senior team I have worked with in the past eighteen months. I will pose it to the reader now. *If you were starting your company today, with full knowledge of the technology now available and of the technology arriving in the next thirty-six months, would you build the company you currently work in?*

The honest answer, for almost every leader I have asked, is no. Not because the company they work in is a bad company. Most of them are genuinely good companies, built by good people, generating real value for real customers. But they were built for the world that is ending. Their org chart is shaped for it. Their incentives are shaped for it. Their data architecture is shaped for it. Their training programmes are shaped for it. Their procurement is shaped for it. Their suppliers are shaped for it.

The honest leader's task in this decade is not to wish their company into a different shape, which is impossible, but to begin the work of reshaping it deliberately, while the existing business continues to fund the rebuild. This is the hardest leadership task there is, because it requires you to defend the world that is ending — to keep it operating, to keep it solvent, to keep its people employed, to keep its customers served — while simultaneously building the world that is beginning, often with a different team, often with a different set of suppliers, often with a different set of metrics, and almost always against political resistance from the parts of the organisation whose careers were built inside the old shape.

This book is, in one reading, a manual for that simultaneous defence and construction. It does not pretend the work is easy. It does pretend that it is necessary, and that the time for it is now.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Name three assumptions in your business that belong to the world that's ending.**

   Take an hour with one trusted colleague. Write the assumptions down by hand. Most of them will be invisible because they have been correct for the entire careers of the people running the company. The exercise is to surface them.

2. **Identify the most uncomfortable conversation your leadership team has been avoiding about AI, and schedule it within the week.**

   Avoiding it does not protect anyone. The conversation is going to happen. It will go better if you initiate it.

3. **Block one full day this quarter to read, walk, and write about what you would build if you were starting your company today.**

   Do this alone, off-site, with no slides. The output should be a single page in your own handwriting. Bring the page to your next leadership offsite. Defend it.

CHAPTER TWO

2

The Five Forces Crossing the Threshold

Each is transformational alone. Together, they are civilisational.

—

"Compounding is the eighth wonder of the world.
He who understands it, earns it; he who doesn't, pays it."
— Attributed, Perhaps Wrongly, To Albert Einstein

◆ ◆ ◆

In April of 2024 I was on a panel at a research convening in Bengaluru, sitting next to a friend who reviews grant applications for one of India's larger philanthropic foundations. Between sessions she said something that I have repeated to other people many times since, with attribution, because it captured a thing I had been struggling to articulate. She said: "Last year, the grant proposals I read fell into clean buckets. Health proposals. Climate proposals. Education proposals. Defence proposals. This year, the proposals are arriving with the same five technologies in them — AI, sensors, spatial, biology, energy — regardless of which bucket they belong to. The buckets have not changed. The toolkit has."

The toolkit has. That is the inflection.

Every major shift in human capability has been the result of a small number of enabling technologies arriving at the same time and reinforcing each other. The agricultural revolution was the plough, the wheel, irrigation, and grain storage. The industrial revolution was the steam engine, the railway, the factory, and metallurgy. The information revolution was the transistor, the integrated circuit, the personal computer, and the internet. In each case, the individual technologies are interesting; the convergence is what makes a civilisation.

We are inside such a convergence now. Five technologies are simultaneously crossing usability thresholds within the same five-year window. This chapter walks each of them at the level a senior leader needs to make capital allocation decisions — not at the level a research scientist would. The five are: artificial intelligence, quantum computing, spatial and immersive computing, programmable biology, and energy and climate technology. Each is transformational alone. Together, they are civilisational.

◆ ◆ ◆

I. Five forces, one window

Artificial intelligence is the most advanced of the five. By the close of 2025, more than ninety percent of Fortune 500 companies and a similar share of large Indian enterprises had integrated a generative AI system into at least one workflow. By mid-2026, the leap from generative to agentic AI — from systems that answer to systems that act — was visible in production at thousands of enterprises globally. The economic implication of this shift, which I treat in detail in Chapter Five, is the largest single productivity event since the personal computer. It is not the future. It is the present, unevenly distributed.

Quantum computing is the youngest of the five but the curve is steepening. Google's Willow chip in late 2024 crossed an error-

correction threshold that physicists had pursued for thirty years. By 2026, quantum advantage on specific problem classes — molecular simulation, certain optimisation problems, certain machine-learning subroutines — has moved from theoretical to demonstrated. The commercial frontier is still three to seven years away for most categories, but the cryptography migration that quantum forces — the move to post-quantum cryptographic standards across critical infrastructure — is a present-tense problem. Any organisation holding data that must remain confidential into the 2030s is, in 2026, custodian of cryptography that may not survive the decade.

Spatial and immersive computing is the most misunderstood of the five. The brand "metaverse" collapsed in 2022 and 2023 under the weight of premature consumer marketing. The technology underneath did not collapse. It graduated. By 2026 spatial computing is in routine use across surgical training, pilot training, military simulation, virtual production for film and OTT, real-estate sales walkthroughs, industrial digital twins, and immersive learning. The devices have become wearable; the software has become serviceable; the use cases have become specific. The brand is dead. The infrastructure has arrived.

Programmable biology is the quietest of the five and the most consequential at the level of what it means to be human. CRISPR has moved from laboratory to clinic; mRNA platforms have proven themselves at planetary scale; longevity science has produced the first credible interventions; AI-designed proteins have moved from journal article to drug candidate. The cost of sequencing a genome has fallen below a hundred dollars. The cost of synthesising a custom DNA strand has fallen by a factor of a thousand in fifteen years. We are crossing the line, in this decade, between reading biology and writing it.

Energy and climate technology is the substrate beneath the other four — none of the others happen at scale without it. Solar electricity is now cheaper than fossil-fuel electricity in two-thirds of the world. Grid-scale storage has crossed the threshold at which intermittency stops being the structural barrier that solar's critics had assumed. Fusion has produced credible engineering progress in two pathways. Each AI training run consumes electricity equivalent to a small town for a week; each immersive computing experience streams gigabits to wearables; each quantum machine requires cryogenic cooling. The five forces are thermodynamically inseparable. You do not get the others without this one.

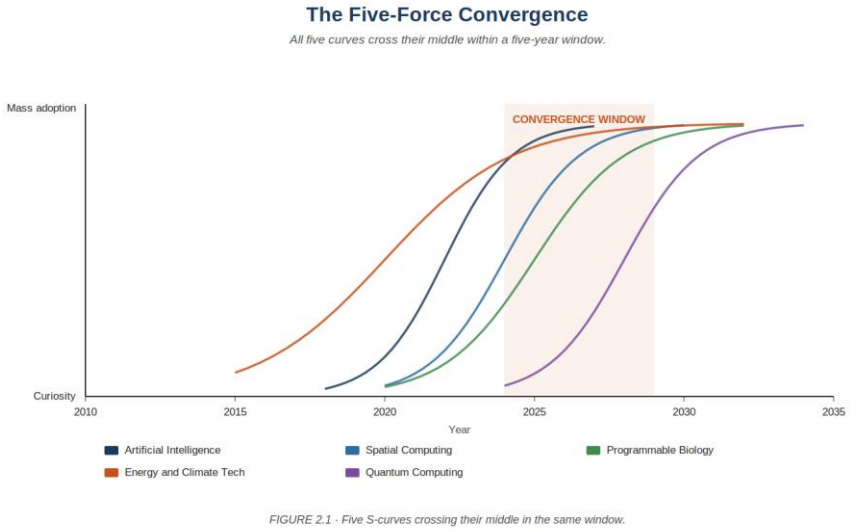


Figure 2.1

◆ ◆ ◆

II. Three macro tectonics on which the forces land

Technology does not arrive into a vacuum. It arrives into demographic, geopolitical, and climatic conditions that shape what gets built, by whom, for whom. The convergence window opens

against three macro tectonics that no leader can plan around individually but that all leaders must plan inside.

The demographic flip. The global working-age population is at a peak this decade and begins to shrink in much of the developed world thereafter. China's working-age population is contracting. Most of Europe is in shrink mode. Japan crossed the line a generation ago. Against this, India's median age is twenty-eight, and the country's demographic dividend window is roughly fifteen years. Africa's population will reach two and a half billion by 2050, half of it under twenty-five. Where the workers will be in 2040 is no longer a forecast; it is arithmetic. Where the consumers will be is the same arithmetic. Companies and capital that build for this arithmetic will own the second half of the century.

The geopolitical reset. The unipolar moment that began in 1991 ended somewhere between 2017 and 2022 — historians will argue about the date. The replacement is multipolar. Tech-bloc decoupling — chips, AI models, standards — has moved from rhetoric to industrial policy. Trade routes are redrawing themselves around the Red Sea, the Pacific, the Arctic. Capital is following alignment as much as it follows return. India sits, geographically and politically, at the intersection of every important alignment of this period, which is a strategic asset few of its leaders have fully internalised.

The climate compression. Net-zero deadlines, which were aspirational a decade ago, have hardened into capital-allocation deadlines. The energy transition will require roughly five trillion dollars of new infrastructure spend per year for two decades. Resource scarcity — water, lithium, rare earths, copper — is a planning constraint, not a footnote. Heat-stressed cities are redesigning themselves in real time. Climate adaptation is no longer a sustainability programme; it is core capital expenditure for any asset-heavy business with a fifteen-year horizon.

The technology convergence lands on top of these three tectonics. The leaders who plan as if technology can be considered separately from demographics, geopolitics, and climate are planning for a world that does not exist.

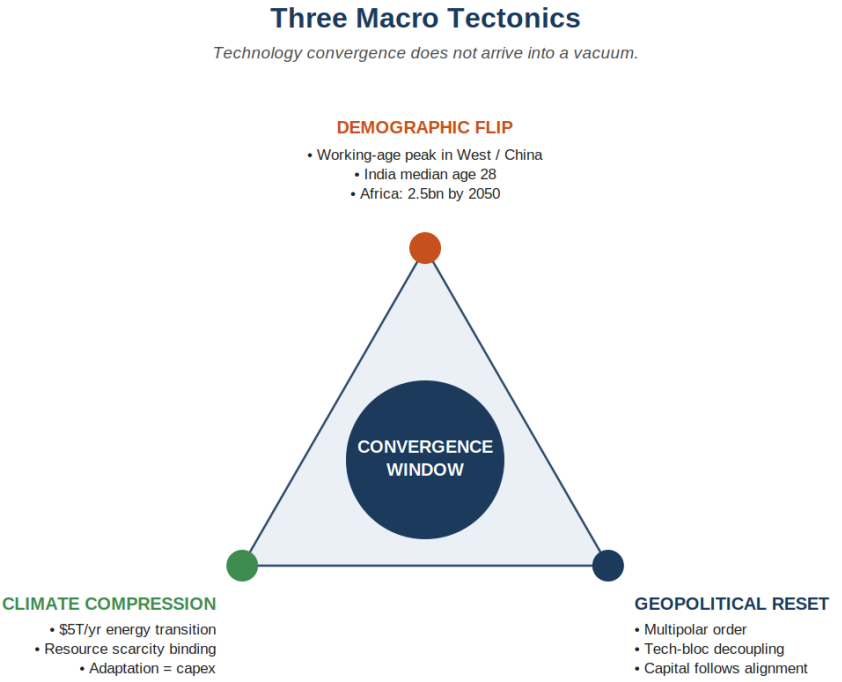


Figure 2.2

◆ ◆ ◆

III. India's structural moment

Most of what I have written so far applies globally. What follows applies particularly to India, which is the country I am writing from, the country I am writing for, and the country whose vantage on this convergence is, in my honest opinion, structurally better than that of any other large economy in the world.

India's advantages this decade are not lucky. They are structural. Begin with demographics: 1.4 billion people, the world's largest market, median age twenty-eight, a fifteen-year demographic dividend window. Add digital public infrastructure: Aadhaar, UPI, ONDC, DigiLocker, Account Aggregator, the open-network agricultural and health stacks now in roll-out. No other large economy has built a public-good digital stack of comparable depth, and most countries are now studying what India built rather than the other way around. Add talent: the largest pool of AI engineers, software developers, and design talent entering a workforce in the world. Add a frugal innovation DNA forged by three decades of building for a price-sensitive market — the country that learned to deliver complex services at one-tenth global cost is now positioned, with AI as a multiplier, to do the same to entire global industries. Add an open society and open-standards posture, and a young population fluent in English and at home with multinational technology partners.

These advantages are not enough on their own. India's risks this decade are also structural. We do not own enough compute capacity, foundation models, or AI capital. Indian data trains foreign models, and the value extracted from that training is captured elsewhere. The skills bifurcation between the top one percent of engineers, who thrive globally, and the bottom eighty percent, who may face displacement faster than re-skilling can catch them, is an active social risk. Regulatory ambiguity around data protection, AI rules, and intellectual property slows enterprise adoption. Infrastructure execution — data centres, energy capacity, fibre — is the bottleneck, not vision.

The next eighteen months will, in my view, decide whether India leads the convergence decade or merely participates in it. I do not say this to be dramatic. I say it because the leverage in such windows is non-linear, and the windows close.

◆ ◆ ◆

IV. The third currency

There is a third currency, beyond money and beyond imagination, that the convergence window requires. It is the currency of strategic patience combined with operational urgency. The two are uncomfortable to hold simultaneously. Strategic patience says: do not place the bet on next quarter's hype cycle, the technology will mature on a longer arc than the headlines suggest. Operational urgency says: do not delay foundational moves, the cost of delay is now larger than the cost of imperfect action, the headlines will be wrong but the direction of travel is real.

The leaders who navigate this decade well will be the ones who can hold both at once. They will be patient with what to build and impatient with when to start building. They will be humble about which technology will dominate which category and ruthless about beginning the institutional changes — data architecture, talent, governance, ethics — that every winning configuration will require. They will not be the leaders who shouted loudest at any single technology conference. They will be the leaders whose companies, in 2032, are quietly running on substrates that the rest of the industry is still trying to assemble.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **For each of the five forces, write one sentence describing where your industry sits on its curve.**

   Five sentences. Hand-written, unedited, one sitting. The exercise reveals, immediately, which forces you have actually thought about and which you have only nodded at.

2. **Identify which of the three macro tectonics — demographic, geopolitical, climatic — most directly affects your top three customer segments, and what your strategy is for each.**

   If you do not have a strategy for at least one of the three, you do not yet have a 2030 strategy.

3. **Commission a 90-day internal scan of all five forces from your strategy or innovation function, to be presented to the board.**

   The scan does not need to be exhaustive. It needs to be honest. Most boards have not yet seen one.

CHAPTER THREE

3

The Great Value Migration

What customers buy is shifting underneath
every business in the room.

—

"The customer never buys what you think you sell."
— Peter Drucker

◆ ◆ ◆

The CEO of an Indian premium automotive company told me, late one evening at her office in Pune, that her best customers had stopped using the word "car" to describe what they had bought from her. They used the word "drive." They used the phrase "how I move." The younger ones, increasingly, used the phrase "my mobility." She had noticed this in customer interviews three quarters before our conversation. It had unsettled her. She had asked her marketing director to investigate. The marketing director had returned with an analysis that explained the language change but missed, she felt, the more important question that the language change was pointing toward.

The more important question, she said, was this: if the customer no longer thinks she is buying a car, am I still a car company? If I am still a car company, but my customer is no longer buying a car, then somewhere between us a category has gone missing. And if a

category has gone missing between me and my best customers, what does that say about my business?

I thought about her question for several weeks. The chapters that follow this one are, in many ways, my attempt at an answer. But the foundation of the answer, the thing that has to be argued before everything else can be argued, is the migration of value itself. What customers are willing to pay for has changed. The change is not a small one. It is, in my view, the single most important commercial fact of the next decade.

This chapter is about that migration.

◆ ◆ ◆

I. The Pine and Gilmore staircase, twenty-five years late

In 1999, the business academics Joseph Pine and James Gilmore published a book called *The Experience Economy* in which they argued that economic value had been migrating, over centuries, from commodities to goods to services to experiences. Bulk coffee beans (commodity) became a can of ground coffee (good) became a cup of coffee at a café (service) became the specific experience of a cup of coffee at Starbucks, sipped in a particular chair, while a particular kind of music played, in a particular smell of espresso and wood. Pine and Gilmore were right. They were also early, by about twenty-five years.

What they correctly identified was the staircase. What they could not have predicted was how long the climb to the top step would take, and which technologies would make the top step finally affordable at industrial scale. The technologies that have arrived in the past five years — generative AI, agentic AI, spatial computing, sensor mesh, on-device inference — are the technologies that finally allow experience to be delivered to a customer at the unit cost of a service. That is the unlock. Before this unlock,

"experience" was a category reserved for the wealthy or the exceptional. After this unlock, it becomes the default.

The implication is that the traditional hierarchy of value — commodity, then good, then service, then experience — is being rewritten across every industry whose margins used to depend on the first three tiers. A car manufacturer in 2030 is not selling a car; it is selling an integrated journey of ownership across fifteen years of driving, servicing, software updates, autonomous upgrades, eventual resale, and the social signal of the brand throughout. A hotel is not selling a room; it is selling the specific emotional and cultural texture of a stay in a particular place, with a particular history, among a particular community. A hospital is not selling a procedure; it is selling the whole arc of a patient's relationship with their health, from first anxiety to final recovery and beyond.

This is what the CEO in Pune was sensing, before she had a vocabulary for it. Her customers had moved up the staircase. Her company had not. She is not unusual. Most companies in most industries are, in 2026, one step below where their best customers already are.

◆ ◆ ◆

II. Five rows of the new value table

It is helpful, when arguing this point with a leadership team, to make it concrete with five rows. I use the same five rows, more or less, in every workshop I run. They are in the figure below.

[Figure 3.1 — to be supplied by illustrator. THE GREAT VALUE MIGRATION: FROM PRODUCT TO EXPERIENCE. A two-column table. Left column header: "What you used to sell."…]

Each of the five rows is a several-billion-dollar industry that is, in 2026, in the early innings of re-platforming. Each represents a company that is currently winning its category on the left of the table, and a company that is currently winning its category on the

right. The companies on the right are smaller in 2026, by revenue. They will be larger by 2032. This is not a prediction; it is an observation about which configurations have higher gross margins, lower customer acquisition costs, higher recurring revenue, and stickier customer relationships. The arithmetic favours experience.

I want to walk each row briefly, because the abstraction matters less than the operational detail.

The car becomes mobility. The vehicle is one component. The integrated experience includes the route intelligence (where am I going, what is the traffic, what is the optimal time to leave), the in-vehicle workspace (can I take a meeting in here, can my colleagues be present in spatial AR), the service relationship (does the manufacturer know my car's wear pattern before I do), the second-life value (what is the car worth in 2034, and is that knowable), and the brand of the experience (am I a person who chooses this kind of mobility). Companies that have re-platformed across all five components have customer lifetime values four to seven times higher than companies still selling cars.

The house becomes a living environment. Real estate is, of all the categories I work on, the one where the experience migration is least understood and most lucrative. I treat this in detail in Chapter Seventeen. Suffice it to say here that the developer who continues to sell square feet in 2030 will have lost the category to the developer who has learned to sell — and to monetise on a recurring basis after the sale — the integrated experience of living in the building.

The textbook becomes a learning journey. The single most under-monetised relationship in twentieth-century commerce was the publisher's relationship with the learner. The book was sold once. The learner's growth happened entirely outside the publisher's view and outside its revenue. The technologies of the next decade

make that growth visible, addressable, and monetisable. Education companies that understand this will have the lifetime customer value of a financial services firm.

The bank account becomes embedded well-being. Banking, in the world that is ending, was a destination — the customer came to the bank's product, completed the transaction, and left. Banking, in the world that is beginning, is invisible — embedded inside every other transaction the customer makes, anticipating, recommending, protecting. The destination bank loses; the embedded bank wins.

The software license becomes capability-as-a-service. The B2B software industry has been on this migration longer than any other category — first from licenses to subscriptions, then from subscriptions to consumption-based pricing. The next step, which the agentic shift makes possible, is from consumption to outcome. The customer pays for the result the software produced, not for the software itself. This is the largest re-pricing of the enterprise software category since its invention.

Five rows, five industries, one pattern. The pattern is universal.

◆ ◆ ◆

III. The compounding problem

There is a deeper economic reason that the experience side of the table wins, not just intuitively but mathematically. It is that *experiences compound, and products depreciate*.

A product, the moment it is sold, begins to lose value. The car drives off the lot at twenty percent below its sale price. The phone is worth half its retail in eighteen months. The software, even the best software, is being written over by its successor as it ships. The economics of the product side of the table are economics of depreciation. The seller's task is to build new products fast enough to outrun the depreciation of the old ones. This is a tiring economic

posture, and it is the posture most twentieth-century companies were built for.

An experience, by contrast, compounds. The longer the customer remains inside it, the more the experience knows about her, the more it adapts to her, the more memory it accumulates with her, the more identity she invests in it, the more difficult it becomes for a competitor to reach her with a comparable offer. The economics of the experience side of the table are economics of compounding. The seller's task is to keep the customer inside the experience for as long as possible and to deepen the experience over time, both of which are positive-feedback activities. The longer the relationship runs, the more profitable each year of it becomes.

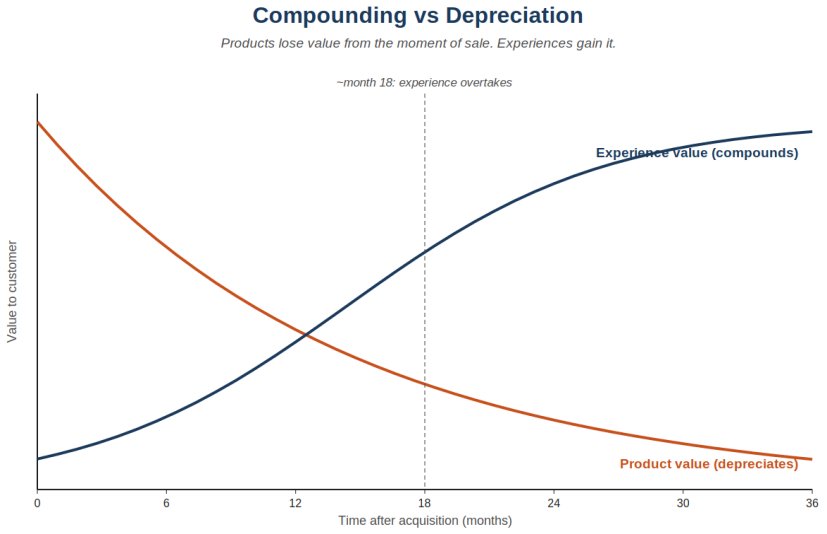


FIGURE 3.2 · After eighteen months, experience compounds faster than products depreciate.

Figure 3.2

This is the mathematical heart of why the experience side wins. It is also why the firms that re-platform earliest gain an advantage that is, in many categories, structurally hard to reverse. The customer relationships they begin compounding now will not be

available, on the same terms, to a competitor that re-platforms in 2029.

◆ ◆ ◆

IV. What this asks of leadership

It asks four things, and each of them is hard.

It asks for *category re-statement*. The CEO who continues to describe her company as a "car company" or a "publishing house" or a "construction firm" or a "hospital chain" is mis-stating the category to her own people, her own investors, and her own customers. The category re-statement is uncomfortable because it implies that some of the company's existing assets — factories, sales channels, supplier relationships — are no longer the central assets of the business. The honest leader does this anyway.

It asks for *organisational redesign*. The functions that won the world that is ending — engineering, manufacturing, supply chain, sales — were structured around the product as the unit of value. The functions that win the world that is beginning — design, intelligence, service, narrative, trust — are structured around the experience. The mismatch produces friction, and the friction does not resolve itself. Someone has to redesign the organisation deliberately. Chapter Thirty-Two treats this in detail.

It asks for *new metrics*. The leadership team that continues to run the company on units sold, gross margin per product, and quarterly revenue will not see the experience migration in its dashboards. The metrics that catch it — recurring relationship value, memory metrics, time-saved-for-customer, trust scores, experience NPS — have to be added deliberately. Most companies do not yet have them.

It asks for *capital reallocation*. The investments that built the world that is ending were investments in product creation. The investments that build the world that is beginning are investments

in data, intelligence, design, trust, and continuous service capacity. The capex/opex shape of an experience-first company looks meaningfully different from that of a product-first company, and the CFO has to be brought into the migration as a primary architect of it, not as a late-stage gatekeeper.

These four asks — category, organisation, metrics, capital — are the throughline of the rest of this book. Every chapter that follows is, in some way, an examination of one or more of them.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Write a one-paragraph description of the experience your business is actually selling, not the product.**
   If you cannot do this in one paragraph, the experience does not yet exist clearly enough in the company's mind to be sold.

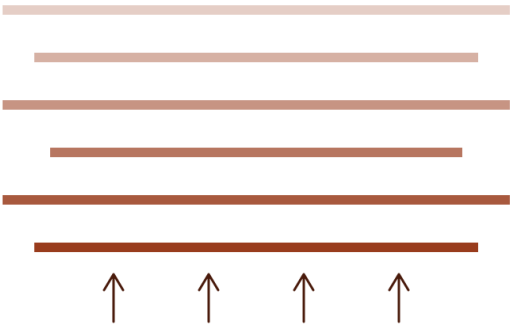
2. **Ask your top customer to describe what she buys from you.**
   Listen for whether she uses a product noun or an experience verb. The answer will tell you which side of the table she is already on.

3. **Identify one revenue line you currently book as product that should, within eighteen months, be re-platformed as experience.**
   Begin the re-platforming this quarter. The re-platform takes longer than you think it will. The window for finishing it is shorter than you assume.

PART TWO

II

The Forces Beneath



The four technology forces that make the convergence economically possible — examined only at the depth a leader needs to allocate capital and attention.

THE WORLD THAT'S BEGINNING

The four technology forces that make the convergence economically possible — examined only at the depth a leader needs to allocate capital and attention.

CHAPTER FOUR

4

Where AI Actually is in 2026

Cutting through the hype: what's real,
what's near, what's still distant.

—

"The future is already here — it's just
not very evenly distributed."

— William Gibson

◆ ◆ ◆

I sat in a board meeting last March in which two directors, both intelligent, both senior, both well-meaning, gave the company entirely opposite advice about the same technology in the same hour. The first director, who had read three articles that week and had a cousin in San Francisco, argued that the company should immediately reduce its workforce by half because AI was about to replace it. The second director, who had used a chatbot the previous Sunday and found that it confused two of his children's names, argued that AI was a glorified autocomplete and the company should ignore it for another year. Both were directors of the same company. Both were giving each other identical confident smiles across the boardroom table. The chairman looked at me.

I have come to believe that the most useful thing I can do as a builder, in a room like that, is not to give a confident answer of my own but to give the room a calibrated picture of what AI can

actually do, in production, this year — and what it cannot. The unhelpfulness of the conversation in that boardroom was symmetric. Both directors were arguing from impressions. Neither was arguing from a calibrated map.

This chapter is the calibrated map.

I will divide the territory into three horizons, because three horizons is what most leadership teams can hold in their heads at once. Real today, by which I mean visible in production at scale right now and unlikely to retreat. Emerging in the next twelve to twenty-four months, by which I mean visible in serious pilots, on credible roadmaps, with the second-generation product already shipping. Distant, by which I mean three to five years out at minimum, with significant uncertainty about timing and shape. The map is not comprehensive. It is the map I personally use when I am asked to make capital decisions or hiring decisions on behalf of my own company.

◆ ◆ ◆

I. Real today

The capabilities I list here are not aspirational. Each of them is operating, in 2026, inside enterprises I have personally worked with or audited.

Drafting, summarising, and translating at near-human quality. The single most universally absorbed AI capability in the enterprise is the production of written content — first drafts of emails, contracts, policy documents, marketing copy, internal memos, customer responses, training material, board notes. The quality is not always better than what a careful human would produce. The speed is, by a factor of ten or more. The economic effect is that knowledge work which used to consume an hour is now consuming six minutes, with the human reviewing rather than authoring. This effect is real, it is large, and it is accelerating.

Code generation across full software stacks. At the leading software firms, between thirty and forty percent of the code shipping to production in 2026 was authored, in first draft, by AI. This figure is conservative — it counts only what was directly suggested by an integrated coding assistant. It does not count the code that was rewritten by an engineer using AI assistance. The genuine fraction of AI-touched production code at top firms is meaningfully higher. The implication for software economics is the largest unit-cost shift in the industry's history.

Image, video, and voice generation indistinguishable from real media. By 2026, the threshold of "is this generated or is this captured" has been crossed for most consumer applications. The advertising industry has absorbed this completely; seventy percent of creative work at major brands now involves a generative pipeline. The film industry's high-end has absorbed it more selectively; mid-tier OTT in India absorbs it almost entirely. Voice agents capable of conducting end-to-end customer conversations in Indian languages, with appropriate accent and register, are now in production at major BFSI players.

Domain copilots in legal, medical, design, and customer service. In the law firms, the hospitals, the architecture studios, and the customer-service centres I work with, there is now, in production, a domain-trained AI that participates in every meaningful workflow. The junior associate who used to research case precedent has an AI partner that does it in seconds. The triage nurse who used to take a patient history has an AI partner that drafts it before the nurse arrives. These copilots do not replace the professional. They re-allocate her time. The professional's job becomes higher-leverage. Her output per hour multiplies. Her capacity for the unambiguously human parts of the job expands.

This is what real today looks like. It is not science fiction. It is what is on the laptop of the colleague sitting next to you at work, whether or not that colleague has admitted it.

◆ ◆ ◆

II. Emerging in the next twenty-four months

Three capability shifts are visible in pilots in 2026 and will be visible in production by 2027 or 2028. Any leadership team that does not have an explicit strategy for each of them is, in my view, unprepared.

Agentic workflows in routine operation. I treat the generative-to-agentic shift in detail in the next chapter. The short version is that AI is moving from systems that answer questions to systems that complete multi-step goals — book the trip, file the contract, reconcile the account, write the report, debug the code, run the campaign — using whatever tools and information sources it needs along the way. The pilots in 2026 are showing five-to-ten-times productivity gains on whole processes, not the thirty percent gains we became used to from generative assistance. By 2028, agentic workflows will be the dominant form of AI deployment in the enterprise.

Personalised, lifelong learning companions. The first generation of personalised learning AIs in 2024 and 2025 was largely an interactive textbook with a chat interface. The second generation, now in pilot, is meaningfully different — it has memory across years, it adapts the curriculum to the learner's interests rather than the institution's, it teaches in the learner's preferred language, and it produces credentials based on demonstrated capability rather than seat time. By 2028, every serious educational institution, including the public-school system in India, will have to make a strategic decision about whether to integrate this capability or compete against it.

AI-native interfaces beyond text and chat. The chat box was the interim interface for AI, the way the command line was the interim interface for personal computing. The native interfaces are arriving. They are voice-first, ambient, multi-modal, and tied to spatial computing. By the end of 2027, the dominant consumer AI experience will not be a text dialogue but a continuous, voice-and-vision conversation embedded in glasses, earbuds, vehicles, and home environments. I treat this in detail in Chapter Twenty-Six.

◆ ◆ ◆

III. Distant

Three capabilities sit in the distant horizon. They are likely. They are not certain. The leadership team that bets the company on them in 2026 is being foolish; the leadership team that ignores them entirely in its 2030 strategy is being equally foolish.

Artificial general intelligence at human-expert level across most cognitive tasks. The technical community is genuinely divided on whether AGI in any rigorous sense will arrive in this decade. My honest view is that the question is partly a definitional one, that something close to general competence at most professional cognitive tasks will arrive between 2028 and 2032, and that whether we choose to call this "AGI" or not will matter less than what it can do.

Embodied AI at scale. Humanoid and task-specific robots will move from research demonstrations to large-scale industrial and domestic deployment. The first credible mass-market humanoid will, in my view, ship in 2027 or 2028, but the curve from there to a billion deployed units will be a decade-long climb similar to the climb of the smartphone.

AI scientists generating original research at human-expert level. The first AI-generated papers accepted to top journals in chemistry, biology, and certain branches of mathematics are already appearing. The transition from "AI as research assistant" to "AI as research

generator" will, in my view, complete in this decade in narrow domains and in the next decade more broadly. The implications for the cost and pace of scientific progress are profound, and I treat them briefly in Chapter Thirty-Four.

[Figure 4.1 — to be supplied by illustrator. THE THREE HORIZONS OF AI CAPABILITY. A three-column table. Column 1: REAL TODAY. Column 2: EMERGING (12–24 months). Colu…]

◆ ◆ ◆

IV. The strategic implication

The capability is doubling roughly every six months at the frontier. Compounding moves slowly until it doesn't. The dominant risk to any enterprise in this decade is not over-investment in AI. It is under-imagination — the failure to forecast forward at the rate AI is actually moving and to plan accordingly.

What this asks of leadership is a different planning rhythm. The annual strategic plan, set in the first quarter and reviewed in the fourth, is no longer adequate. The strategic horizon must be re-forecast every quarter, because the capability available to the company at the start of the year and the capability available at the end of it are no longer comparable. This is uncomfortable for boards used to annual cadences. It is, I think, non-negotiable.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit one strategic plan in your organisation that was written before September 2024 — the inflection moment when agentic AI moved from labs to pilots — and identify, line by line, where it overstates AI and where it understates.**

   Most plans do both, in different sections.

2. **For each of the three horizons, name one capability your industry must have an explicit position on, and write the position in one sentence.**

   Three capabilities. Three sentences. Bring them to your next leadership meeting.

3. **Cancel one AI vendor demo this quarter and replace it with a hands-on workshop using a live model with your top fifteen leaders in the room.**

   Demos teach you about a vendor. Workshops teach you about your company.

CHAPTER FIVE

5

From Generative to Agentic

Generative AI answers. Agentic AI acts.
The economic implications are not the same.

—

"The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday's logic."

— *Peter Drucker*

◆ ◆ ◆

A finance team at a medium-sized investment management firm in Mumbai began, in September 2024, to use a generative AI to summarise the earnings calls of the companies in its portfolio. The use case was unremarkable. Every analyst on the team had, before the AI, listened to forty earnings calls a quarter and produced summaries by hand. The AI made the summaries faster and arguably better. The team welcomed the change. By Christmas, the analysts were producing twice the coverage in the same hours. By February, the head of research had quietly approved the use of an agentic system that did more than summarise — it monitored, scored, and flagged. By April, without anyone signing off explicitly, the same agentic system was, on certain low-conviction trades within strict risk limits, executing the trades itself. The first that the head of compliance learned of this was when, in late May, she looked at the trade log and found a transaction she had not

authorised, that the system explained to her, in clear English, with full justification, citing five sources.

She came to me, in some distress, to ask whether this was a violation. The honest answer was: probably not, in any specific regulatory sense, but nearly certainly a governance failure of the firm — because no one had drawn the line between summarisation, monitoring, scoring, recommending, and executing. The line, in the firm's own policy, did not yet exist. The technology had crossed it before the firm had drawn it.

The firm is not unusual. It is, in 2026, the median firm.

What had happened is that the firm had crossed, without ceremony, from generative AI into agentic AI. The shift sounds incremental. It is not. It is, in my view, the single largest economic event since the personal computer, and it deserves the chapter that follows.

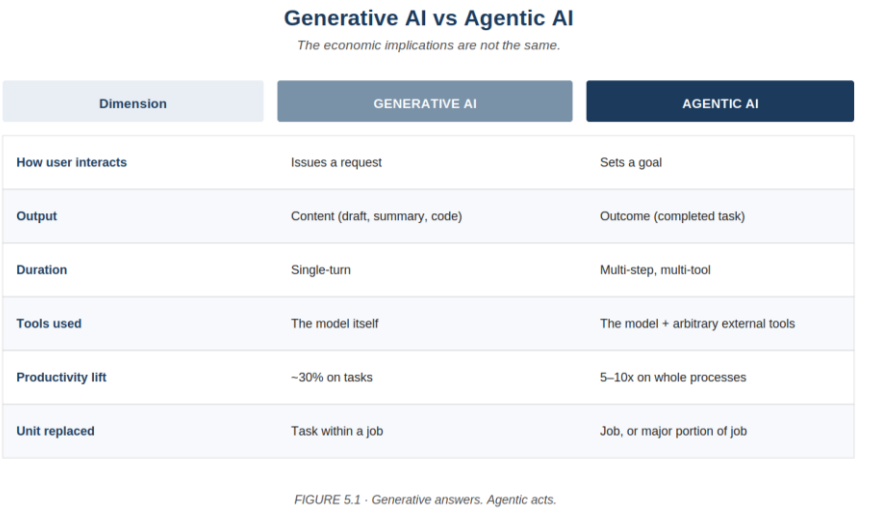
◆ ◆ ◆

I. The structural difference

Generative AI answers. The user issues a request — write this email, summarise that document, draft this code, propose those scenarios — and the AI returns content. The user then decides what to do with the content. The interaction is single-turn. The output is a document or a draft. The productivity gain is meaningful but bounded; in most enterprise tasks, it is in the range of twenty to forty percent. Critically, the AI replaces a *task* within a job. The job remains, recomposed, in the hands of the human.

Agentic AI acts. The user sets a goal — book the trip with these constraints, reconcile last quarter's accounts, run the marketing campaign for this product across these channels, file the response to the regulator before Friday — and the AI plans, executes, and reports. It uses tools. It browses websites. It calls APIs. It writes documents. It books meetings. It sends emails. It pays invoices. It

files paperwork. Where it cannot complete a step on its own, it pauses, asks the human for input, and continues. The interaction is multi-turn and multi-tool. The output is not a draft; it is a completed outcome. The productivity gain is in the range of five to ten times on whole processes — not on tasks. And critically, the AI does not replace a task within a job. It replaces, in many cases, the *job* — or large portions of the job — leaving the human to set the goal, verify the outcome, and own the accountability.



Generative AI vs Agentic AI

The economic implications are not the same.

Dimension	GENERATIVE AI	AGENTIC AI
How user interacts	Issues a request	Sets a goal
Output	Content (draft, summary, code)	Outcome (completed task)
Duration	Single-turn	Multi-step, multi-tool
Tools used	The model itself	The model + arbitrary external tools
Productivity lift	~30% on tasks	5–10x on whole processes
Unit replaced	Task within a job	Job, or major portion of job

FIGURE 5.1 · Generative answers. Agentic acts.

Figure 5.1

The difference between the two columns of that table is the largest single discontinuity that any enterprise in this decade will navigate. It is a discontinuity because the productivity gain on the right is not five times the gain on the left — it is *a different category* of gain, applied to a different unit of work, with a different set of governance, accountability, and employment implications.

◆ ◆ ◆

II. The economic implication

The generative AI economy of 2023 to 2025 was, in retrospect, a benign one for the labour market. Every job got a little easier. Every

knowledge worker could do a little more. Salaries did not collapse, hiring did not stop, the economy absorbed the productivity gain through expansion. The agentic AI economy of 2026 onward will be, in my honest view, less benign. The productivity gain is too large to be absorbed only through expansion. Some portion of it will be expressed as reduced headcount in some categories of work — not all categories, not uniformly, but in clusters where agentic systems have already crossed the threshold of reliability for whole workflows.

I want to be honest about which clusters are at most risk and which are not. The most at-risk clusters in 2026 are, in my view, the categories whose work is largely structured, repeatable, and transactional, and whose accountability is largely procedural rather than fiduciary. First-pass legal review of standard contracts. Tier-one customer support for known issues. Routine bookkeeping. Standard market research. Routine compliance reporting. Junior coding for well-specified features. Routine financial analysis. The work in these categories is not gone, but the headcount required to do it is meaningfully less than it was, and the gap is widening.

The least at-risk clusters are the categories whose work involves genuine ambiguity, fiduciary accountability, contextual creativity, or empathic presence. Senior negotiation. Patient-facing care. Master-craft work in any field. Brand stewardship. Strategic synthesis under uncertainty. Final-mile sales of high-trust products. Crisis response. Mentoring and teaching. These categories do not shrink. In many cases they grow, because as the surrounding categories shrink, the human attention available for the irreducibly human work expands.

The middle, between these two clusters, is where most of the workforce sits, and where the leader's responsibility is greatest. It is the responsibility I treat in detail in Chapter Thirty.

◆ ◆ ◆

III. The governance implication

Most enterprise governance in 2026 was written for the generative AI era. The policies say, in effect: the AI may suggest, but a human must approve. The policies are silent on what happens when the AI does not merely suggest, but acts. The Mumbai investment firm I described at the start of this chapter discovered the gap by accident. Most enterprises will discover it the same way, unless they redraw the line deliberately.

The line that needs to be drawn is the line between recommendation and action — the line between *the AI suggests a trade and the human executes* and *the AI executes the trade and the human reviews*. The line is not drawn by technology. It is drawn by policy. And the policy has to specify the line at multiple levels: which decisions the AI may execute autonomously, within what risk limits, with what audit trail, with what human review cadence, with what escalation, with what failure modes, with what accountability ownership. None of this is exotic. All of it is non-negotiable for any organisation that intends to deploy agentic systems in the next twenty-four months.

The leaders who do this work in 2026 will, in 2028, be operating agentic systems with confidence. The leaders who do not will, in 2028, either be exposed to incidents they could have prevented, or be paralysed by a justified fear of the systems they have refused to govern.

◆ ◆ ◆

IV. The strategic implication

The strategic question that the agentic shift forces every leader to answer, sooner or later, is this: for which processes in my company is the human currently doing work that an agent could do reliably, and what should I do with that fact?

The wrong answer is to ignore it, because the agent will arrive whether the leader is ready or not. The wrong answer is also to fire the humans, both because the displacement is socially destructive and because the humans contain knowledge the agent does not. The right answer, in almost every case, is to redesign the process — to allow the agent to do the work that is structured and repeatable, to elevate the humans into the work that is contextual and consequential, and to capture the productivity gain in the form of expanded service, expanded coverage, expanded quality, or expanded margin.

[Figure 5.2 — to be supplied by illustrator. THE PRODUCTIVITY LIFT CURVES. A line chart with two curves over time. The blue curve, "Generative AI Lift," asymptotes a…]

This redesign is the work of the next two years for every serious enterprise. The leaders who lead it, rather than reacting to it, will set the terms on which their organisations participate in the convergence economy that follows. The leaders who do not will find that their participation is set by the agents themselves, and by the smaller, faster, more agentic firms that overtake them.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify one repetitive, internal, end-to-end process in your business and replace it with an agentic workflow as a pilot.**

   Measure end-to-end time and cost before and after. Publish the result internally. The pilot is the foundation for every subsequent agentic deployment.

2. **Define your organisation's policy on where humans must remain in the loop, in writing, before you scale agentic deployments.**

   The policy must specify the line between recommendation and action, the risk limits, the audit trail, and the accountability ownership. If you cannot write the policy, you are not ready to deploy.

3. **Begin charging customers, in at least one product line, for outcomes rather than effort.**

   Outcome-based pricing is the commercial expression of agentic capability. The line items the customer pays for change. The relationship deepens. The margin compounds.

CHAPTER SIX

6

The Metaverse, Graduated

The brand died. The technology graduated.
Spatial computing is now infrastructure.

—

"We shape our tools and thereafter our tools shape us."
— Marshall McLuhan, often quoted, often correctly

◆ ◆ ◆

A pilot trainer in Hyderabad told me, in 2025, that he had not used a physical flight simulator for any new pilot trainee in eighteen months. He was not telling me this because he thought it was remarkable. He was telling me this because I had asked him, over coffee, how the new generation of pilots was different from the one he had himself come up through. He looked at the question, thought for a moment, and said: "They expect their training to be in VR. The physical simulator feels old to them, the way the analogue cockpit feels old to us." He did not realise, until I pressed him, that what he had just described was the quiet completion of a transformation that was, three years earlier, the loudest and most over-promised theme in technology marketing.

The metaverse, the public and the press were told between 2021 and 2023, would change everything. Then, between 2022 and 2024, the metaverse — the public and the press were told — was dead. The brand collapsed under the weight of its own marketing

budget. Meta lost forty billion dollars on Horizon Worlds. The articles that had announced the metaverse's arrival began to announce its demise. Headline editors moved on. Most leaders moved on with them.

The technology, meanwhile, did not move on. It did the unglamorous thing. It went to work. The pilot trainer in Hyderabad is one of millions of professionals around the world who, by 2026, are using spatial computing every day for serious purposes. They do not call what they are doing "the metaverse," because the word has acquired a residue of failed marketing. But what they are doing is what the metaverse was always supposed to be: presence at distance, immersion in environments that do not yet physically exist or cannot yet physically be visited, and digital embodiment of work, training, design, and play.

The metaverse did not die. It graduated. This chapter is about where it actually lives now, and why it matters more than the headlines suggest.

◆ ◆ ◆

I. The misconception

Three things died between 2022 and 2024, and they are easy to conflate. The brand "metaverse" died, because Meta's specific bet on Horizon Worlds did not produce the consumer adoption Meta had projected, and the financial press, justifiably, made an example of the company's losses. The hype cycle around consumer VR died, because the headsets of 2021 were too heavy, the content was too thin, and the proposition that consumers would spend hours a day in a virtual social space turned out to be wrong, at least in the form Meta had imagined. And the over-promised use cases — virtual real estate speculation, NFT-driven communities, virtual concerts as cultural mainstream — died, because they were always, in retrospect, over-promised.

What did *not* die, but was conflated with what died, was the underlying technology stack. Spatial computing devices continued to ship. The hardware improved. The software ecosystems matured. The use cases that had genuine economic logic — training, simulation, design, virtual production, digital twins, immersive collaboration, immersive learning — continued to grow, often faster after the brand collapsed than before. By 2026, the install base of spatial computing devices for professional use crossed one hundred million units globally. The brand died. The infrastructure arrived.

The wrong question, in 2026, is "will the metaverse arrive?" The metaverse has arrived. It does not look like the picture that Meta painted in 2021, which is exactly why most observers think it has not arrived. It looks like infrastructure, which is to say, it looks like the parts of the world that no one notices because they are working.

◆ ◆ ◆

II. Where it actually lives now

Six domains, as far as I can tell, are where spatial computing is now infrastructure rather than experiment. Each of them is a several-billion-dollar category. None of them appears in the consumer headlines that declared the metaverse dead.

Spatial computing devices in wide professional use. Apple's Vision Pro shipped in 2024 and the second generation in 2025; Meta's Quest line continues to evolve; AR glasses from a half-dozen credible vendors are now shipping at consumer price points. The consumer mass market is still ahead. The professional mass market has arrived. By 2027, more than a hundred million units will be in the field globally.

Virtual production for film and OTT. Eighty percent of premium filmed content in 2026 — *The Mandalorian* is the example everyone knows; the truth is most major Indian OTT productions of the past

two years — uses immersive virtual sets rather than location shooting or traditional green-screen. The economics are decisive: the cost is lower, the iteration is faster, the creative range is wider.

Industrial digital twins. Every major asset-heavy industry — manufacturing, real estate, energy, transportation, mining, ports — now operates digital twins of its physical assets. The twin allows the operator to simulate, predict, optimise, and train without disturbing the physical asset. The category was largely a research curiosity in 2018. By 2026 it is standard practice.

Training and simulation across high-stakes professions. Pilots, surgeons, soldiers, plant operators, emergency responders, and increasingly factory floor workers are now trained primarily in spatial environments before they touch the physical equipment. The pilot trainer in Hyderabad is the rule, not the exception. The cost reduction over physical simulation is between five-fold and fifty-fold. The skill transfer is, in most studies, equivalent or better.

Immersive learning in K-12 and higher education. History taught in walkable historical environments. Biology taught inside the cell. Chemistry taught at the level of molecular interaction. The early adopters of immersive learning in Indian education — both public and private — are reporting comprehension and retention gains that are difficult to ignore. This category will, in my view, be one of the largest of the coming decade, and I treat it in detail in Chapter Fifteen.

Live events, immersive shopping, and tele-presence at scale. Virtual concerts that draw genuine paying audiences. AR product try-ons that demonstrably reduce returns. Tele-presence meetings that, for the first time, feel meaningfully more like being in a room than a video call. These are smaller than the previous five categories, but they are growing, they are profitable in pockets, and they are the on-ramp to consumer adoption.

[Figure 6.1 — to be supplied by illustrator. WHERE THE METAVERSE ACTUALLY LIVES NOW. A six-cell grid, one cell per domain (devices, virtual production, digital twins…]

◆ ◆ ◆

III. The detached hype cycle

There is a useful way to look at this. Every major technology of the past forty years has had the same shape: a peak of inflated expectations, a trough of disillusionment, a slope of enlightenment, a plateau of productivity. Gartner has charted this for decades. What is unusual about the metaverse is that the brand and the technology decoupled during the trough. The brand did not climb the slope of enlightenment. The technology did, alone, while the brand remained in the trough as a cautionary tale.

This decoupling has made the technology under-priced as an investment thesis and under-discussed as a strategic priority for the past three years. Both, in my view, are now beginning to correct. The investors who placed quiet bets on industrial spatial computing in 2023 and 2024 are now collecting returns that the consumer-VR investors of 2021 never saw. The leaders who ignored the headlines about the metaverse's death and continued to build digital twins and immersive training programmes are now operating capabilities their competitors will spend the rest of the decade trying to replicate.

[Figure 6.2 — to be supplied by illustrator. THE HYPE CYCLE, DETACHED. The standard Gartner-style hype-cycle curve. The "brand" curve falls into the trough and stays…]

◆ ◆ ◆

IV. Why this matters for the convergence economy

Spatial computing matters for the argument of this book because it is the layer at which experience finally becomes physical.

Without spatial computing, "experience" remains a metaphor for service excellence; the customer's senses are not actually engaged in any new way. With spatial computing, experience becomes literal. The customer is somewhere. She is moving through a place. She is seeing what was previously invisible. She is collaborating with people who are not physically with her. She is learning by doing rather than by reading. She is rehearsing the future before she enters it.

Every category in the value-migration table I drew in Chapter Three becomes meaningfully more powerful with a spatial computing layer. The car becomes a moving spatial workspace. The house becomes a sensor-rich, navigable digital twin that the family can rehearse before moving in. The textbook becomes a place the learner walks through. The bank account becomes a financial environment the customer inhabits rather than a transactional list. The software license becomes an immersive operating environment for a team.

This is not science fiction. Each of these is in pilot, somewhere, in 2026. The strategic question for the leader is whether to be among the first to deploy in her category, or among the firms that learn the medium five years late from the firms that did.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Visit one operational deployment of spatial computing in a non-tech industry within the next thirty days.**

   A factory, a hospital, a training centre, a real-estate sales office, an OTT production set. The visit will be more useful than ten conferences.

2. **Identify one customer-facing or operational asset that should now be mirrored in a digital twin, and commission the twin within the year.**

   The first twin is the hardest. After it, the rest are projects rather than experiments.

3. **Add "spatial" to your product roadmap as a first-class consideration, not an experimental footnote.**

   Every roadmap review should ask, of every product, what its spatial expression is or will be. Most products will have one, sooner than their teams expect.

CHAPTER SEVEN

7

Quantum, the Silent Revolution

Quantum is where AI was in 2018 — almost ready, mostly invisible, soon non-negotiable.

—

"Anybody who is not shocked by quantum mechanics has not understood it."

— Niels Bohr

◆ ◆ ◆

In April 2025 the chief information security officer of one of India's larger private banks called me, after a board meeting at which she had been asked, for the first time, about her institution's quantum readiness. The question had come from a director who had read a recent article about Google's Willow chip and the United States National Institute of Standards and Technology's post-quantum cryptography standards. The director had asked, reasonably, whether the bank's customer data — twenty-five years of it, accumulated and encrypted and stored — was safe against decryption by a quantum computer. The CISO, to her professional credit, had answered honestly. She had said: I do not yet know.

She came to me to find out. The conversation we had was not, as she had expected, a conversation about quantum physics. It was a conversation about regulatory horizons, encryption inventory, data lifetime obligations, vendor roadmaps, internal capability

building, and the cost of a transition that, depending on the scope, would be either a discrete project or a multi-year programme. The question was not "is quantum a threat to my data today?" The question was "is quantum a threat to my data within the lifetime obligations of the data, and if so, what is my migration path?"

The answer, for her bank, was yes — and the migration path was a three-year programme. The answer, for most banks I have since spoken to, is the same. The answer, for any organisation that holds confidential data with a confidentiality obligation extending into the 2030s, is the same. Quantum is a present-tense problem. The fact that quantum advantage is still a few years away does not change the fact that the migration to quantum-secure cryptography must begin now, because the migration takes years, and because adversaries can collect encrypted data today and decrypt it later.

This chapter is about what a leader actually needs to know about quantum in 2026 — without the physics — to make good decisions.

◆ ◆ ◆

I. What changes when bits become qubits

A classical computer represents information as bits, each of which is either zero or one. A quantum computer represents information as qubits, each of which can be in a superposition of zero and one simultaneously, and can be entangled with other qubits across space such that measuring one instantaneously determines the other. The intuition that needs to be carried forward from this paragraph is not the physics. It is the consequence: quantum computers can, on certain problem classes, explore an exponentially larger solution space in a given time than classical computers can. They are not faster at everything. They are dramatically faster at specific things.

The specific things are, broadly, six. Each of them matters for at least one industry that the reader of this book may run.

Drug discovery and molecular simulation. The fundamental task in drug design is simulating how molecules interact. Classical computers can simulate small molecules well and large molecules badly. Quantum computers, in principle and increasingly in practice, can simulate large molecules accurately. The implication for pharmaceutical R&D timelines is, by some estimates, a hundred-fold acceleration in the front end of drug discovery.

Optimisation problems at scale. Logistics, supply chain, traffic management, financial portfolio construction, energy grid balancing — each of these is, mathematically, an optimisation problem. Classical methods solve them approximately. Quantum approximate optimisation algorithms, today, are demonstrating advantages on certain instances. By the late 2020s, this category will be one of the largest commercial applications of quantum computing.

Materials and energy science. The design of new battery chemistries, new superconductors, new catalysts for fusion containment, and new carbon capture materials is, fundamentally, a materials simulation problem of the same shape as drug discovery. Quantum advantage here means new materials reaching market faster, with better properties, at lower cost.

Cryptography. The most discussed quantum implication, and the one that has cryptographic policy moving fastest, is that a sufficiently capable quantum computer will break the public-key cryptography that secures most internet traffic today. The post-quantum cryptography standards finalised by NIST in 2024 are now being rolled out across critical infrastructure globally. The migration is non-trivial. Every organisation with confidential data outliving the next decade must own this transition.

AI acceleration. The genuine fusion of the convergence is here. Quantum machine learning algorithms — still early — promise to train models that classical compute cannot reach, in particular for

reinforcement learning and certain classes of pattern recognition. The convergence of quantum and AI is, in my honest view, the single most consequential technological frontier of the 2030s.

Financial modelling. Risk simulation, derivatives pricing, fraud detection — banks have been investing in quantum capability since 2019, and by 2026 most large banks have at minimum a research relationship with a quantum hardware provider. The early commercial advantage will arrive in narrow applications first; the broader transformation is a 2030s story.

◆ ◆ ◆

II. The cryptography migration that cannot wait

Of the six domains above, five are opportunity-shaped and one is risk-shaped. The risk-shaped one — cryptography — is the one that requires action in 2026, regardless of the leader's interest in quantum more broadly.

The reason is structural. An adversary capable of collecting encrypted data today can store it and decrypt it later, when sufficient quantum capability becomes available. This means that any data with a confidentiality obligation extending past, conservatively, 2030, is at risk *now*, even though the decryption itself is not yet possible. The data is being collected. The decryption is on a timer. The migration to post-quantum cryptography is the only defence, and it is a multi-year programme for any large enterprise.

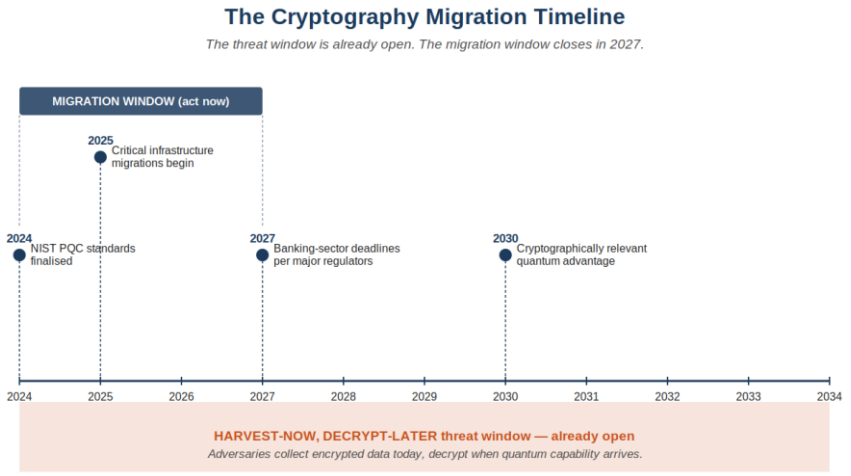


Figure 7.1

The migration is technically straightforward — replace existing public-key algorithms with NIST-approved post-quantum alternatives — and operationally extensive. The leader's task is not to understand the cryptography. The leader's task is to commission the inventory, allocate the budget, set the deadlines, and audit the progress. The CISO can manage the rest.

◆ ◆ ◆

III. What to do about the other five

For the five opportunity-shaped domains, the leader's posture in 2026 should be neither aggressive nor passive. It should be exploratory, with deliberate learning investments and at least one credible partnership.

The aggressive posture — committing significant capital to quantum-specific applications today — is, in my view, premature for most enterprises. The hardware is not yet at commercial-advantage scale for most categories. The talent is scarce and expensive. The applications are still being defined. The companies

that bet large on quantum-specific products in 2024 and 2025, with rare exceptions, are now revising downward.

The passive posture — ignoring quantum until commercial advantage arrives — is also wrong, because the lead time on capability building, partnership development, and use-case definition is too long to be compressed into the eighteen months between commercial advantage and competitive necessity. By the time the headline arrives, the firms that started in 2026 will be three years ahead of the firms that started after the headline.

The exploratory posture is what I recommend. Commission a one-page quantum-readiness assessment. Identify, for each business unit, the one optimisation, simulation, or cryptography problem most likely to benefit from quantum. Build a small internal capability — typically two to four people — that can engage with vendors, evaluate pilots, and translate between the physics and the business. Fund one credible partnership with a quantum hardware or software provider, sized to give your team genuine access without committing to a specific vendor too early. Track the field quarterly. Be prepared to scale aggressively when the threshold is crossed in your category.

◆ ◆ ◆

IV. The strategic horizon

Quantum is, in my view, where AI was in 2018. The technology was real, the inflection was visible to specialists, the commercial applications were narrow and contested, and the leaders who paid attention then were positioned to lead by 2023. The leaders who ignored AI in 2018 spent 2023 to 2025 trying to catch up, often expensively.

I do not know exactly when quantum's 2023 will arrive. My honest forecast is somewhere between 2028 and 2032 for the first category-defining commercial advantages, with cryptography running on

its own faster timeline. The implication for leaders in 2026 is not to bet the company on quantum today. The implication is to ensure that, when the quantum equivalent of 2023 arrives in your industry, your company is positioned to participate rather than to scramble.

[Figure 7.2 — to be supplied by illustrator. SIX DOMAINS WHERE QUANTUM MATTERS FIRST. A six-cell grid: drug discovery, optimisation, materials, cryptography, AI acce…]

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Commission a one-page quantum-readiness assessment for your organisation within the quarter.**

   The assessment should cover three things: cryptographic exposure, opportunity domains, capability gaps. The page is the start, not the end.

2. **Begin migrating critical cryptographic infrastructure to post-quantum standards now, not when quantum arrives.**

   The harvest-now-decrypt-later threat is already operational. The migration is a programme, not a project.

3. **Identify one optimisation problem you currently solve poorly that may be a candidate for early quantum advantage in your industry.**

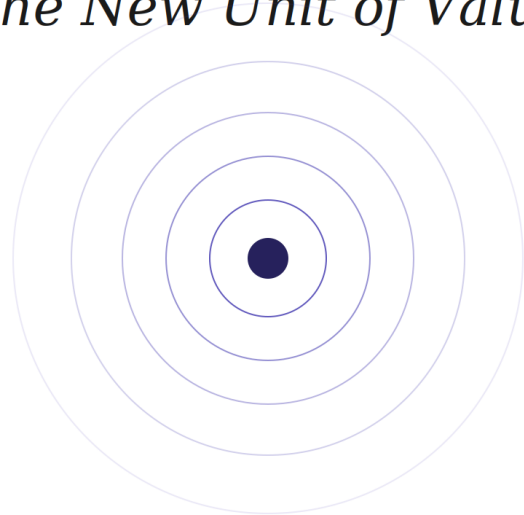
   Begin a conversation with a credible quantum partner about it. The conversation costs nothing and positions you to move quickly when the technology crosses threshold in your category.

PART THREE

III

PRODUCT → EXPERIENCE

The New Unit of Value



The heart of the book. Why experience eats product, what the new currencies are, how the experience stack actually works, and why India has a thirty-year head-start.

THE WORLD THAT'S BEGINNING

The heart of the book. Why experience eats product, what the new currencies are, how the experience stack actually works, and why India has a thirty-year head-start.

CHAPTER EIGHT

8

Why the Unit of Value has Shifted

These are not opinions. They are arithmetic.

—

"The customer's perception is your reality."

— *Kate Zabriskie*

◆ ◆ ◆

I once spent an afternoon in a conference room with the senior leadership of a consumer electronics brand that had, over twenty years, built a respected portfolio of audio products. They had asked me, more or less, to explain why their newest product — which was, on every spec sheet, the best they had ever made — was struggling against a competitor whose product was, on every spec sheet, technically inferior. I asked them to bring two units to the meeting, theirs and the competitor's, both unboxed.

The competitor's box was an experience. The packaging took a minute and a half to open with care. Inside, the speaker was wrapped in cloth, accompanied by a hand-signed thank-you card from the founder, a small story-book about the artisans whose acoustic theory had inspired the design, and a QR code linking to the playlist the founder listened to while testing the speaker. The speaker itself was beautiful, but no more beautiful than my client's. It paired with the phone in three seconds. It greeted the user by the

name on the phone. It asked, in a calm voice, whether to play the founder's playlist or to learn from her listening habits.

My client's box was a box. It contained the speaker, a power cable, a manual in eight languages, and a warranty card.

My client had spent four years engineering a better speaker. The competitor had spent four years engineering a better experience. The market was making its choice, and I did not need to add anything to the analysis the room was already, with mounting discomfort, making for itself.

I begin this chapter with that scene because it is the cleanest illustration I know of the central argument: experience eats product. Not because experience is more emotional, more romantic, or more fashionable. Because the underlying economics — five distinct forces — make it so. This chapter walks each of the five.

◆ ◆ ◆

I. The marginal cost of products is collapsing

The first force is the simplest, and it has been quietly true for two decades. The marginal cost of producing a physical good — and even more so, a digital good — has been collapsing for a long time. Automation has compressed it. Global supply chains have compressed it further. AI-driven design and additive manufacturing have, in 2026, brought it to a level that would have been inconceivable thirty years ago.

The implication is that *what* you produce is becoming a smaller share of the value the customer pays for. The competitive frontier has migrated up the stack. If the marginal cost of the speaker itself is falling toward zero — and it is, slowly, in real terms — then what the customer is paying for is increasingly the things that surround the speaker: the curation, the brand, the design language, the service relationship, the integration into her life. These things are

not on the spec sheet. They are not the product. They are the experience.

The same arithmetic operates in software, where the marginal cost of code is now near-zero and falling further as AI generates increasing fractions of new code. It operates in financial products, where the marginal cost of executing a transaction is asymptotic to nothing. It operates in education, where the marginal cost of delivering content is the cost of the bandwidth. In every category where the marginal cost of the thing itself is collapsing, value migrates inevitably to the layer above the thing — and that layer is experience.

This first force is, in some ways, the foundational one. Without it, experience would be an aesthetic preference. With it, experience is an economic necessity.

◆ ◆ ◆

II. Customer choice has become overwhelming

The second force is psychological, and it has accelerated in the past five years. The customer in 2026 has access to more choice in every product category than any consumer in human history. Twenty thousand audio brands. Ten thousand restaurants on a delivery app. Three million books. Two hundred streaming services. Five hundred banks. Forty operating systems and counting.

The decision research literature is unambiguous on what this does to the customer. Above a threshold, additional choice ceases to add value and begins to subtract it. The customer is exhausted by the comparison. The cognitive cost of choosing exceeds the marginal benefit of the better choice. The customer reaches for whatever feels easiest, most curated, most recommended by a source she trusts, most aligned with the identity she is constructing. She does not, in most cases, conduct a feature-by-feature comparison even when one is available to her.

This is, for the firms that have spent twenty years optimising for feature-by-feature comparison, a difficult fact to absorb. It says, in effect, that the company that produces the best feature-by-feature product may not be the company that wins the customer. The company that wins the customer is the company that has reduced her cognitive load — that has curated for her, recommended for her, designed for her, anticipated for her, simplified for her. That curation is, again, experience.

This is why brands that look, on the spec sheet, redundant — the third premium grocery delivery, the fifth high-end audio brand, the eighth meditation app — can build durable businesses. They are not competing on the spec sheet. They are competing on the experience of choosing them, of using them, of being the kind of person who chose them.

◆ ◆ ◆

III. AI now enables personalisation at scale

The third force is the technological unlock. Personalisation, in the world that is ending, was a privilege of the wealthy. The bespoke tailor, the personal banker, the family physician, the concierge — all were forms of personalisation that operated at a unit cost far higher than the median customer could afford. The result was that the median customer received a generic version of every service: the suit that nearly fit, the savings account that nearly suited, the medical visit that nearly addressed her actual concern.

Generative AI, used well, collapses this cost asymmetry. The unit cost of producing a personalised experience — a custom learning programme for a specific student, a custom financial plan for a specific household, a custom care pathway for a specific patient, a custom shopping experience for a specific customer — is now, with AI, within an order of magnitude of the cost of producing a generic one. By 2028, in many categories, it will be lower.

The implication is that the entire generic middle of every service category is structurally vulnerable. The category that used to offer a generic product to most customers and a personalised product to the wealthy will, by the end of this decade, offer a personalised product to most customers and an obsolete generic product to no one. The companies that re-platform around personalisation early will absorb the migration. The companies that retain the generic product as their default will lose share continuously as customers migrate to the personalised competitor.

This third force is the one that converts experience from an option into a structural inevitability. Without it, experience would be an upmarket strategy. With it, experience is the new default for the entire market.

◆ ◆ ◆

IV. Networks compound experiences while products depreciate

The fourth force I introduced in Chapter Three, but it is worth restating in its proper sequence. Products depreciate. Experiences compound. The longer the customer remains inside an experience that knows her, adapts to her, accumulates memory with her, and integrates her preferences into the next iteration, the higher the switching cost becomes — not because the company has erected an artificial barrier, but because the experience itself has become more valuable to the customer than a competitor's experience could be on day one of switching.

This compounding is asymmetric. The first month of the customer's relationship with an experience produces a small amount of compounding. The third year produces a great deal of it. The new entrant who tries to win the customer in year three is not competing against a product spec; she is competing against three years of accumulated context and habit that the incumbent has and she does not. The economics of customer acquisition shift.

The economics of customer retention shift. The economics of the entire category shift.

This is why platform businesses — once they have crossed the threshold of compounding — are so structurally durable. It is also why the experience-first companies of the next decade will look more like platforms than like brands, and why the competitive moats of 2032 will be made of memory and adaptation rather than features and price.

◆ ◆ ◆

V. Customers demand meaning, not just function

The fifth force is cultural, and it is the most contested. The argument is that the post-utility consumer asks not "is this functional?" — that question has been answered for most product categories — but "why should I buy this?" — and the answer must be emotional, identity-based, story-driven, meaningful in the literal sense.

I have heard this argument resisted on the grounds that it is a luxury phenomenon, applicable to consumers wealthy enough to choose on grounds beyond function. The data, in my experience, does not support this resistance. The phenomenon appears across price points and across markets. The customer in tier-three urban India making a thousand-rupee purchase is, increasingly, asking the same identity-based questions the customer in Mumbai making a lakh-rupee purchase is asking. The questions she is asking are: does this fit who I am becoming, does this fit who my children should see me as, does this align with what I believe matters, does this contribute to a story I want my life to tell?

These questions are unanswerable on a spec sheet. They are, however, eminently answerable in the design of an experience — in the brand voice, the use of the customer's vernacular, the values the company is observably staking, the community the customer is

being invited into, the story the experience is telling about the customer's place in the world. This is where the experience economy becomes most distant from product economics, and where the firms that have learned to build meaning differentiate most decisively from those that have not.

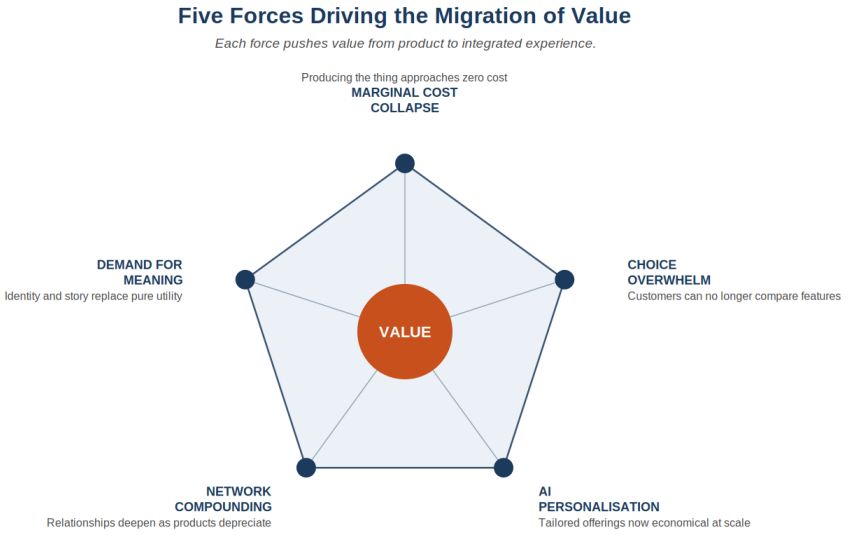


FIGURE 8.1 · Five distinct forces. One direction of travel.

Figure 8.1

◆ ◆ ◆

VI. The arithmetic, taken together

Each of the five forces, in isolation, would push value up the stack from product to experience. The forces operate together. Their effects compound. Each makes the others more powerful. The marginal cost collapse makes personalisation economical. Personalisation amplifies the value of network compounding. Network compounding deepens the meaning the customer constructs around the experience. The demand for meaning increases the customer's willingness to absorb choice overload by deferring to a curated experience. Each loop reinforces the next.

This is why I describe the migration as arithmetic rather than opinion. It is not because experience is in fashion or because the consultants of 2026 are recommending it. It is because the underlying economics of every category in which marginal cost has fallen, choice has expanded, AI is available, networks compound, and customers seek meaning — which is, conservatively, eighty percent of all consumer-facing categories — make experience the higher-margin, higher-retention, higher-LTV strategy. The arithmetic favours experience. The arithmetic is not negotiable.

[Figure 8.2 — to be supplied by illustrator. MARGINAL COST COLLAPSE: A CENTURY OF PRODUCT MANUFACTURING. A line chart showing real-cost-of-production indices for rep…]

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Track the marginal cost trajectory of your top three products over the past ten years.**

   If the trajectory is downward — and it almost certainly is — your competitive frontier has already migrated up the stack, whether or not your strategy reflects it.

2. **Test whether your customers experience choice fatigue with your product line.**

   A short customer survey, an unprompted focus group, a careful look at conversion rates by SKU count. If they do, you are competing on choice when your customer is asking you to compete on curation.

3. **Identify one experience layer your customers would value enough to pay a recurring fee for, and design it.**

   The recurring revenue from the experience layer, in most categories, exceeds the margin from the product itself within three years.

CHAPTER NINE

9

The Three New Currencies

Money is becoming a commodity.
Three other currencies become the new moats.

—

"The way to a man's heart is through his stomach.
The way to a customer's loyalty is through her time."
— A paraphrase, with apologies to the original

◆ ◆ ◆

A senior executive at one of India's largest advertising holding companies told me, late one evening in Bandra, that the most valuable customer relationships in her firm's history had not begun with a click. She said this as though she were confessing something. She had spent thirty years optimising for the click, with results she was professionally proud of. But she had recently looked, in detail, at the customer journeys behind the firm's largest accounts, and the journeys did not begin where her metrics dashboard said they began. They began in a magazine read on a flight, a recommendation from a friend whose taste the customer trusted, a brief encounter with a brand at an industry event, an unsolicited and unpaid mention in an article. By the time the click happened, the decision had largely been made. The click was not the start of the journey. It was the end.

She had, she said, been measuring the wrong currency. She had been measuring money — what the customer paid, what the campaign cost, what the click was worth — and she had been measuring it precisely. What she had not been measuring was attention, time, and trust — the three currencies that, in retrospect, had actually been the input to the customer's decision.

I have come to believe that her observation generalises. Money, in the convergence economy, is becoming a commodity. The cost of capital is volatile but available. The cost of executing a transaction is collapsing. The cost of producing the product is, as I argued in the previous chapter, falling toward zero. What is not a commodity, what is not collapsing, what is not available in any quantity at any price, is the customer's attention, the customer's time, and the customer's trust. These three are the currencies of the next economy. This chapter is about them.

◆ ◆ ◆

I. Attention, the scarcest resource on Earth

The first new currency is attention, and it is already, by some careful measures, the scarcest resource on Earth. The average customer in 2026 is exposed to roughly ten thousand commercial messages per day across every interface in her life — phone, laptop, billboard, podcast, television, email, social, in-store, in-vehicle. The cognitive bandwidth available to absorb any of these messages is the same as it was for her grandmother in 1956, when the figure was roughly three hundred messages per day. The arithmetic is impossible. Most of the messages are not absorbed at all; they pass over the customer like light.

The implication is that the customer's attention, when she gives it, is a deliberate gift of disproportionate value. The brand that earns her attention has been chosen against thousands of competitors who have failed to. The brand that holds her attention for thirty seconds has succeeded where the median brand fails in two. The

brand that wins her sustained attention — the kind of attention she gives a friend, a teacher, an artist, a faith — has won something that money cannot directly purchase, and that compounds in ways money does not.

The competitive implication is that every interface a brand operates becomes a stewardship question, not just a real-estate question. The notification that interrupts the customer for a transactional reminder costs her attention; if the brand has not earned the right to that interruption, the cost will be repaid in a slow erosion of the relationship. The advertisement that intrudes on the customer's focus costs her attention; if the advertisement does not return more value than it took, the customer's tolerance for the brand will quietly diminish. The product that continually demands attention to maintain a state — ten apps, six logins, four updates, three notifications — costs the customer attention every day, and competitors that demand less will eventually win her by virtue of not asking for what the customer is no longer willing to give.

The brands that win the attention currency in this decade will be the brands that ask for less and return more — the brands that respect the customer's bandwidth as a finite resource that the brand has been given temporary custody of. The brands that lose will be the ones that continue to behave as though attention is free.

◆ ◆ ◆

II. Time, the gift that AI can finally give

The second new currency is time, and it is the one that AI most directly enables. For the first time in commercial history, technology can produce, at scale, the sensation of having time given back to the customer. The agent that handles an hour of routine email has given the customer an hour. The smart-home system that pre-emptively manages the household's logistics has given the family back the half-hour they would have spent

coordinating. The AI tutor that adapts to the learner's pace has given the parents back the evenings they would have spent on homework supervision. The agentic financial assistant that has reconciled the family's accounts and paid the bills has given the household back the Sunday afternoon.

These returns of time are not abstractions. They are the highest-value thing AI does, and they are widely under-priced as a competitive lever. Most enterprise AI deployments in 2026 are still optimised for productivity inside the enterprise — efficiency for the company, not time-back for the customer. This is, in my view, a significant strategic error. The company that gives the customer back five minutes a week, every week, for a year, is a company the customer cannot easily replace, because the alternative is to relinquish those five minutes for something that may or may not return them.

I have begun to recommend to my clients a new operating metric: *minutes returned to the customer per week*. The metric is harder to measure than the conventional productivity metrics. It is also, in the long run, more predictive of retention. The brand that gives me my time back wins my loyalty. The brand that takes my time wastes my willingness to remain its customer. This is not a soft observation. It is the central commercial logic of the AI era applied to the customer side of the relationship rather than the operational side.

◆ ◆ ◆

III. Trust, the ultimate premium in a synthetic age

The third new currency is trust, and it is the one that is becoming structurally most scarce and therefore commercially most valuable.

The 2020s have introduced a phenomenon that earlier decades did not have to manage at scale: the indistinguishability, at the level of the customer's senses, between authentic and synthetic content.

Voice clones that pass for the original speaker. Deepfake video indistinguishable from genuine recording. AI-generated images that pass for photographs. AI-generated text in the voice of a public figure. AI agents that convincingly impersonate human customer-service representatives. AI-generated reviews. AI-generated testimonials. AI-generated news.

The cumulative effect is that the customer's default trust in any unverified content, any unknown sender, any unfamiliar voice, any ambiguous source, is structurally lower in 2026 than it was in 2016. This is rational. The customer cannot verify, by inspection, what is real. She has to rely on something else, and that something else is *trust* — accumulated, audited, durable trust — placed in specific brands, specific institutions, specific people, specific verification infrastructure.

The brands that have earned this trust — by transparent practice over time, by consistent honesty, by accountability when they have erred, by deliberate verification infrastructure — are now in a structurally privileged position. Their messages reach the customer because she does not first have to authenticate them. Their products are believed because she does not first have to research them. Their recommendations are absorbed because she does not first have to discount them. Trust is, in the synthetic-content era, a competitive advantage that money cannot buy and that AI cannot fake — at least not durably, not without being eventually exposed.

The brands that do not have it have, in 2026, an asymmetric problem. The cost of building it is not large in any single quarter. The cost of building it is large *cumulatively* — it is the cost of consistent practice over years, of a corporate posture that rewards the boring honesty over the exciting exaggeration, of governance that treats every customer interaction as evidence in the trust ledger. Most companies have not done this work. The few that have are, in this decade, beginning to collect the compounding return.

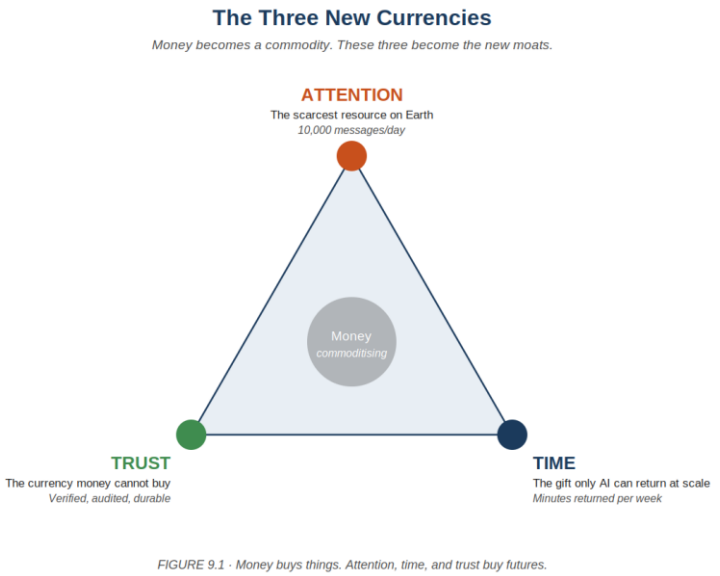


Figure 9.1

◆ ◆ ◆

IV. The compounding loop between the three

The three currencies are not independent. Each amplifies the other two, and the loop between them is what produces the durable competitive moats of the next decade.

A brand that respects attention earns the right to deliver more communication, which gives it more opportunities to demonstrate that it gives time back, which deepens the customer's trust, which lowers the friction on every subsequent interaction, which conserves attention for the next round. The loop runs forward. The customer relationship deepens at a rate the customer barely notices and the competitor cannot easily disrupt.

The opposite loop is also possible, and is more common in 2026 than companies admit. A brand that disrespects attention burns through the customer's tolerance, which forces it to demand more time to compensate for the friction, which erodes the customer's

trust in the brand's intentions, which raises the friction on every subsequent interaction, which forces still more demands on attention. This loop also runs forward, in the wrong direction. The customer relationship hollows at a rate the company barely measures and the customer is increasingly willing to leave behind.

The choice, for any leader in 2026, is not between these two loops in some abstract sense. It is between them in every specific interaction with every specific customer, and the cumulative results determine which loop the brand will be running in 2030.

[Figure 9.2 — to be supplied by illustrator. THE ATTENTION-TIME-TRUST COMPOUNDING LOOP. A circular diagram with three nodes (attention, time, trust) and arrows betwe…]

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Calculate the average minutes of customer attention your business currently consumes per week.**

   Notifications, emails, app sessions, ads, friction in service interactions. Then ask whether you have earned them. If the answer is "approximately," you have not.

2. **Identify one feature, notification, or process you could remove that would give your customer back five minutes a week.**

   Remove it this quarter. Tell the customer you have removed it. Watch what happens to the relationship.

3. **Publish a public trust dossier for your business — what you measure, what you disclose, what you stand by.**

   The publication itself is the act of trust-building. Companies that cannot publish such a dossier do not yet have the substance of one. Companies that can, deepen their trust currency by the publication itself.

CHAPTER TEN

10

The Experience Stack

Most incumbents compete on Layer 1. Disruptors win at Layers 2–3. Category-defining brands earn Layers 4–5.

—

"The whole is greater than the sum of its parts."
— Aristotle, Metaphysics

◆ ◆ ◆

In the second chapter of *Reimagining Reality* I argued for a four-layer model of human experience: sensation, interaction, memory, and meaning. The model has worked well as an introductory frame, and I have used it in dozens of workshops over the past three years. But the more I have used it, the more I have come to believe that the four-layer model — built primarily for descriptive purposes — is not the right model for *operational* purposes. It under-specifies the technological layer at which the experience economy of the 2020s is actually being constructed.

I have come to prefer, and use in working sessions with leadership teams, a five-layer model. The five layers are: product, intelligence, interaction, emotion, meaning. The model is borrowed in spirit from the OSI model that organised our thinking about networking infrastructure for forty years — each layer abstracted above the one below, each layer a legitimate site of competition, each layer

hosting different kinds of capability. I will defend the model for the rest of this chapter.

◆ ◆ ◆

I. The five layers, defined

Layer 1 is product. The thing itself — the device, the service, the physical place, the software interface, the document. This is the layer at which most of twentieth-century competition occurred. A better screen, a faster engine, a more functional service, a more attractive store. The product layer is real, and remains the foundation; experience is not a substitute for a working product. But the product layer alone is, in the convergence economy, the layer at which the largest share of competitive activity occurs and the smallest share of competitive advantage is durable.

Layer 2 is intelligence. What does the experience know about the customer? Memory, context, prediction, adaptation. The customer's history with the brand, her current state, her likely next need, her expressed and unexpressed preferences, her household context, her professional context, her cultural context. The intelligence layer is what AI principally enables, and it is the layer at which the most significant competitive shift of the 2020s is occurring. A product without intelligence is a generic offering. A product with intelligence is, at minimum, a personalised offering — and the difference is structural, not cosmetic.

Layer 3 is interaction. How does the experience respond to the customer in real time? Is it conversational or transactional? Is it responsive to her state of mind, her time of day, her physical location, her present mood? Does it feel like an attentive companion or a vending machine? The interaction layer is where the intelligence layer becomes manifest to the customer. The customer rarely perceives the intelligence directly. She perceives the interaction the intelligence enables.

Layer 4 is emotion. How does the experience make the customer feel? Surprise, delight, calm, pride, belonging, security, competence, joy. The emotion layer is the one that most product teams under-engineer and that most category-defining brands have learned to engineer with extraordinary care. The Apple unboxing is a Layer 4 design. The Disney park entrance is a Layer 4 design. The IKEA breakfast is a Layer 4 design. None of these is an afterthought. Each is the result of deliberate, expensive, sustained engineering of a particular emotional experience that the brand has decided is its signature.

Layer 5 is meaning. Why does this matter to the customer as a human? Identity, story, belonging, alignment with values, contribution to the life she is constructing. The meaning layer is the highest, the rarest, and the most durable. Once a brand has earned a place in the meaning layer of a customer's life, the customer's loyalty is not a choice she revisits — it is a part of who she is. The brands that occupy this layer are few. They tend to be the brands the customer would defend in conversation, identify with publicly, hand down to her children, and build small parts of her life around. Apple lives here. Tata, in India, lives here. Disney, IKEA, certain religious institutions, certain cultural institutions, certain civilisational practices live here. The list is short, and the position is hard-won.

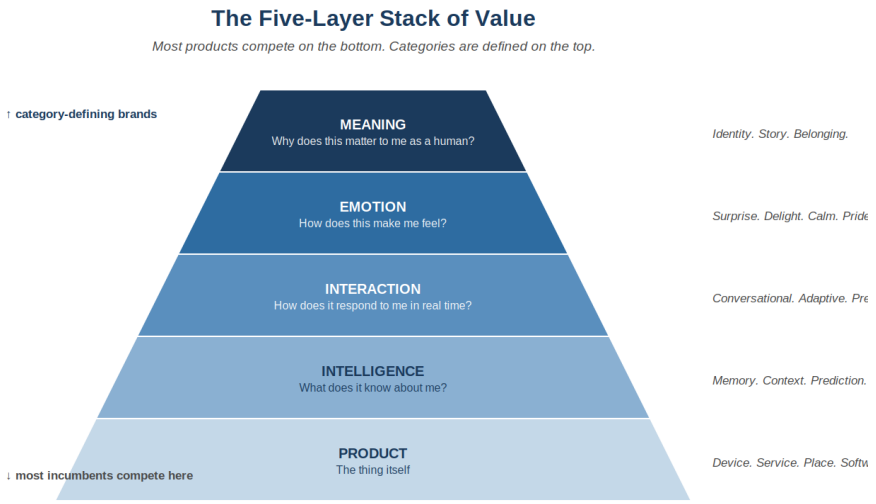


Figure 10.1

◆ ◆ ◆

II. Where everyone is competing right now

Most incumbents in most categories, in 2026, are still competing primarily at Layer 1. They are running their organisations, their R&D budgets, their go-to-market strategies, and their performance metrics around the product layer. This is not because they are stupid. It is because the product layer is the layer at which their organisations were built, the layer at which their data is structured, and the layer at which their competitive comparisons are conducted. It is also the layer at which competitive advantage is most ephemeral, and at which margin pressure is most relentless.

The disruptors who are most successfully attacking incumbents in 2026 are competing primarily at Layers 2 and 3 — intelligence and interaction. They are not, in most cases, building better products. They are wrapping equivalent or marginally inferior products in dramatically better intelligence and interaction. The bank that

knows the customer's financial situation, anticipates her next need, and responds in conversational tone, in her language, in her preferred channel, beats the bank with the better account-feature comparison. The healthcare provider that remembers the patient, contextualises her current visit against her history, and responds in a way the patient experiences as continuous, beats the provider with the better physical facility. The retailer that knows the customer, anticipates her purchases, curates for her, and interacts as a trusted advisor beats the retailer with the larger SKU count.

The category-defining brands, the ones that will be the fixed points of the next economy, are competing at Layers 4 and 5 — emotion and meaning. They are designing the feelings their customers have at specific moments. They are constructing the meaning their customers are invited to share in. These brands are not numerous. They are not any specific set of brands forever. But the position they hold is the position other brands cannot reach by spending more money on the bottom layer of the stack. Layer 5 is not for sale.

◆ ◆ ◆

III. How to read your own competitive position

There is a diagnostic exercise I run with clients that I will reproduce here. Take the flagship product of your business. Score it, on a scale of one to ten, on each of the five layers. Score it honestly — not as you wish to be, but as your customer would score it if she knew the framework. Most leadership teams, when they do this exercise honestly, find that they score eight or nine on Layer 1, three or four on Layers 2 and 3, and one or two on Layers 4 and 5. The product is good. The intelligence is thin. The interaction is transactional. The emotion is neutral. The meaning is absent.

Now take the most threatening new entrant in your category, and score it the same way. The pattern, in my experience, is consistent: the disruptor scores six or seven on Layer 1 (often genuinely

behind the incumbent on the product itself), eight or nine on Layers 2 and 3 (this is where they win), and four or five on Layer 4 (developing). The disruptor is not winning by being a better product. The disruptor is winning by being a better experience around an adequate product.

Now take the brand in any adjacent category that you most admire as a customer, and score it. The pattern, again, is consistent: solid on Layer 1, strong on 2 and 3, deliberate and impressive on 4, present (sometimes uniquely so) on 5. This is the brand you would not switch from for a small price difference, because the brand is not asking you to compare it on price. The brand is asking you to compare it on who you are when you choose it.

[Figure 10.2 — to be supplied by illustrator. WHERE COMPETITORS SIT ON THE STACK. A radar chart with five axes (one per layer). Three overlaid polygons: incumbent (he…]

◆ ◆ ◆

IV. How to climb the stack

The climb is the work of the next several years for every serious enterprise, and it is uneven across organisations. Some can climb from Layer 1 to Layer 2 within a year, given the data they already have and the agentic capabilities now available. Others have to first restructure their data architecture before they can climb at all. Most can begin to engineer Layer 4 — emotion — within twelve months, with the right design talent and the willingness to redesign at least one customer touchpoint deliberately. Layer 5 — meaning — cannot be engineered within a single year. It requires a sustained, multi-year alignment of brand, behaviour, product, communication, and conduct, and it is built through trust accumulated over time. It cannot be bought. It cannot be hired. It can be lost in a single bad quarter and rebuilt only over many good ones.

The leader's task, in 2026, is not to attack all five layers at once. The task is to honestly assess where the company sits, identify the one layer that is now most under-served by the company's investment relative to its competitive opportunity, and commission a single project this year that demonstrably moves the company up that layer. The compound effect of doing this every year for five years is a transformation. The compound effect of not doing it is a slow erosion that is, by year five, very difficult to reverse.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Map your top product against all five layers and score yourself one to ten on each.**

   Do this with three independent leadership-team members, then reconcile. The disagreements are where your competitive blind spots live.

2. **Identify the single layer at which your most threatening new entrant competes, and respond.**

   The response is not to copy. The response is to decide whether to compete on the same layer, leapfrog to a higher one, or differentiate by occupying a layer the entrant cannot.

3. **Commission a single project this quarter that demonstrably moves you up one layer.**

   One layer in one quarter is the right cadence. Five layers in one quarter is theatre. One layer per quarter, sustained, over five years, is transformation.

CHAPTER ELEVEN

11

Memory is the Real Metric

What people remember, not what they click.

—

"People will forget what you said, people will forget what you did, but people will never forget how you made them feel."

— *Maya Angelou*

◆ ◆ ◆

A retail executive in Mumbai, whose career had been built on customer satisfaction metrics, told me that she had recently audited her own data and found a result she could not initially believe. The customers who scored a nine — out of ten — on her brand's net promoter score at point of sale were, three months later, no more likely to remember the brand name than customers who had scored a four. The point-of-sale satisfaction was real. The thirty-day memory was random. The metric on which her entire customer experience programme had been built was, she was forced to conclude, predicting nothing about whether the customer would actually return.

She was, of course, not alone. She was experiencing what the psychology research has been telling us for fifty years and what most of the consumer-marketing world has been politely declining to absorb. The metrics that we measure most easily — clicks, conversions, satisfaction-at-moment — are not the metrics that

predict the customer's future behaviour. The metric that predicts future behaviour is memory: what the customer actually remembers, how she narrates the experience to others, what shape the memory takes in her broader life story. And memory, it turns out, is constructed by a small number of moments that look very different from the average of all moments.

This chapter is about that gap. It is also about what to do about it.

◆ ◆ ◆

I. The peak-end rule, and what it actually says

In 1993, the psychologist Daniel Kahneman and his collaborators published a study of how people remember experiences of varying duration and intensity. The result was counter-intuitive and has been reproduced many times since. People do not remember the average of an experience. They remember it as approximately the average of two specific moments: the emotional peak (whether positive or negative) and the end. Everything else fades. This is the peak-end rule, and it has profound implications for how an experience should be designed.

The implications are easy to state and surprisingly difficult to act on. First, *the peak matters far more than the average*. An experience that is mostly mediocre but has one extraordinary moment will be remembered as extraordinary. An experience that is uniformly good but has no peak will be remembered as adequate. Most product teams, optimising the average, are actively suppressing the peaks they should be amplifying. Second, *the end matters disproportionately*. An experience that is excellent throughout and ends poorly will be remembered poorly. An experience that is mediocre throughout and ends beautifully will be remembered as good. Most companies, in my observation, end their customer experiences badly — the post-purchase experience, the offboarding, the cancellation, the goodbye — because they have nothing in particular to gain from the customer at that point. They

are wrong. They have everything to gain at that point, because the end is half the memory.

The implications for business are operational. The customer's memory of her relationship with a brand is dominated by the emotional peak of the relationship and by how it ends. Almost everything in between fades. If the peak is unmemorable and the end is poor, the customer will not return regardless of how good the average was. If the peak is extraordinary and the end is dignified, the customer will return regardless of how mediocre the average was.

◆ ◆ ◆

II. Why the moment metrics are so misleading

The moment metrics — click-through, conversion, point-of-sale satisfaction, post-call survey, NPS at touchpoint — are easy to collect and almost completely uncorrelated with what the customer will remember thirty days later. This is not because the metrics are invalid; the customer really did experience the moment they captured. It is because the moment in which the metric was collected is, in most cases, not one of the two moments the customer will remember. The metrics capture the average. The customer remembers the peak and the end.

This is why a company can have a high satisfaction score and a low retention rate. It is why a company can have a low complaint rate and a low repurchase rate. It is why customers can rate every interaction with a brand as adequate and still leave the brand for a competitor whose interactions were no better in any individual moment but whose peak moment was unforgettable. The metrics dashboard, optimised for the average, was hiding the absence of a peak — and the absence of a peak was the actual cause of the brand's decline.

The retail executive in Mumbai, when she absorbed this, did the most useful thing a leader can do with this insight. She reorganised her measurement programme. The point-of-sale NPS continued to be collected — it remained a legitimate diagnostic of moment quality — but it was no longer the primary outcome metric for her customer experience function. The new primary outcome metric was thirty-day spontaneous brand recall, complemented by a quarterly sample of "describe your most memorable interaction with this brand in the last year." Both metrics were harder to collect. Both were vastly more predictive of repurchase. Within eighteen months her team's prioritisation had shifted from optimising the average to engineering the peaks.

◆ ◆ ◆

III. Engineering peaks deliberately

Once a leadership team accepts that the peak is what the customer will remember, the question becomes operational: how do you engineer a peak deliberately?

The first answer is that you have to decide *what feeling the peak should produce*. This is more difficult than it sounds. Most companies have not articulated, in writing, what specific emotion they want the customer to feel at the peak of the experience. Surprise? Delight? Pride? Calm? Belonging? Awe? Recognition? The feelings are not interchangeable. The design of an experience that produces awe is not the design that produces calm. The leadership team must choose, and the choice must be specific enough to brief a designer.

The second answer is that you have to *engineer the moment in which the peak will occur*. The peak will not happen on its own. Most touchpoints in most customer journeys are flat — they neither amplify nor diminish the customer's experience. The peak is constructed at a specific moment — the unboxing, the first use, the first failure-and-recovery, the surprise upgrade, the personal

acknowledgement, the gift, the rite of passage — and the construction is a deliberate act of design. The peak does not happen because the company is well-meaning. It happens because a small team somewhere has decided where in the journey the peak will occur, what emotion it will produce, and what specific elements of the customer's encounter will produce that emotion.

The third answer is that the peak has to be *repeatable*. A peak that occurred once for one customer is a story, not a metric. A peak that occurs reliably for every customer of a given segment, every time, is a competitive asset. The reliability is what distinguishes a category-defining brand from a brand with occasional good moments. Engineering a repeatable peak requires standardisation of the elements that produce the emotion — the script, the timing, the artefact, the surroundings, the language. This is mundane work. It is also what produces durable competitive advantage at the emotion layer.

◆ ◆ ◆

IV. Designing the end

The other half of memory is the end, and most companies neglect it. The reason is intuitive: the company has no further commercial goal at the end of a relationship. The customer is leaving. Why invest in her experience when there is no transaction at stake?

The reason to invest is precisely because the end is half the memory. The customer who leaves with dignity, with her data preserved, with a graceful acknowledgement of the relationship, with a reasonable mechanism for return, will *remember the brand well*. She will recommend it to friends who are evaluating it. She may, in time, return. She will, in any case, not become an active detractor.

The customer who leaves badly — frustrated by exit friction, condescended to by save-attempts, mistreated by the offboarding

process — will remember the brand badly, regardless of how good the rest of the relationship was. She will become an active source of negative word-of-mouth. She will not return. She will tell the story of her bad ending in places where the company cannot see and cannot respond. The end is the half the company most easily ignores and the half the customer most permanently remembers.

The discipline of designing the end well is rare. The companies that practice it — typically the same companies that have learned to engineer peaks well — accumulate a reservoir of departed customers who continue to speak well of the brand. This reservoir is, in my experience, a meaningful contributor to category-defining brand status. The brands at Layer 5 of the experience stack tend, almost without exception, to have invested in their endings.

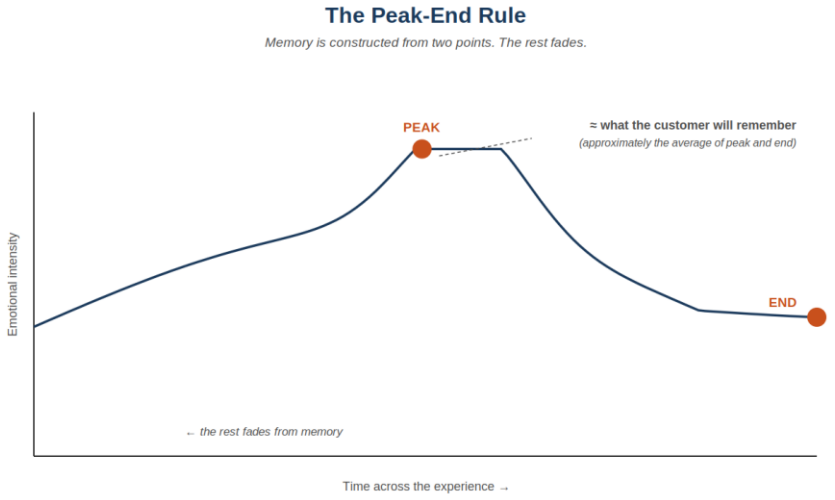


Figure 11.1

[*Figure 11.2 — to be supplied by illustrator. MOMENT METRICS vs MEMORY METRICS. A two-column comparison. Left column: metrics most companies track today (NPS, CSAT, c…*]

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Add at least one memory-based metric to your customer dashboard within the quarter.**

   Thirty-day spontaneous brand recall, willingness to narrate, end-state quality. The metric will be harder to collect than NPS. It will be vastly more predictive.

2. **Audit your top customer journey for its peak and its end — both deliberately.**

   If the peak is not engineered, design it. If the end is not engineered, design it. Most companies need to do both.

3. **Stop optimising the average of any customer experience that has a memorable peak or end.**

   The optimisation of the average dilutes the peak and softens the end. Both are mistakes. Engineer the two moments. Allow the rest to be merely adequate.

CHAPTER TWELVE

12

Bharat as Experience Laboratory

Civilisations that have thought hardest about experience will build best.

—

"In India, even the gods are characters in a story,
and the story is what makes them gods."
— *A. K. Ramanujan, paraphrased from his lectures*

◆ ◆ ◆

A foreign delegation visiting one of our virtual-darshan installations in Vrindavan in 2024 watched, for nearly an hour, a continuous queue of devotees passing through a small enclosure in which a high-resolution immersive rendering of a temple sanctum was being displayed. The delegation had come to see a technology project. They were watching, instead, a steady stream of people of every age and class — schoolchildren in uniform, businessmen in suits, elderly women in saris, tourists with selfie-sticks — entering the enclosure quietly, standing for one to three minutes in front of the rendering, and emerging with the specific expression on their faces that I have come to recognise as the expression of darshan completed. The expression is calm, often tearful, often slightly dazed. It is, in the technical sense, the expression of an experience that has been received and absorbed. The delegation, which was led by the head of an American immersive-technology firm, eventually said the thing that delegations from outside India often

say after watching such a scene: "We have never seen this many people deliberately, joyfully, give their attention to a screen."

I did not have the heart to correct him on the screen part, because the technology was not what he was actually responding to. What he was responding to was a three-thousand-year-old experience architecture, of which the immersive technology was merely the most recent expression. The architecture was the experience. The technology was the new vehicle for the architecture. The vehicle would, in time, change again. The architecture would persist, as it has persisted across generations of vehicles, because the civilisation that built the architecture has thought longer about experience as a first-class category than any other civilisation on Earth.

This chapter is the argument that India's structural advantage in the convergence economy of the next decade is not a function of demographics, English fluency, IT services, or cost — though all of these matter. The structural advantage is civilisational. The civilisations that have thought longest about experience as a first-class human category are positioned to build the new economy of the convergence decade better than the civilisations that have thought primarily about the product as the unit of value. India is one such civilisation. So are several others — Japan in some specific dimensions, Italy in others, parts of Latin America in still others — and the leaders who recognise this will build differently and better.

◆ ◆ ◆

I. Five Indian experience architectures

I want to walk five of them, because the abstraction does not do the argument justice. Each is, in the technical sense, an experience architecture — a set of structural choices about how a human being is invited to move through, to feel within, and to remember a particular kind of human encounter.

Darshan. The ritual of seeing and being seen by the divine, practised across South Asian religious traditions for at least three millennia, is — in the analytic vocabulary of the experience economy — a multi-sensory, embodied, placeless-yet-placed encounter that is designed to produce a specific emotional and spiritual state in the participant. The architecture is precise. The participant's path of approach is structured. The sensory environment is engineered — incense, bells, the heat of oil lamps, the specific frequencies of mantra. The moment of darshan itself is a peak. The departure is a deliberate end. The return is a continuous practice, not a one-time event. Every element of what experience designers in the West are now learning to construct deliberately has, in the darshan tradition, been operating at scale for thousands of years.

The ramleela and other classical theatres. From Sanskrit drama through the regional theatres of Maharashtra, Bengal, Karnataka, Kerala, the Indian dramatic tradition has assumed, for two thousand years, that the audience is a participant rather than a spectator. The ten-day Ramleela of north India is not a play. It is an immersive, durational, communal re-enactment in which the community itself is the performer, the audience, the set designer, and the ritual context, in shifting combinations across the ten days. The category that the Western entertainment industry is now reaching for under the label "immersive theatre" is one that the Indian theatre tradition has been operating in continuously for centuries.

Tirth-yatra. The Indian pilgrimage tradition is a journey-as-experience architecture of extraordinary sophistication. The traveler does not simply arrive at the destination; the journey itself is the experience, structured by stops, encounters, austerities, rituals, and small communal acts that produce an internal transformation by the time the destination is reached. The destination is not the experience; the entire arc is the experience.

This is, in the language of customer journey design, the longest-form experience architecture in continuous human use.

The Indian wedding. The multi-day, multi-event Indian wedding is, in commercial scale and in experiential design, one of the most sophisticated experience products humans have built. Each ceremony is a designed peak, with a specific emotional register and a specific role for the participants. The arc of the four to seven days is engineered, deliberately or by tradition, to produce a sequence of intensifications culminating in a peak and a dignified end. The economic scale of the Indian wedding industry — by some estimates fifty billion dollars a year — is the financial expression of an experience design tradition that families have refined for generations.

Raga and the classical music tradition. The Indian classical music tradition built, over a thousand years, a structured framework in which a musician constructs an experience for an audience over a multi-hour performance, beginning with a slow introduction, building through deliberate intensifications, reaching a peak that audiences can feel viscerally and remember decades later, and resolving into a quiet end. The framework — alap, jor, jhala, gat — is, when one looks at it carefully, a peak-end-rule architecture refined across a thousand years of audience research conducted by performers in conversation with their listeners.

[Figure 12.1 — to be supplied by illustrator. FIVE INDIAN EXPERIENCE ARCHITECTURES, THEN AND NOW. A two-column table. Left column lists the five architectures with br…]

◆ ◆ ◆

II. What this advantage actually is

The advantage is not that Indian companies will build better cars, or better software, or better consumer electronics. The advantage is that, when the technology of the convergence decade allows

experiences to be built at scale across every category of human life, the firms whose intuitions about experience are deepest will design better experiences. India's intuitions about experience are unusually deep, because the civilisation has been working on the problem for a long time.

I want to be careful about the boundary of this argument. I am not claiming Indian moral or cultural superiority, which would be silly. I am not claiming that every Indian company is poised to build great experience businesses, which would be empirically false. Most Indian companies, in 2026, are still organised around product economics, not experience economics, and they will struggle in the decade ahead exactly as the equivalent companies in any other country will. What I am claiming is more specific: the *civilisational substrate* in India is unusually well-suited to the experience economy, and the firms — Indian or otherwise — that draw deliberately from this substrate will build better products of the experience economy than firms that do not.

The drawing-from is not automatic. It is a deliberate creative act. It requires a designer who has read Sanskrit drama as well as service design textbooks. It requires a product team that has visited a Ramleela alongside their visits to design conferences. It requires a brand team that has thought about darshan as a precedent for the Layer 4 emotion engineering of the company's signature peak. It requires, in short, a posture toward Indian civilisation that treats it as a working creative resource rather than as a museum curiosity. The Indian companies that adopt this posture, and the global companies operating in India that hire creative leadership from this tradition, will find themselves with capabilities that competitors elsewhere cannot easily replicate.

◆ ◆ ◆

III. Three concrete implications for builders

The argument, to be useful, has to translate into specific decisions that builders can make in 2026. I want to name three.

First: read the tradition as design source material. Most Indian designers and product builders I work with — including, for many years, myself — have not read the Indian experience tradition systematically as a source of design ideas. We have read Norman, we have read Pine and Gilmore, we have read Christensen, we have read the published case studies of Apple and Disney and Pixar. We have not, on average, read the Natyashastra or the Tirthankalpataru or the wedding manuals or the texts on raga performance. The omission is, I think, a strategic mistake. The Indian texts are full of usable design wisdom. The company that hires a humanities scholar onto its design team — the way good firms now hire data scientists — will be astonished by what becomes accessible.

Second: design for vernacular India before English India. The argument that India's civilisational substrate is rich does not apply uniformly across the country's English-speaking elite. It applies most densely in the vernacular, the regional, the rural, the traditional — in the parts of the country where the experience traditions have been most continuously transmitted. Companies that design for this India first, and translate to English India second, will draw more directly from the substrate. Companies that design for English India first will, often inadvertently, dilute the very civilisational advantage they could have drawn upon.

Third: build experience architectures that work across generations rather than across quarters. The Indian experience traditions did not arise from companies optimising for quarterly returns. They arose from communities and institutions that took a multi-generational view of what an experience should accomplish. The companies that build with this longer time-horizon — and the investors who fund them on terms compatible with that horizon — will produce

experience architectures that compound across decades. The companies that bring conventional product-economics quarterly thinking to the work will produce thin, short-lived experience offerings that do not compound.

[Figure 12.2 — to be supplied by illustrator. THE CIVILISATIONAL STACK. A three-layer vertical diagram. Bottom layer: tradition (the experience architectures of a par…]

◆ ◆ ◆

IV. A small note for non-Indian builders

This chapter has been written largely for Indian readers. The argument it makes, however, is not exclusively Indian. Every civilisation has experience architectures that are deeper than its product-economics surface allows it to acknowledge. Japan has architectures around hospitality (omotenashi), aesthetics (wabi-sabi, ma), and craftsmanship that produce experience advantages most non-Japanese firms have not yet internalised. Italy has architectures around food, craft, family, and place that produce experience advantages French and German firms know to study. Mexico, Iran, Senegal, Indonesia, Egypt — each has architectures of its own. The general principle is that the civilisational substrate of a place is a competitive resource for the firms operating there, and the firms that read it carefully build better experience products than firms that ignore it.

For Indian readers, this is an unusually valuable opportunity, because the substrate is unusually rich, the global competition for experience-economy primacy is wide open, and the technology to build at industrial scale has finally arrived. For non-Indian readers, the same logic applies to your own substrate. Read your tradition. Build from it. The convergence decade will reward firms that have done so.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you are an Indian business, identify one civilisational architecture you under-use as design source material — and commission a project that uses it.**

   The civilisational architecture is sitting next to your design team, available for free, and not being read.

2. **If you are a global business operating in India, hire one humanities scholar onto your design team this year.**

   The hire will be inexpensive, will produce uncommon insights, and will give you access to a pool of design intelligence your competitors are not drawing on.

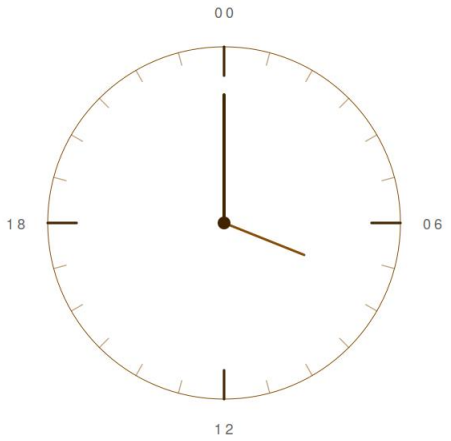
3. **Commission a single experience-design project that draws explicitly from a non-Western tradition.**

   Document the project's outcomes. The company that does this once tends to find it does not stop doing it, because the results are too useful.

PART FOUR

IV

A Tuesday in 2030



From abstract trends to a Tuesday morning in 2030. Work, learning, healthcare, the home, money and identity, examined hour by hour and domain by domain.

THE WORLD THAT'S BEGINNING

From abstract trends to a Tuesday morning in 2030.
Work, learning, healthcare, the home, money and identity,
examined hour by hour and domain by domain.

CHAPTER THIRTEEN

13

Aarav, Meera, Two Kids, Lucknow

Every detail in this chapter is technologically achievable in five years.

—

"The future is here. It's just not very evenly distributed."

— William Gibson, again, because the line will not stop being relevant

◆ ◆ ◆

I want to walk a single day. The day is a Tuesday in October 2030. The family is composite — built from elements of clients, friends, students, and my own household — but every detail in the day is technologically achievable within five years of the time of writing. Some of the details are already here in my own life in 2026, in earlier or rougher form. None of them is science fiction. The composite picture is meant to bring the abstractions of the previous twelve chapters into the texture of an ordinary, urban, Indian, middle-class life.

The family is Aarav and Meera, both in their late thirties, with two children — a daughter aged twelve, a son aged nine. They live in Gomti Nagar, in Lucknow, in a building that was completed in 2027 by a developer who treated the property as a platform rather than as a sale. Aarav is a senior product manager at a financial services firm headquartered in Mumbai but with most of its workforce now distributed; Meera is a doctor, running her own clinic about thirty minutes from the apartment. The children are in

two different schools, both of which have been substantially redesigned in the past three years.

I will walk the day from morning to evening, briefly. The chapters that follow this one will examine each domain in detail. This one is the picture, not the analysis.

◆ ◆ ◆

06:30 — Wake

The home wakes Aarav, not with an alarm, but with the slow, sequenced lighting of his bedroom across a fifteen-minute window calibrated to his sleep cycle and to the day's schedule. The system has noticed that his REM cycle ended eight minutes ago and that this is the optimal window. The air-purification system, predicting today's PM 2.5 spike based on weather data and on the construction activity scheduled three blocks east, has been pre-emptively running for ninety minutes. The coffee machine has begun its cycle. Aarav does not think about any of this. He thinks about the meeting he has at ten and the thing he wants to say to his daughter about her science fair project before she leaves for school.

07:15 — The Agent's Brief

He brushes his teeth. As he does, his AI assistant — running on the discreet earpiece he wears and on the bathroom mirror that doubles as a display — briefs him in his own voice, in his own characteristic phrasing, on what his agentic systems have done overnight on his behalf. Forty-three emails received; thirty-eight handled, four flagged for his attention with proposed responses already drafted, one held back as requiring his judgment. Three calendar conflicts identified; two resolved, one needing his decision. The market summary in his own accent. The two news items his profile says he will care about today. The brief is two minutes. He nods, dictates two short instructions, and walks out of the bathroom.

07:45 — Family Breakfast

Breakfast is human. The home's intelligence has receded into the background. Meera is in conversation with their daughter about a science fair project that involves mapping local water quality. The son is showing his father a drawing of a creature he invented yesterday in his metaverse history class — an Akbar-era court official, somewhat exaggeratedly large, whom he met during a walking tour of the Mughal court. Aarav is amused. The drawing looks like a creature. The boy insists it is a faithful likeness. Both of them are right, after a fashion.

08:00 — Mixed Commute

Aarav goes to the firm's Lucknow office two days a week. This is one of those days. The car drives itself for the twenty-five-minute commute. Aarav uses the time, which used to be lost time, in a spatial-AR collaborative session with three colleagues — one in Mumbai, one in Bengaluru, one working from her home in Pune — who appear to be sitting around the car's interior with him. They review a pricing proposal. He suggests a change. The change is captured. The session ends as the car pulls into the office parking. He has used the commute for thirty-five minutes of work that used to require an hour at his desk.

09:30 — Meera's Morning

Meera, meanwhile, has driven herself to her clinic — she prefers to drive — and has begun the morning's patient list. The clinic's AI has prepared, for each patient, a one-page contextual brief: history, recent vitals from the patient's wearable, latest test results, the patient's specific concerns from the chat she had with the AI intake system the previous evening, and a first-pass differential diagnosis that the AI labels clearly as suggestion. Meera reviews each brief in fifteen seconds. She then sees the patient. Her time with the patient is unbroken by note-taking; the conversation is transcribed by the

room's ambient AI and structured into the clinic's records by the time she steps out. Her capacity has tripled in three years. So has the warmth of her interactions, because she is no longer dividing her attention with the keyboard.

10:00 — The Quarterly Business Review

Aarav's ten o'clock is the firm's quarterly business review, and it is the meeting that, three years ago, used to consume an entire day of his and his team's preparation. It now consumes ninety minutes total. The AI has, the previous evening, summarised the firm's quarterly performance against plan, drafted three scenario analyses for the next quarter, costed each scenario, and produced the visualisation deck. The meeting itself is humans debating the scenarios — the trade-offs, the implications, the choice. The AI has prepared the data and the options. The humans make the decisions and own the consequences.

13:00 — Meera's Sixty Patients

Meera has, by lunch, seen thirty patients and will see thirty more before the clinic closes at six. Three years ago she would have seen eighteen in a full day. The increase is not because she is rushing. It is because the AI has absorbed the parts of her day that used to consume her — intake, history, transcription, first-pass diagnosis, prescription drafting, billing, follow-up scheduling. Her conversations with patients are longer and more present than they used to be.

16:00 — The Children's Learning

The daughter is in her room with her AI tutor, working on the science fair project — measuring local water samples, analysing them against open-government datasets, generating hypotheses. The tutor is teaching her, in a Hindi-English mix she finds comfortable, about confounding variables. The pace is hers. The

tutor knows when she is bored and when she is challenged. Her parents are not asked to supervise the homework, and have not been for a year and a half.

The son, meanwhile, is with his classmates in his metaverse history class. The class is walking through Mughal Lucknow as it appeared in 1722. His classmates are present as their avatars — six children from across India. The teacher, who is in Hyderabad in physical reality, walks them through the bazaar. They stop at the spice merchant. They watch how transactions occurred. They listen to the language. The son, later, will tell his father about the merchant who sold cardamom from Kerala that had travelled four months to reach Lucknow.

19:00 — The Evening

Aarav comes home. Meera comes home. The home, having anticipated the family, has the lighting and the temperature ready. Dinner has been prepared by a service the family uses three nights a week. The four of them eat together. The phones are off, by family agreement, between dinner and the children's bedtime. The conversation is human. The technology, in this hour, is invisible.

21:30 — The Parents' Hour

The children are in bed. Aarav and Meera have an hour together. They watch a film that the home's recommendation system has chosen, based on what they have liked over a long period and on a current mood-state inference that they trust without entirely understanding. They drink tea. They talk briefly about money — about the small experimental investment they made last week in a tokenised real-estate fund, about the children's college savings, about Meera's clinic's new equipment. The conversation is short. The systems have done most of the housekeeping. The decisions left to them are the ones that actually matter.

[*Figure 13.1 — to be supplied by illustrator. A TUESDAY IN 2030: THE HOUR-BY-HOUR DAY. A vertical timeline of the day from 06:30 to 22:30, with each hour marked and t…]*

◆ ◆ ◆

What this picture is, and is not

The picture is composite. The technologies are not. Each of them — the sleep-cycle wake, the agentic morning brief, the spatial collaborative commute, the AI clinical co-pilot, the agentic enterprise meeting prep, the AI tutor, the metaverse class, the recommendation, the tokenised investment, the home that has receded into the background — exists somewhere in the world in 2026, in early or rough form. The picture is what these technologies look like when they have matured to ordinary use.

The picture is also not utopian. It does not pretend that every Indian family in 2030 will live this way. Most will not. The penetration of the technologies will be uneven across geography and class for many years. The picture is what the affluent middle class of urban India will live through — perhaps a hundred million households — by 2030, and what is genuinely available to a wider population somewhat later.

The remaining chapters of this part — work re-composed, learning unbundled, healthcare as continuous practice, the home as platform, the trinity of money, identity, and ownership — examine each of the domains in detail. This chapter is the picture. The chapters that follow are the analysis.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Write your own version of this scene for your customer's life in 2030.**

   One page. Hour by hour. Imagine specifically. The exercise will tell you, very fast, where your business currently sits in your customer's day and where it should sit by 2028.

2. **Identify which moment in that scene your business already participates in, and which moment it should be participating in by 2028.**

   Most businesses participate in the customer's day in one or two moments. Most businesses should be participating in three to five.

3. **Identify the moment in the day at which your business is structurally absent — and decide whether to enter it.**

   Some moments are not yours to occupy. Others are wide open and waiting for the firm that can earn the customer's invitation.

CHAPTER FOURTEEN

14

Work Re-Composed

Roles dissolve. Outcomes endure.

—

"The illiterate of the 21st century will not be those who cannot read and write, but those who cannot learn, unlearn, and relearn."

— Alvin Toffler

◆ ◆ ◆

The HR director of one of India's larger private banks called me in for what she described, on the phone, as a difficult conversation. When I arrived at her office in Bandra-Kurla Complex, she had laid out, on a long table, the bank's organisation chart printed on twelve large sheets of paper joined together. She walked me along the chart for nearly half an hour, level by level. At the end of the walk she said: "I have been looking at this for three weeks. Every time I look, I see more roles whose primary function is to shuttle information between the level above and the level below them. I have counted forty-three such roles, just in the part of the bank I am responsible for. None of them is a bad employee. Each of them is doing what the design of the role requires. The design of the role no longer requires what it used to."

She is not unusual. She is, in 2026, ahead of most of her peers. The recognition that information-shuttling middle management has become structurally redundant is one that every large enterprise

will have to absorb in the next three years. The bank she ran was, in this sense, an early indicator. The organisation she would build over the following twenty-four months — flatter, more outcome-shaped, with skills graphs replacing CVs and AI agents handling the information shuttling — was the organisation most large enterprises will have to build by the end of the decade.

◆ ◆ ◆

I. Three eras, one arc

It helps to look at the work-redefinition arc across three eras. The first era — call it pre-2020 — had fixed roles, fixed titles, hierarchical management, annual performance reviews, and a career conceived as a single ladder up which the employee climbed for forty years. The second era — 2020 through 2026 — introduced remote and hybrid work, gig and portfolio careers, manager-as-coach instead of manager-as-supervisor, continuous feedback in place of annual reviews, and a career conceived as a lattice rather than a ladder. The third era — 2027 through 2030 and beyond — composes work into fluid teams organised around outcomes, with AI augmenting every role, humans setting context and AI executing within that context, skills graphs replacing CVs as the unit of professional identity, and a career conceived as a continuously curated portfolio of capabilities that the worker takes with her across many assignments.

The transition from era two to era three is the significant one. Most enterprises in 2026 are operating partly in era one and partly in era two, with experiments in era three. The transition is the work this chapter is about.

◆ ◆ ◆

II. Why information-shuttling roles dissolve first

The roles whose primary function is information shuttling — the layers of middle management whose job is to summarise the level

below and report it upward, and to translate the level above and disseminate it downward — are exactly the kind of work that AI does well. The synthesis is structured. The translation is repetitive. The judgment, in the median case, is procedural rather than strategic. AI agents can, in 2026, perform most of this work at higher quality, higher speed, and at orders of magnitude lower unit cost than human middle managers.

The implication is not that middle management dissolves entirely. It is that the *information-shuttling* portion dissolves, and that the remaining portion — coaching, contextual judgment, talent development, conflict resolution, organisational leadership — concentrates into a smaller number of roles played by people who do those things genuinely well. The leader's responsibility, as this transition unfolds, is twofold: to redesign the shuttling roles deliberately rather than letting them dissolve under the embarrassment of redundancy, and to invest in the elevation of the roles that remain.

◆ ◆ ◆

III. The skills graph replaces the CV

The CV is a twentieth-century artefact that has long outlived its usefulness. It is a list of past employments, with titles, dates, and brief descriptions of responsibilities. It is, in almost every case, a poor predictor of what a person can actually do today.

The replacement is the skills graph. The skills graph is a structured, continuously updated representation of a person's capabilities — what she can do, at what level of mastery, in what contexts, validated by what evidence. It is updated by the person's actual work, by the projects she completes, by the validations she receives from collaborators and from systems, and by the certifications she earns. It is portable across employers, but the depth of the validation is highest within the systems that have observed her work directly.

By 2030, the skills graph will be the dominant representation of professional identity for knowledge workers under thirty-five. The CV will continue to exist, but as a summary artefact rather than a primary one. Hiring decisions, project staffing decisions, and career development decisions will be made primarily against the skills graph.

◆ ◆ ◆

IV. The organisation as a fluid composition

The organisation that emerges from this transition is not a hierarchy of fixed roles connected by reporting lines. It is a continuously composed assembly of capabilities — humans plus AI agents — organised around current outcomes, with the composition shifting as the outcomes shift.

In practice, this looks like a small core of permanent leadership accountable for outcome ownership, surrounded by fluid teams that come together for the duration of a specific outcome and disband when the outcome is achieved or redirected. The teams are smaller than their predecessors — often a third or a quarter the size — because the agents absorb the bulk of execution work that used to require additional human capacity. The teams are also faster, because the iteration loops are shorter when the agents are doing most of the construction work.

The leader's task in this organisation is not primarily to manage people. It is to compose teams, set context, define accountability, and verify outcomes. Coaching matters more. Composition matters more. Direct supervision matters less.

Work Redefined: From / Through / To

The transition from era two to era three is the significant one.

DIMENSION	FROM (pre-2020)	THROUGH (2020–2026)	TO (2027–2030+)
Roles	Fixed roles, fixed titles	Hybrid; manager as coach	Fluid teams around outcomes
Where work happens	Office	Hybrid; remote-first for many	Anywhere; spatial-collaborative
Performance review	Annual, top-down	Continuous, peer-included	Outcome-based, real-time
Career model	Single ladder, 40 years	Lattice, multiple moves	Portfolio of capabilities
Skills representation	Job title and CV	CV plus LinkedIn	Skills graph, continuously updated
Team composition	Humans only	Humans + tools	Humans + agents, fluidly composed
AI involvement	None	Generative augmentation	Agentic execution; humans set context

FIGURE 14.1 · Three eras of work. The shift to era three is the inflection.

Figure 14.1

◆ ◆ ◆

V. The implications for India's workforce

India has an unusually strong position in era three of work, and an unusually exposed one as well. The strong position is that India is the world's largest source of knowledge workers entering the labour market, that it has built extensive digital infrastructure for skills validation, that its workforce has demonstrated remarkable adaptability across previous transitions, and that the country's frugal-innovation DNA is well-suited to the cost economics of agentic-AI-augmented work.

The exposed position is that the largest single category of Indian employment is the service-economy middle, much of which is the kind of structured, repeatable, transactional work that agentic AI does well. The displacement of this work, if it happens too fast, is a social-cohesion problem of significant scale. The leader's responsibility — and the country's policy responsibility, which I treat in Chapter Thirty — is to slow the displacement curve and

accelerate the reskilling curve so that the two cross sustainably rather than catastrophically.

The bank's HR director, by the time we ended our walk along the organisation chart, had committed to a reskilling programme for the forty-three role-holders whose jobs were structurally redundant. The programme was expensive. It worked, mostly. Twenty-eight of the forty-three were re-platformed into roles they could not have held before. Eight retired with dignity. Seven left the bank for other work. None was discarded. The programme was the right thing to do morally and the right thing to do commercially.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify three roles in your organisation whose primary function is information shuttling between layers.**

   Re-architect or remove. Either action is acceptable. The status quo is not.

2. **Begin building a skills graph for your top two hundred people, parallel to the CV.**

   The first version will be rough. The second version, six months later, will be useful. The third, a year later, will be transformational.

3. **Define one team this quarter as outcome-based rather than function-based.**

   Compose it deliberately. Let it dissolve when the outcome is achieved. Document what worked and what did not. The next composition will be better.

CHAPTER FIFTEEN

15

Learning Unbundled

The two-hundred-year-old factory model of schooling will not survive this decade.

—

"Education is not the filling of a pail, but the lighting of a fire."

— *attributed to W. B. Yeats*

◆ ◆ ◆

I sat in the back of a grade-five classroom in Lucknow in 2024, observing a daughter of one of my colleagues at work on a tablet. She was learning algebra. The tutor on her tablet was an AI that addressed her by name, that knew her boredom thresholds better than her parents did, that taught her in the Hindi-English mix she found comfortable, that explained quadratic equations using cricket statistics because the AI had noticed she liked cricket, and that paused, gently, when her attention waned, to ask whether she would prefer a short break or a different kind of problem. I watched her, on and off, for twenty minutes. She was, in the technical sense, learning. She was also enjoying herself, in a way that I had not seen children enjoy mathematics in any classroom I had visited in the previous decade.

I went home that evening and tried to explain to my own daughter what I had seen. She, then in college, listened politely, and then

said: "Papa, that was my entire childhood. Nothing happened. The teachers stood at the front of the class. They talked. We took notes. We took exams. None of it was personal. It was the same for everyone whether we were ready or not." She was, of course, exaggerating somewhat. There had been good teachers. There had been moments of personal attention. But the structural shape of her schooling — the standard curriculum, the class of forty, the time-bound progression, the memorisation reward — had been the dominant fact of her education. The grade-five girl in Lucknow had a different dominant fact.

◆ ◆ ◆

I. The factory model and what it produced

The model of schooling that dominated the past two centuries was, structurally, a factory. It was designed in the late nineteenth and early twentieth century, in tandem with industrialisation, to produce a particular kind of citizen and worker. The citizen was literate, numerate, broadly informed, capable of receiving instructions, willing to work in coordinated groups, and trained to defer to expertise. The worker was the same. Both were produced by a standardised curriculum delivered to age-graded cohorts in time-bound classes by a single teacher, evaluated by standardised examinations, and certified by degrees that signaled completion.

The factory model worked well, by historical standards. It produced widespread literacy. It produced the workforces of the industrial economies. It produced an extraordinary expansion of opportunity for people who would, in earlier eras, have had no access to formal education at all. It is not a model to be condemned. It is a model that was extraordinary for its time and that has, for some time now, been outliving its original purpose.

The original purpose has changed. The economy of the convergence decade does not need standardised workers. It needs creators, integrators, judgment-makers, contextual problem-

solvers, lifelong learners, and humans capable of working alongside intelligent systems in roles that no curriculum committee could have anticipated five years before. The factory model, optimised to produce a kind of citizen the economy no longer needs, is now actively miscalibrated. The students it produces, on average, are worse-prepared for the work that exists than they were thirty years ago, even as the curriculum they receive has remained largely the same.

◆ ◆ ◆

II. Five shifts that constitute the new model

There are, by my count, five structural shifts that the new model embodies. Each of them is happening, in 2026, somewhere in the world. None of them is yet at scale in mainstream public education. The work of this decade is to bring them to scale.

Standard curriculum becomes adaptive learning paths. The old model assumed every student needs the same content in the same sequence. The new model recognises that students have different starting points, different interests, different learning rates, and different optimal pathways through a body of knowledge. AI tutors construct these pathways dynamically, ensuring that no student is bored by content too easy and no student is lost by content too hard. The pathways converge on agreed mastery standards but diverge on how each student reaches them.

The class of forty becomes one-on-one AI tutoring with cohort-based human community. The old model concentrated one human teacher's attention across forty students for a fixed period. The new model gives each student continuous one-on-one tutoring from an AI that knows her, while preserving the cohort experience — the friendships, the discussions, the shared projects, the social-emotional development — through human teachers and small-group work. The teacher's role shifts from primary instruction to curation, mentoring, and the parts of teaching that AI cannot do.

Memorisation becomes contextual reasoning. The old model rewarded the retention of facts. The new model rewards the ability to reason about novel problems using whatever information is available. The skill being trained is not recall but judgment — what to ask, what to trust, what to do, in conditions where the facts are infinite but the questions are not.

Degrees become verifiable skills graphs. The credential that signals what a student knows is no longer principally the degree from a particular institution at a particular point in time. It is the continuously updated skills graph that captures what the student can demonstrate, validated by the institutions, the employers, and the AI systems that have observed her work directly.

School-then-work becomes lifelong learning. The old model assumed a discrete period of education followed by a discrete period of work. The new model assumes that learning continues across the entirety of working life, that workers cycle in and out of formal learning multiple times across decades, and that the employer, the institution, and the worker share responsibility for keeping the worker's capabilities current.

[Figure 15.1 — to be supplied by illustrator. THE FACTORY MODEL vs THE ADAPTIVE MODEL. A two-column comparison table. Left column: factory model attributes (standard…]

◆ ◆ ◆

III. India's leapfrog opportunity

India has a particular, asymmetric opportunity here. The factory model of schooling is most entrenched in the countries where it has been operating longest — Western Europe, North America, parts of East Asia. The institutional inertia in these countries against re-platforming education is significant. Existing institutions have buildings, faculty, accreditation systems, parental expectations,

and budget structures organised around the factory model. The transition is harder for them, not easier.

India's institutional inertia is also significant, but the country has two structural advantages that the older systems lack. The first is the scale of unmet demand: hundreds of millions of children for whom the factory model is currently not delivering even what it promised, and for whom the new model offers a path to genuine education. The second is the country's demonstrated capacity for leapfrog adoption — the way it skipped landline telephony for mobile, skipped retail banking for digital banking, skipped legacy media for streaming, skipped credit cards for UPI. Education can be one more leapfrog, if the institutional will is there.

The leapfrog requires three things: serious public investment in the underlying infrastructure (devices, connectivity, AI tutoring platforms in Indian languages), serious institutional redesign of the public school and university systems, and serious engagement from the private sector. Each is happening in patches. None is yet happening at the scale the opportunity requires. This is, in my view, one of the most consequential public policy choices India will make in this decade.

◆ ◆ ◆

IV. The institutional choice for leaders

For the leader who runs an educational institution, the question is whether to retrofit existing structures or build new ones alongside them. My honest recommendation, based on watching dozens of institutions attempt this transition, is to build alongside. The retrofitting tends to produce a hybrid that has the worst features of both models — the rigidity of the factory model with the cost of the new technology, without the benefits of either. The building-alongside tends to produce, over five to seven years, a parallel structure that can absorb the existing institution's strengths while leaving behind its constraints.

For the leader who runs a corporate function, the question is similar. The corporate university, the in-house training programme, the L&D function — these were built for the factory model of learning, and they are increasingly inadequate for the convergence-era workforce. The redesign is the same: build alongside rather than retrofit.

For the parent or the learner herself, the question is whether to wait for the institutions to redesign themselves or to begin learning by the new model directly. My honest recommendation is the latter. The new model is available now. The platforms are improving every quarter. The institutions will, in time, catch up, but the time may be longer than the years a particular learner has at her disposal.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit the most recent learning programme in your organisation against the five shifts.**
   Where does it conform to the factory model? Where does it conform to the adaptive one? The audit is the start of the redesign.

2. **Sponsor one cohort of underserved learners through an AI-native programme this year.**
   The cohort is small. The cost is bounded. The learning — yours, not theirs — is significant.

3. **Stop hiring on degree filters for at least one role this year.**
   Replace the degree filter with a skills graph filter or a demonstrated-capability assessment. The candidates who emerge will be better-suited and more diverse than the ones the degree filter was producing.

CHAPTER SIXTEEN

16

Healthcare, Continuous and Personalised

From treating disease to managing health.

—

"The greatest medicine of all is teaching people how not to need it."

— Hippocrates, more or less

◆ ◆ ◆

In a small clinic on the outskirts of Lucknow, in late 2025, I watched Dr. Anjali Mishra — a friend of mine, a general physician — see sixty patients in a single day. The figure was, she told me afterward, three times what she had been seeing in 2022. She was less tired at the end of it than she had been on a thirty-patient day three years earlier. The patient interactions had been longer than they used to be, not shorter. She had taken a forty-minute lunch break, which she had not done in years. The clinic had treated more patients, more thoroughly, with more attention, with less burnout.

She walked me through the day's tooling. The intake AI, in conversational Hindi or English, had taken each patient's history before the patient sat down with her. The wearable data had streamed from the patient's device, with her consent, into a dashboard Anjali could glance at in five seconds. The differential diagnosis suggestions had been pre-computed and clearly labelled.

The transcription had been ambient. The prescription drafting had been one click. The follow-up scheduling had happened automatically. The billing had reconciled in the background. None of these tools had replaced her judgment. Each had absorbed a layer of administrative work that used to interrupt her attention from the patient. Her capacity had tripled. So had the patient experience.

◆ ◆ ◆

I. From episodic to continuous

The twentieth-century model of healthcare was episodic. The patient was healthy until she became symptomatic. She presented at a clinic or a hospital with a complaint. The provider diagnosed and treated. The patient went home. She returned only when the next symptom appeared. The relationship between patient and provider was a series of discrete encounters with long, dark gaps between them, during which the provider had no visibility into the patient's health and the patient had no continuous support.

The model that is replacing it — call it continuous health — assumes that the patient is, in some sense, always in the system. Her wearable streams vitals. Her intelligent home monitors her sleep, her movement, her air quality, her hydration. Her AI health companion engages with her daily — sometimes minimally, sometimes in depth — about how she is feeling, what she is eating, what she is concerned about, what is changing. The clinical encounter, when it happens, is a punctuation in a continuous record rather than a fresh start from an empty file. The provider's attention, when it is given, is given in the context of months or years of data that previous-era providers did not have.

This shift has consequences across every dimension of the system. Diagnostics shift from reactive to predictive: the AI analysing the patient's continuous data flags the irregularity weeks before it becomes a symptom. Treatment shifts from one-size-fits-all to

personalised: the recommendation is calibrated to the patient's specific physiology, history, and context, drawing on AI models that can integrate orders of magnitude more variables than a human physician can hold in mind. Care shifts from clinic-bound to home-bound: many encounters that used to require a clinic visit no longer do, because the wearable, the AI, and the at-home diagnostic have absorbed the function. And the relationship shifts from transactional to continuous: the patient is not seeing a stranger in a fifteen-minute slot but interacting with a system that knows her over years.

◆ ◆ ◆

II. The care experience, not the procedure

The category re-statement that I argued for in Chapter Three applies sharply to healthcare. The old hospital sold a procedure: an appointment, a surgery, a test, a discharge. The new healthcare provider sells the patient's whole arc of relationship with her own health — from prevention through detection, treatment, recovery, and continuing wellness. The procedure is one event inside a much longer, more deliberate, more sustained experience that the provider is now responsible for designing and delivering.

This is, in 2026, the single biggest blind spot of the Indian private healthcare industry. The country has built a remarkable infrastructure of high-quality hospitals, often run by genuinely talented clinicians, often delivering procedures of world-class standard at a fraction of Western costs. The infrastructure is excellent. The category re-statement has not happened. Patients still arrive sick, get treated, and leave to fend for themselves. The relationship between hospital and patient ends at the discharge desk. The continuous-health competitor — and there will be such competitors, many of them — will eat the procedure-bound hospital's revenue from the high-margin top, by selling the relationship that comes before and after the procedure.

The hospital that re-platforms early — that builds a continuous-health offering alongside its procedural one, that treats the procedure as a punctuation in the patient's relationship rather than as the relationship itself — will own the patient's lifetime healthcare value. The hospital that does not will, by 2032, find that its procedural business has become a low-margin commodity, with the high-margin relationship value captured by digital-first competitors.

◆ ◆ ◆

III. The longevity revolution

Layered on top of the continuous-health shift is the longevity revolution that programmable biology now makes possible. In the past five years, the science of healthy lifespan — distinct from total lifespan — has progressed from a contested research field to a domain with a small but credible set of evidence-based interventions. The AI-augmented investigation of aging biology has accelerated. The first generation of longevity-focused clinics is now operating in major cities globally, including in India. The proposition is no longer "extend your life by three years." It is "extend the years in which you are functionally healthy by ten or fifteen."

The proposition resonates particularly strongly in India for cultural reasons. Indian families take seriously, as a matter of multi-generational responsibility, the question of whether parents and grandparents will be able to live in dignity, function independently, and remain present in family life into their late eighties and nineties. The longevity revolution speaks directly to this concern. The market for evidence-based longevity services in India is, in my view, one of the largest under-recognised opportunities of the next decade.

[Figure 16.1 — to be supplied by illustrator. FROM EPISODIC TO CONTINUOUS CARE. A two-column comparison. Left: episodic model attributes (reactive, clinic-bound, gene…]

◆ ◆ ◆

IV. The trust question

Healthcare is, of all the domains in this book, the one in which the trust currency is most decisive. The patient is, in any healthcare interaction, vulnerable. She is asked to share her most intimate data. She is asked to accept advice about her body. She is asked to undergo procedures she does not fully understand. The whole transaction depends on trust — in the provider, in the institution, in the technology, in the data handling, in the long-term incentives of the system.

The trust question becomes harder, not easier, as healthcare becomes more AI-augmented. The AI has access to more of the patient's data. The AI's recommendations are less interpretable to the patient. The AI is, in some institutional contexts, optimising for things the patient does not fully see. The provider's role, in this context, becomes increasingly that of the trust steward — the human who guarantees, by her training and her relationship with the patient, that the AI's recommendations are aligned with the patient's interests.

The institutions that will thrive in continuous health are the institutions that take this trust stewardship seriously and visibly. They will publish, openly, what data they collect and how they use it. They will document, transparently, how their AI systems work and what they optimise for. They will compensate their providers in ways that align with patient outcomes rather than procedure volumes. They will, in short, design their entire operation around the trust currency, recognising that in healthcare more than in any other category, the brand that earns trust durably owns the future.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you run a healthcare institution, define your continuous-health offering this year.**

   Even a minimal version — a wearable program for high-risk patients, an AI follow-up service, a longitudinal health record made available to the patient — begins the migration.

2. **Audit your patient experience for the moments after the procedure.**

   Most institutions end their attention at discharge. The patient's experience continues for years. The institution that earns those years owns the patient.

3. **Publish, openly, your data and AI practices to your patients.**

   The publication is the act of trust-building. Institutions that cannot publish do not yet have practices worth publishing. Institutions that can, deepen their trust currency by the publication itself.

CHAPTER SEVENTEEN

17

The Home That Lives

Real estate becomes the most under-priced category in the convergence economy.

—

"We shape our buildings; thereafter they shape us."
— Winston Churchill

◆ ◆ ◆

A real-estate developer who has been a friend and a client for nearly fifteen years told me, over a long lunch in Lucknow in 2024, that he was beginning to wonder whether he was in the right business. He had built, by that point, more than three thousand apartments across north India, mostly in the affordable and mid-premium segments. Each of them had been sold once, to its first buyer, after which the developer's relationship with the apartment had largely ended. He had begun, in his late fifties, to look at the arithmetic of this differently. He calculated that the average apartment, sold for one crore in 2024, would generate household expenditures across its first thirty years of life of approximately fifteen crore — utilities, maintenance, services, repairs, security, hospitality, entertainment, and the various technologies the household would consume within the apartment's walls. He had captured one crore. The other fifteen had gone to other firms.

He wanted, he said, to figure out how to capture more of the fifteen, and he wanted to do it not by gouging the homeowner but by genuinely improving the experience of living in the apartment over its life. The conversation we had that afternoon became the foundation of a project we have worked on together for the past two years. The project is one of the most interesting things I have ever worked on, and the rebuild it has produced is, I believe, the template for the next generation of real estate.

◆ ◆ ◆

I. The category restated

The old real-estate category was the apartment. The unit of value was square footage, configured in a particular layout, in a particular location, at a particular price. The transaction was the sale, executed once. The developer's revenue ended at the registration. The homeowner's expenditure on living in the apartment continued for thirty years, but the developer had no further share of it.

The new real-estate category is the living environment. The unit of value is the integrated experience of the home across the years the household occupies it. The transaction is the beginning of a multi-decade relationship in which the developer continues to participate in the household's expenditure on its own living, in exchange for continuously improving the experience of living. The square footage is one component of the offering. It is no longer the offering itself.

The shift, when stated this way, sounds modest. Operationally it is a complete re-platforming. The developer's organisation, designed to plan, build, sell, and walk away, now has to be redesigned to plan, build, sell, *and continue to deliver* across the household's occupancy. The skills required are different. The capital structure is different. The customer relationship is different. The metrics are different. Most existing developers cannot make this transition

without redesigning their entire enterprise. The ones that succeed will own the next twenty years of the category. The ones that do not will become low-margin construction firms competing increasingly on price with whoever else can build a building.

◆ ◆ ◆

II. What the home that lives actually does

Six layers of capability, when fully integrated, constitute what I have come to call the home that lives. None of them is exotic in 2026. The integration of all six is, however, exotic, and remains so.

Layer one: an integrated sensor mesh. Air quality, temperature, humidity, light, occupancy, sound, water, energy. The mesh is not visible to the residents; it is built into the fabric of the apartment. It generates a continuous, granular picture of the home's environment.

Layer two: an intelligence layer that learns the household. The patterns of when each member is home, when each is away, what the household's preferences are at different times of day and across the week. The learning is local; the data does not leave the home unless the household chooses to share it for a specific service.

Layer three: an actuation layer that acts on the household's behalf. The lighting, heating, cooling, and air-purification respond automatically to predicted needs. The morning routine is anticipated. The evening routine is anticipated. The household does not configure the home; the home configures itself around the household.

Layer four: a service marketplace integrated into the home. Cleaning, cooking, maintenance, deliveries, healthcare visits, tutoring sessions, repair services. The household orders services through the home's interface; the developer's platform takes a share of the recurring transaction. The services are vetted, rated, and

accountable to the platform, which gives the household a quality assurance that the open market does not.

Layer five: a community layer. The building's residents are connected through a community platform that supports shared resources, neighbourhood events, mutual help, local commerce, and emergency coordination. The platform turns the building from a collection of strangers into a functional community, which is genuinely valuable — and rare in modern Indian urban life.

Layer six: a continuous improvement layer. The home's software and capabilities are upgraded continuously, the way a software product is upgraded, not the way a building used to be upgraded. The residents wake up, every few months, with new capabilities they did not have before, delivered remotely.

[Figure 17.1 — to be supplied by illustrator. THE HOME-AS-PLATFORM STACK. A six-layer pyramid: sensor mesh, intelligence, actuation, services, community, continuous i…]

◆ ◆ ◆

III. The economics

The developer's revenue model, in this configuration, includes the original sale plus a recurring fee for the home-platform services across the household's occupancy. The recurring fee is modest in any single month and substantial across thirty years. The platform also takes a share of the service marketplace transactions that flow through it, which compounds the recurring revenue. By the developer's own modelling on the project we worked together on, the lifetime revenue per household is between three and five times the original sale price, with the recurring component contributing the larger share. The margins on the recurring revenue are higher than the margins on construction. The capital required for the recurring business, once the building is built, is small.

The arithmetic is, in other words, dramatic. The developer who continues to sell square footage in 2030 will be running a low-margin commodity business. The developer who has re-platformed to the home-as-platform model will be running a high-margin recurring-revenue business with construction as one of its inputs rather than its output. The two firms will look entirely different. They will value entirely differently in capital markets. They will compete entirely differently for talent.

◆ ◆ ◆

IV. The hardest part: trust

The hardest part of building a home-as-platform is not the technology. The technology is available and is improving every quarter. The hardest part is the trust the household must place in the developer to manage their data, their privacy, their domestic intimacy, and their dependencies on the platform. The developer is asking the household for an unprecedented level of access to the home's interior life. The access is, when properly used, a benefit to the household. When improperly used, it is a violation.

The developers that succeed will be the ones that earn this trust deliberately, transparently, and verifiably. They will publish their data practices. They will give the household genuine control over what is collected and what is shared. They will build technical architectures in which the most intimate data does not leave the home. They will accept independent audits of their practices. They will, in short, treat the trust currency as a foundation rather than as a marketing claim.

The developers that do not will, sooner or later, encounter a violation — a leak, a misuse, a regulatory action, a public incident — that destroys the trust capital it took years to accumulate. The home-as-platform business, more than any other in this book, is one in which a single trust failure can end the entire enterprise. This is a feature of the model, not a bug. The leader who is not prepared

to manage trust as a fiduciary obligation should not enter the category.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you are a developer, identify one existing project that can be retrofitted into a home-as-platform pilot this year.**

   The pilot does not need to be at scale. It needs to be honest. The learning will inform every subsequent project.

2. **If you are a homeowner, audit your home for the six layers and prioritise the next one to add.**

   Even without a developer's involvement, much of the home-as-platform stack is now available to individual households. The early adopters are building the demand the developers will eventually serve.

3. **Begin treating the homeowner not as a one-time customer but as a thirty-year relationship.**

   The accounting changes. The CRM changes. The capital allocation changes. The orientation of every team in the company changes. This is the migration. Begin it.

CHAPTER EIGHTEEN

18

Money, Identity, Ownership

The trinity that decides the next economy.

—

"Money is a record of human activity.
Identity is its custodian. Ownership is its boundary."
— A line I wrote in a notebook some years ago, attribution unclear

◆ ◆ ◆

I had a conversation in Bengaluru, in late 2024, with the chief executive of one of India's better-known fintech firms. He had been thinking, he said, about why his industry had grown so fast in India over the past decade and so slowly in most other markets. The conventional answers — UPI, Aadhaar, low cost of acquisition, demographics — were all true but incomplete. The deeper answer, he had come to believe, was that India had, perhaps inadvertently, built a coherent infrastructure for the three things that the next economy would actually run on: a digital identity layer, a digital money layer, and the beginnings of a digital ownership layer. The three were not always built with the convergence economy in mind, but they had ended up being the substrate on which the convergence economy could be built.

The conversation has stayed with me. I have come to believe he was right. The three — money, identity, ownership — are the trinity on which the next economy will be built, and the country

that has built them best, with the right balance of openness and protection, will host disproportionately many of the next economy's businesses. India is, in 2026, in an unusually strong position on all three. This chapter is about why, and about what the leader needs to understand about each.

◆ ◆ ◆

I. Money becomes programmable

Money, in the world that is ending, was a passive instrument. It sat in an account. It moved from one account to another when the account-holder instructed it to. It did not, on its own, do anything. The intelligence around money — what to spend it on, when to invest it, whether to repay it — sat outside the money itself, in the head of the account-holder or in the systems she used.

Money, in the world that is beginning, is programmable. It moves on its own, under conditions the account-holder has specified. It allocates itself across investments based on rules the account-holder has set. It pays itself out at moments the account-holder has scheduled. It refuses transactions the account-holder has disallowed. It earns yield in instruments the account-holder has approved. The intelligence is no longer outside the money; it is in the money.

In India, this transition is being accelerated by the maturation of UPI, the Account Aggregator framework, the Open Network for Digital Commerce, and the progress on a central bank digital currency. The infrastructure is in place. The application layer that uses the infrastructure to deliver programmable money to the household, the small business, the corporation, and the government is now being built. The companies that build well, in this layer, will own a significant share of the next decade's financial services revenue.

◆ ◆ ◆

II. Identity becomes self-sovereign

Identity, in the world that is ending, was issued by institutions and held by them. The bank issued an account number. The government issued a passport. The school issued a certificate. The employer issued a credential. The individual was a passive recipient of these institutional records. To prove who she was, she had to ask the institution to vouch for her — and the institution might or might not, depending on its current priorities.

Identity, in the world that is beginning, is self-sovereign. The individual carries a portable, cryptographically verifiable record of who she is and what she has done. The records are issued by institutions, but they are held by her, in her own digital wallet, in standards that any other party can verify without contacting the issuing institution. She can present credentials selectively — proving her age without revealing her birth date, proving her income range without revealing her exact salary, proving her qualifications without exposing her academic transcript. The institution loses control over the use of the record. The individual gains it.

This shift is, in 2026, partly built and largely conceptual. Aadhaar is one of the foundational components, although not yet a self-sovereign one in the strictest sense. DigiLocker is another. The Verifiable Credentials standards, now globally adopted, are the protocol layer. The applications that use these layers to give the individual genuine self-sovereign identity are now in early stages. By 2030, in my forecast, self-sovereign identity will be the dominant pattern for credential management, at least in the digitally mature economies, and India will be among the leaders.

◆ ◆ ◆

III. Ownership becomes tokenisable

Ownership, in the world that is ending, was concentrated in fixed-share structures. The house was wholly owned. The company was owned in shares of standardised denomination. The art was owned by one collector. The intellectual property was owned by one firm. The transaction costs of fractionalising ownership were high enough that fractional ownership was, in most cases, not worth pursuing.

Ownership, in the world that is beginning, is tokenisable. Any asset — real estate, equity, art, intellectual property, contracts, royalties, future cash flows — can be represented as digital tokens of arbitrarily small denomination, traded on programmable infrastructure, with ownership records maintained on shared ledgers that no single party controls. The transaction costs of fractionalisation collapse. The pool of possible owners for any given asset expands enormously. The asset becomes more liquid, more accessible, and more efficient.

The implications are larger than they at first appear. Real estate fractionalisation alone, if it scales, could redistribute participation in property wealth from the few who can buy a whole apartment to the many who can buy a thousand-rupee fraction of one. Equity fractionalisation could allow the small saver to participate in the high-growth start-up at a unit cost that current securities regulation makes impossible. Intellectual-property fractionalisation could allow the songwriter to retain micro-shares of her catalogue while still raising capital. Each of these implications threatens an existing intermediary — the property broker, the broker-dealer, the music publisher — and creates an opening for new categories of business.

The Indian regulatory posture on tokenisation, in 2026, is cautious but improving. The Reserve Bank's pilots on tokenised deposits and tokenised bonds are encouraging. Regulatory clarity will, in

my forecast, arrive by 2027 or 2028 in serious form. The companies that have done the work of building tokenised-ownership infrastructure during this regulatory ambiguity will be the ones positioned to scale when the clarity arrives.

[Figure 18.1 — to be supplied by illustrator. THE TRINITY OF THE NEXT ECONOMY. A triangular diagram with three vertices: money (programmable), identity (self-sovereig…]

◆ ◆ ◆

IV. Why the trinity matters together

The three are not independent. They reinforce one another. Self-sovereign identity allows individuals to participate in tokenised-ownership markets without surrendering personal data they should not need to. Programmable money allows tokenised assets to settle and yield without the friction of traditional banking. Tokenised ownership creates new instruments that programmable money can flow through and new claims that self-sovereign identity can verify.

When the three are mature and integrated, the resulting economy looks meaningfully different from the one we have. The household can hold programmable money, self-sovereign identity, and fractional ownership in dozens of assets, all in a single integrated wallet, all under the household's control. The friction of participating in capital markets, of accessing credit, of monetising one's own data, of receiving payments for one's own work, drops by an order of magnitude. The categories of small business that become viable expand. The categories of investment that become accessible to the median household expand.

The companies that build well in this trinity — which is to say, the companies that can hold money, identity, and ownership coherently for their customers — will be the platform companies of the next decade. They will be banks in the broadest sense, but they

will not look like the banks of 2025. They will look like integrated life-platforms in which the financial layer is one component of a continuous relationship.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit, in your own organisation, where you stand on each layer of the trinity — money, identity, ownership.**

   Most organisations have made progress on one and lagged on the others. The audit makes the lag visible.

2. **Identify one customer experience that would be transformed by treating money as programmable rather than passive.**

   The redesign begins with the experience, not the technology. The technology will follow.

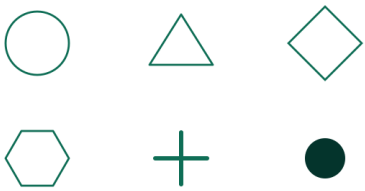
3. **Engage with the regulatory development of self-sovereign identity and tokenised ownership in India.**

   The regulations of the next two years will shape the next twenty. The leaders who participate in shaping them will be better-positioned than the leaders who wait for them to be finalised.

PART FIVE

V

The New Product Categories



Six categories that do not exist as distinct industries in 2026 and will be obvious by 2032.

THE WORLD THAT'S BEGINNING

Six categories that do not exist as distinct industries in 2026 and will be obvious by 2032.

CHAPTER NINETEEN

19

Six Categories That Don't Yet Exist

The taxonomy of what the convergence is producing.

—

"The future is already here — it's just not very evenly distributed. And it's not always called what it will eventually be called."

— A loose paraphrase, again, of Gibson

◆ ◆ ◆

Every technological epoch has produced its own category vocabulary. The industrial epoch gave us the corporation, the factory, the railway, the assembly line, the mass-market brand, the chain store. The information epoch gave us the personal computer, the internet, the search engine, the social network, the smartphone, the app store. Each of these category names is now so familiar that we forget there was a moment when it did not exist as a category — when the thing it now names was an unnamed combination of older categories, and when the people building it were not yet sure what to call it.

We are in such a moment now. The convergence of AI, spatial computing, sensors, biology, and energy is producing categories of product that do not yet have stable names, that are not yet recognised as distinct industries, and that the leaders of 2026 are still describing in awkward circumlocutions because the proper category names have not yet been adopted. This part of the book is

my attempt to name them, to describe them, and to argue for why each will be a several-billion-dollar industry by 2032.

I count six categories. The number is somewhat arbitrary; another writer might draw the lines differently and produce a different count. But the six I will name in the chapters that follow are, in my judgment, the categories that have the clearest internal logic, the largest addressable markets, and the most distinctive set of capabilities required to compete in them. They are: sentient spaces, always-on companions, generative worlds, molecular services, the provable-truth layer, and adaptive everything.

◆ ◆ ◆

I. A brief on each

Sentient spaces are physical environments — homes, offices, retail spaces, vehicles, public spaces — that have been instrumented with sensor meshes, AI intelligence, and actuation capabilities such that they sense, learn, and respond to the humans within them. The home that lives, which I described in the previous chapter, is one example. There will be many.

Always-on companions are AI assistants of a kind we do not yet have at scale: persistent across the user's life, accessed continuously across devices, knowing the user over years rather than over a single session, and serving as the user's primary interface to most of the digital world. The chat-based AI of 2024 and 2025 is the prototype. The mature version is something more like a digital second self.

Generative worlds are AI-generated environments — for entertainment, learning, work, simulation — that adapt to the user, are produced at unit costs that would have been impossible five years ago, and increasingly blur the line between game, film, education, and conversation. The metaverse, properly understood, is generative worlds.

Molecular services are healthcare, agriculture, materials, and consumer products that are designed and delivered at the molecular level, customised to the individual user, drawing on AI-assisted molecular design and on the rapidly falling cost of biological and chemical synthesis.

The provable-truth layer is the infrastructure — cryptographic, social, institutional — that allows information, identity, content, and credentials to be verified as authentic in a world where most digital content is, by default, synthesisable. In an age of indistinguishable fakes, the ability to prove what is real becomes a structurally valuable capability.

Adaptive everything is the cross-cutting category in which existing products of every kind acquire intelligence and adaptation as a default feature — not because the product was designed as an AI product, but because, by 2030, a product that does not adapt to its user will be perceived as broken in the way an unresponsive product is now.

[Figure 19.1 — to be supplied by illustrator. SIX NEW CATEGORIES. A 2x3 grid with one cell per category. Each cell with the category name, a one-line definition, a re…]

◆ ◆ ◆

II. How to read these chapters

Each of the next six chapters treats one category in depth. The structure of each chapter is the same. I describe what the category actually is, in operational terms. I name representative examples that exist in early form in 2026. I describe the technological capabilities that have to converge for the category to mature. I assess the addressable market and the competitive dynamics. I name the firms that are best-positioned to compete and the kinds of capability the leader of an aspirant firm needs to build. I close with three decisions.

The chapters are intended to be read in sequence, but each can be read independently. The chapters that are most relevant to a given reader will depend on which industry she operates in. The leader in real estate will find most use in sentient spaces. The leader in consumer tech will find most use in always-on companions. The leader in entertainment, education, or simulation will find most use in generative worlds. The leader in pharmaceuticals, agriculture, or materials will find most use in molecular services. The leader in media, government, financial services, or any trust-based business will find most use in the provable-truth layer. And every leader, regardless of industry, should read the chapter on adaptive everything, because every product in every category is becoming subject to its logic.

◆ ◆ ◆

III. A note on positioning

A note worth offering before I describe the categories in detail. The leader who reads these chapters will be tempted, naturally, to ask which category her firm should enter. The question is reasonable. The answer, in most cases, is not "enter as a competitor in one of the new categories." The answer, in most cases, is "absorb the relevant capabilities of one or more new categories into the firm's existing positioning." The leader of an established firm in established categories has assets — distribution, customer relationships, brand — that pure new-category competitors do not have. The right strategic posture, in most cases, is to use those assets as the foundation on which to build the new-category capability, rather than to spin up a separate new-category venture from scratch.

This is a defensive observation as much as an offensive one. The leader who waits for one of the six new-category competitors to enter her industry as a disruptor will, in most cases, lose. The leader who absorbs the relevant new-category capabilities into her own

firm before a disruptor arrives, in many cases, can prevent the disruption entirely. The defence and the offence, in this strategic moment, are the same act.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Read the next six chapters and identify the two categories most directly relevant to your industry.**

   The answer is rarely six. It is rarely zero. It is, for most leaders, two or three. Identify them.

2. **For each of those two or three categories, name the capability your firm should be building this year.**

   Not next year. This year. The capability takes time. The window for starting closes faster than the window for finishing.

3. **Identify the firm in your industry that is most likely to enter your category as a disruptor from one of the six new categories — and decide your defensive or offensive posture toward them now.**

   The decision will inform every capital allocation choice for the next thirty-six months.

CHAPTER TWENTY

20

Sentient Spaces

Buildings that sense, learn, and respond.

—

"A room is not a room. It's a relationship."
— A line from a designer I once worked with, whose name I have, regrettably, forgotten

◆ ◆ ◆

When the chairman of a hospitality group I have advised invited me to inspect a property he was designing in Goa, in early 2025, he did not show me the rooms. He showed me the building's intelligence stack. He had taken the unusual step of designing the building's sensor and intelligence architecture before the architects had finalised the room layouts, on the theory that the building's *behaviour* — how it would respond to its guests — was the actual product, and the rooms were simply the physical containers for the behaviour. The architects had not been entirely happy with this inversion of the usual sequence, but they had done the work, and the property that opened ten months later was, in my view, one of the most genuinely sentient spaces I have ever stayed in. The room knew when I had returned from dinner. The lighting adjusted. The temperature adjusted. The minibar showed the items most likely to interest me at the time of evening I had returned. The morning notification came in my preferred language. The departure was

anticipated. None of it felt invasive. All of it felt, the way good service feels, simply attentive.

This chapter is about that category of building.

◆ ◆ ◆

I. The category, defined

A sentient space is a physical environment — a home, an office, a retail space, a hotel, a hospital, a vehicle, a public venue, a classroom — that has been instrumented with three integrated capabilities: a sensor mesh that perceives the environment and the humans within it, an intelligence layer that interprets and learns from what the mesh perceives, and an actuation layer that responds, in real time, to the humans within the space. The three together produce a building that, in the experiential sense, is alive — not metaphorically, but in the operational sense that it senses, considers, and acts.

The category is, in 2026, present in early form across several adjacent industries. The home-as-platform I described in Chapter Seventeen is one form. The smart hotel of the kind the Goa property exemplifies is another. The connected hospital that monitors patients ambient through their stay is another. The sensor-equipped retail space that adapts displays, music, and assistance to the customer in front of it is another. Each of these is a distinct application, but the underlying capability stack is the same, and the firms that build the stack well will be the platforms on which all the applications run.

◆ ◆ ◆

II. The technology stack

The sensor mesh is the foundation. Air quality sensors, temperature, humidity, light, occupancy, presence, sound, motion, energy use, water use, video where appropriate. The sensors are

small, cheap, and increasingly invisible. By 2026, the cost of a comprehensive sensor mesh for a typical home or hotel room has fallen below the cost of the lighting fixtures it integrates with. By 2030, the cost will be a rounding error in any new construction.

The intelligence layer is the differentiation. The sensor data, in raw form, is useless. The value is created by the AI that learns the patterns of the space and the humans within it — when the residents are home, when they are away, what they prefer at different times, how the building should be operated to be both pleasant and efficient. The intelligence is increasingly local, running on edge compute within the building rather than in distant data centres, both for latency and for privacy reasons. The intelligence is also increasingly federated — the building learns from its own occupants but contributes anonymised insights to a broader pool that improves the intelligence for all buildings on the platform.

The actuation layer is the visible result. Lighting that adjusts. Climate that adjusts. Music that adjusts. Information that surfaces when relevant. Suggestions that arrive at appropriate moments. Services that are pre-scheduled. Maintenance that is pre-emptive. The actuation, when done well, recedes into the experience rather than calling attention to itself. The guest does not notice the building responding. The guest notices that the building feels right.

◆ ◆ ◆

III. The hard problems

The hard problems in sentient spaces are not principally technical. The technology is largely available. The hard problems are operational and ethical.

The first hard problem is *what to optimise for*. A sentient space can be optimised for the comfort of its occupants, for the energy efficiency of its operation, for the revenue of its operator, for the

data it can extract, or for any of dozens of other goals. The choice of optimisation target shapes everything about how the space behaves. Spaces optimised for occupant comfort feel different from spaces optimised for operator revenue, even when the underlying technology is the same. The leader who deploys a sentient space without making the optimisation choice deliberately will end up with a space optimised by default for the first metric the operations team measures, which is rarely the right choice.

The second hard problem is *trust*. A sentient space is, by definition, observing its occupants. The occupants are entitled to know what is observed, how the data is used, who has access to it, and what choices they have to opt out. The legal frameworks for this are still developing in most jurisdictions; the ethical obligations are larger than the legal obligations and arrive earlier. Operators who treat their occupants' privacy as a fiduciary obligation will earn trust durably. Operators who treat it as a regulatory minimum will lose trust the first time something goes wrong, and the loss will be hard to recover from.

The third hard problem is *the human in the loop*. Sentient spaces, when they fail — and they will fail — fail in ways that humans cannot easily debug. The lighting suddenly behaves oddly. The temperature is wrong. The recommendation is offensive. The intervention is mistimed. The occupant has no easy way to figure out why. The operators must build, into the design of the space, both technical mechanisms for human override and operational mechanisms for human investigation when things go wrong. Spaces that lack these mechanisms feel, to their occupants, like opaque alien systems. Spaces that have them well-designed feel like attentive servants.

◆ ◆ ◆

IV. The market and the competitors

The addressable market for sentient spaces is, by 2032, in my honest forecast, north of half a trillion dollars annually globally — comprising new construction, retrofit of existing buildings, ongoing platform fees, and the service marketplace transactions that flow through the platforms. The Indian share of this market, given the country's construction velocity and demographic momentum, will be among the largest in the world.

The competitors will come from three directions. The first will be the large technology firms, who will offer horizontal sentient-space platforms that any operator can adopt. The second will be vertically integrated builders — the developers, hotel groups, retail chains, and hospital systems that build their own sentient capabilities into their physical assets. The third will be specialised platform firms whose entire business is the sentient-space stack and which service multiple verticals.

Each of these models has trade-offs. The horizontal tech platforms have the deepest technology but the weakest understanding of any specific vertical. The vertically integrated builders have the deepest vertical understanding but the weakest technology. The specialised platforms try to balance both, with mixed results. My honest forecast is that, by 2032, no single model wins; we will see successful firms in each of the three categories, with the choice of model determining which kinds of operators each firm serves best.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate a physical asset, identify the optimisation target you want a sentient capability to serve, and put it in writing.**

   The writing forces clarity. The clarity prevents the default optimisation that operations teams will otherwise impose.

2. **Pilot a sentient capability in one corner of one asset this year.**

   The pilot is not the full deployment. The pilot is the learning. The full deployment, if you do it later, will be informed by what the pilot revealed.

3. **Publish your data and intelligence practices to your occupants now, before you have to.**

   The publication is the act of trust-building. It is also rare, which is why it is differentiating.

CHAPTER TWENTY-ONE

21

Always-On Companions

The AI that knows you over years.

—

"A friend is one who has the same enemies as you have."
— Abraham Lincoln, who would have been bewildered by this category but might still have approved

◆ ◆ ◆

A senior executive at one of India's larger media houses showed me, in late 2025, the AI assistant she had been using continuously for fourteen months. The assistant, by that point, knew her in considerable depth. It knew her schedule, her contacts, her writing style, her values, her recurring concerns, her family situation, her professional obligations, the books she had read recently, the arguments she was working through. It greeted her in the morning by referencing specifically what she had told it the previous evening. It drafted, in her voice, the emails she had not yet had time to write. It anticipated the meetings of the week ahead and prepared the briefs she would have prepared for herself if she had had the time. It functioned, in her description, as a chief of staff she did not have to manage and that did not, ever, need a vacation.

She told me that she could no longer imagine her professional life without it. She also told me, in the same conversation, that she was uneasy about the depth of the relationship. She was not sure what

would happen if the company that operated the assistant went bankrupt, was acquired, or changed its terms of service. She was not sure what would happen if her relationship with the company ended for any reason. She had, she said, allowed an external party to become structurally important to her professional functioning, and she was not entirely comfortable with that.

The conversation captures, in miniature, both the promise and the unease of the always-on companion category — the second of the six new categories, and the one most likely to define the personal-technology landscape of the next decade.

◆ ◆ ◆

I. The category, defined

An always-on companion is an AI assistant that is present continuously across the user's life — across devices, across contexts, across years — that knows the user in considerable depth from accumulated interaction over time, that operates with substantial autonomy on the user's behalf, and that serves as the user's primary interface to most digital and many physical interactions. The category is defined less by any particular technical capability and more by *persistence and depth* — the assistant is not a tool the user picks up for a session; it is a relationship the user maintains across time, with the relationship deepening as the assistant accumulates context.

The category is, in 2026, present in early form. The chat-based AI assistants of 2023–2025 were the prototype. The agentic AI of 2025–2026 added the ability to act on the user's behalf. The next generation, arriving in 2027 and 2028, adds genuine persistence — the assistant remembers the user across years rather than across sessions, develops a model of the user that is genuinely individual rather than generic, and operates with the kind of deep contextual

understanding that distinguishes a long-time chief of staff from a temp.

◆ ◆ ◆

II. The capability stack

The capability stack of a mature always-on companion has four layers, and only the bottom one is widely available in 2026.

Layer one is the model itself. The reasoning, the language, the multi-modal capability. By 2026, the frontier models are sufficient for most companion functions. The frontier will continue to advance, but the bottleneck has moved up the stack.

Layer two is persistent memory. The companion remembers, across years, what the user has told it, what has happened in the user's life, what preferences the user has expressed, what patterns the user displays. The memory must be retrievable in context, summarisable when needed, and updateable as the user's life changes. The memory architecture is, in 2026, the area of most active research in companion design. By 2028, it will be a solved problem.

Layer three is autonomy with judgment. The companion does things on the user's behalf, but it knows when to ask, when to act, when to escalate, and when to refrain. The judgment is calibrated to the user's preferences, accumulated over time. The companion that errs on the side of asking too often is annoying; the companion that errs on the side of acting too autonomously is, eventually, a liability. Calibration is the differentiator between adequate companions and great ones.

Layer four is portability. The user's relationship with the companion is the user's, not the operator's. The user can move her companion — her data, her memory, her preferences, her trained behaviours — from one operator to another, the way she can move her bank accounts. The portability is necessary for the user's freedom and

for genuine competition in the category. By 2028, in my honest forecast, the regulatory framework for companion portability will be developed in multiple jurisdictions, with India and the European Union leading the work.

[Figure 21.1 — to be supplied by illustrator. THE FOUR-LAYER COMPANION STACK. A vertical pyramid: model (bottom), memory, autonomy, portability (top). Each layer with…]

◆ ◆ ◆

III. The trust question, sharpened

The trust question that runs throughout this book is, for always-on companions, sharpened to its limit. The companion knows, at maturity, more about the user than any other entity in the user's life. It knows things the user has not told her spouse. It knows the patterns the user has not noticed in herself. It has, in some sense, the most intimate relationship the user has with anything outside her own household.

The implications for the operators of companions are significant. The terms on which the user shares this depth of context have to be honest. The operator's incentives have to be aligned with the user's. The data architecture has to ensure that the companion's accumulated context is the user's property, retrievable and portable. The operator's business model has to be one that can survive transparent scrutiny. Operators that cannot meet these standards will, in time, be regulated out of the category or out-competed by operators that can.

The economic implications are equally significant. The relationship between user and companion is, once mature, more durable than any product relationship. The user does not switch companions casually, because the switching cost is the loss of years of accumulated context. The operator who has earned the user's depth, and has earned the right to keep her, has captured the most

durable customer relationship in the digital economy. The operator who has not earned that depth, or who loses the user's trust, has captured a transactional relationship that any competitor can disrupt.

◆ ◆ ◆

IV. The competitive dynamics

The companion category will, in my forecast, consolidate around a small number of platforms. The reason is structural: the value of a companion is largely a function of the depth of its memory of the user, which compounds over time. A user who has used a companion for three years will not switch to a competitor with no memory of her, even if the competitor's underlying model is somewhat better. The first-mover advantages, properly built, are durable.

The competitors in 2026 are largely the firms that built the foundational AI models — OpenAI, Anthropic, Google, Microsoft, Meta, plus the leading Chinese model providers. The competitors by 2032 will, in my forecast, be a smaller number — perhaps three to five global platforms — plus a set of regional or vertical specialists for specific use cases. Indian operators have, in 2026, no foundational-model presence at the global frontier; the path to Indian leadership in the companion category, if there is one, runs through specialisation, language coverage, and trust differentiation rather than through frontier-model competition.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Begin using an AI companion seriously yourself, with the same depth of context-sharing you would give a chief of staff.**

   The product cannot be designed by people who do not use it themselves. The leader who has not committed to her own companion is not equipped to make capital decisions in this category.

2. **If your firm is not an AI company, decide whether your customers will reach you through their companions and what that implies for your interface design.**

   By 2030, the user's primary interface to most services will be through her companion. The firm that has not designed for this is competing on inferior terrain.

3. **If your firm is in the customer relationship business — financial services, retail, media, healthcare — consider whether to build your own vertical companion or to integrate with horizontal ones.**

   The choice is consequential. Both have legitimate business models. Picking the wrong one for your firm's assets is expensive.

CHAPTER TWENTY-TWO

22

Generative Worlds

Environments produced on demand, adapting to the user.

—

"Stories happen to those who can tell them."

— Henry James, who would have understood instantly

◆ ◆ ◆

The director of one of India's larger streaming-content companies showed me, in 2025, an early-stage product his team had been developing for eight months. It was not a film. It was not a television series. It was, in his description, "a place" — a generated environment, set in a fictional Indian town in the 1990s, in which the viewer could move around, talk to characters, listen to ambient conversations, and become a participant in slowly developing storylines. The storylines were generated, in real time, in response to what the viewer chose to do. The characters were AI agents with internal lives, stable personalities, and accumulated memories of their previous interactions with the viewer. The town had its own economy, its own seasons, its own ongoing political subplot, its own romances and grievances. The viewer could enter the town for ten minutes or for two hours. The town continued to live whether the viewer was there or not.

He was, he told me, not yet sure what this product was. It was not a film, because there was no fixed narrative. It was not a game, because there were no win conditions. It was not a metaverse, because there was no social-multiplayer dimension. It was not television, because it was not episodic. It was a category that did not yet have a name, and his marketing team was, he said, struggling with how to describe it to investors and to potential users.

The category I describe in this chapter is the category that the director was, without knowing it, building. I call it generative worlds. By 2032, I expect this category to be one of the largest in the entertainment, education, training, and simulation markets — and to absorb significant portions of those markets entirely.

◆ ◆ ◆

I. The category, defined

A generative world is an interactive environment, produced largely or entirely by AI, that adapts in real time to the user's choices, persists across the user's sessions, and supports rich narrative, social, or experiential interaction within it. The generative world is distinct from the traditional video game in that the content is not pre-authored; it is generated as needed, often differently for each user. It is distinct from the traditional film in that there is no fixed sequence; the user moves through the world as she chooses. It is distinct from the social metaverse in that it does not depend on the presence of other users; the world is populated by AI agents who serve the experience the way actors serve a film.

The category, properly understood, is the natural product of three technological trajectories converging: the maturation of generative AI to the point where compelling content can be produced on demand, the maturation of agentic AI to the point where convincing characters can be sustained over time, and the

maturation of spatial and immersive computing to the point where rich worlds can be experienced viscerally. By 2026, all three trajectories have crossed thresholds at which generative worlds become commercially viable. By 2030, the category will be a mainstream entertainment medium. By 2032, it will be one of the largest categories in the digital economy.

◆ ◆ ◆

II. The applications

The applications of generative worlds extend well beyond entertainment. I will name five.

Entertainment is the obvious application. Films, games, and the new hybrid medium that the streaming director was building. The Indian content industry, with its strength in storytelling and its access to the convergence-era tools, is positioned to be a major producer in this category — if the industry's leaders make the strategic move from authored content to generative content within the next three years.

Learning is the application I find most exciting. The history class my fictional son took in Chapter Thirteen — walking through Mughal Lucknow as it was in 1722 — is a generative world used as a learning environment. The category, deployed at scale in education, can transform comprehension and retention in ways the factory model of schooling cannot. I treated this in Chapter Fifteen.

Training and simulation is the application with the most economic value in the near term. Pilots, surgeons, soldiers, plant operators, emergency responders, and corporate leadership teams can all be trained in generative worlds at unit costs and with skill-transfer outcomes that physical simulation cannot match.

Therapy and mental health is an emerging application that I am watching carefully. Generative worlds, used under the supervision of a trained clinician, can be used to expose patients to feared

situations under controlled conditions, to rehearse social skills, to process trauma in carefully designed environments, and to provide companionship in cases of severe isolation. The category requires significant ethical care; it also represents one of the most genuine therapeutic advances in mental health in a generation.

Social and cultural connection is a longer-term application. Generative worlds can host gatherings — religious, cultural, familial — that bring together people who are physically distant in shared, meaningful spaces. The Indian diaspora, distributed across every major economy, is particularly well-served by this application; a virtual Diwali gathering in a generative environment that resembles a beloved hometown can carry a cultural weight that a video call cannot.

◆ ◆ ◆

III. The production economics

The production economics of generative worlds are dramatically different from those of traditional content. A traditional film of two hours costs, at the high end of the Indian industry, between fifty and three hundred crore. A generative world that delivers, on average, ten hours of high-quality experience per user can be produced, with current technology, for between one and ten crore — and improves continuously as users interact with it, without additional production cost. The unit economics are, in other words, an order of magnitude better, and the experience is several times longer.

This is why I expect generative worlds to absorb significant portions of the traditional entertainment markets. A producer who can deliver ten hours of high-quality experience for one-tenth the cost of two hours of traditional film will, in time, capture the market unless the traditional film offers something the generative world cannot. Traditional film does still offer some things — the

auteur's specific creative vision, the actors' specific performances, the narrative coherence that authored content delivers — and these will remain valuable. But the bulk of the market, in volume, will migrate to the generative side.

◆ ◆ ◆

IV. The hard problems

The hard problems in generative worlds are several.

Coherence over time. The world has to feel like a real place, with consistent characters, consistent history, consistent rules. Generative AI tends to produce content that drifts; the same character can behave inconsistently from one interaction to the next. Solving this is the central technical problem, and the firms that solve it best will dominate the category.

Quality control at scale. When content is generated on demand, the producer cannot inspect every output the way she could inspect a fixed film. The producer must instead build the constraints, the values, and the quality standards into the generative system itself. This is harder than it looks. Operators that fail at it will produce experiences that periodically embarrass them.

Ethical guardrails. A generative world that adapts to its user can adapt in directions the producer did not intend — toward content that exploits, manipulates, or harms the user. The guardrails must be deliberate and robust. Regulators will, in time, require them; producers should impose them on themselves first.

The compute cost. High-quality generative worlds are computationally expensive. The cost is falling, but it remains a significant component of the unit economics. Producers must design carefully for what to generate fresh each session and what to cache, to keep the experience responsive without consuming unsustainable compute.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you produce content of any kind — entertainment, education, marketing, training — commission a generative-worlds pilot this year.**

   The pilot will tell you what your existing capabilities translate into in this medium and what you will need to build.

2. **If you operate in any industry where simulation, training, or rehearsal matters, evaluate a generative-worlds approach against your current methods.**

   In most categories, the unit economics will favour generative worlds within twenty-four months. The leaders who pilot now will be ready when the favour becomes clear.

3. **Begin building, or hiring, the talent that combines storytelling with technical fluency in generative AI.**

   This combination — narrative imagination plus technical execution — is the rarest and most valuable hybrid skill in the entertainment industry of the next decade. Most firms have neither half. Acquiring both is now urgent.

CHAPTER TWENTY-THREE

23

Molecular Services

Healthcare, materials, and food at the molecular level.

—

"Nature is the source of all true knowledge."

— Leonardo da Vinci

◆ ◆ ◆

The founder of a synthetic-biology start-up I have been mentoring told me, with some delight, in early 2026, that the cost of designing a new protein had fallen, in the previous three years, by a factor of approximately a thousand. The design was now done by an AI in hours rather than by a team of researchers in months. The synthesis was done by a service that took the design and produced the physical molecule for delivery within days. The testing was done by automated lab equipment that ran experiments at machine speed. The whole pipeline that, in 2018, would have required a research lab of fifty people and a budget of millions of dollars, could now be operated in 2026 by a team of three with a budget of a few crore. He was using the new economics to design a custom enzyme for an application that, at 2018 cost levels, would never have been worth pursuing. At 2026 cost levels it was a viable product.

He is, in 2026, one of perhaps a few thousand entrepreneurs around the world taking advantage of the same cost collapse. By 2030,

there will be tens of thousands. By 2032, the new economics will have produced a category of business — molecular services — that will resemble, in its commercial structure, what software became in the 1990s and 2000s. The barrier to entry will have collapsed. The number of practitioners will have multiplied. The applications will have proliferated across industries.

◆ ◆ ◆

I. The category, defined

Molecular services are products and services designed and delivered at the molecular level, customised to the individual user, drawing on AI-assisted molecular design and on the rapidly falling cost of biological and chemical synthesis. The category includes personalised pharmaceuticals, personalised nutrition, personalised agricultural products, personalised cosmetics, personalised materials, and the diagnostic, design, and synthesis services that support all of the above.

The category is, in 2026, present in early form across several adjacent industries. Personalised pharmaceuticals — drugs designed for the specific patient's genetic profile and medical context — exist in oncology and are extending to other areas. Personalised nutrition — supplements and meal plans designed for the individual's genome, microbiome, and lifestyle — is a growing consumer category. AI-designed agricultural inputs — fertilisers, pesticides, seed treatments — are entering trials. AI-designed materials are beginning to appear in batteries, catalysts, and structural composites. None of these applications is yet at scale. All are scaling. By 2030, the category will be one of the most significant new industries of the convergence economy.

◆ ◆ ◆

II. The technological enablers

Three technological capabilities, each of which has matured in the past five years, combine to enable molecular services.

AI-assisted molecular design. The ability to predict the structure and function of a molecule from its sequence, and to design novel molecules with desired properties, has advanced enormously since AlphaFold's release in 2021. By 2026, the design of new proteins, small-molecule drugs, materials, and agricultural compounds is routinely done by AI models, with human researchers in the role of curators rather than primary designers. The acceleration of design timelines is, in some categories, by factors of one hundred or more.

Cheap, fast synthesis. The cost of synthesising a custom DNA strand, a custom protein, or a custom small-molecule has fallen continuously over the past two decades. By 2026, the unit costs are low enough that experimental products that would not have been commercially viable five years ago are now routine. The synthesis is done by service providers that operate at scale, the way semiconductor foundries do for chip designers, allowing molecular designers to outsource the manufacturing entirely.

Automated testing and characterisation. The ability to test a new molecule's properties — pharmacological, mechanical, agricultural, or otherwise — has been increasingly automated. AI-augmented robotic labs can run thousands of experiments in the time a human team could run one. The testing is the bottleneck that historically limited the rate at which new molecular products could be developed; the bottleneck has, in 2026, largely been removed.

◆ ◆ ◆

III. The applications, in three buckets

The applications fall into three broad buckets, each with different commercial dynamics.

Healthcare. Personalised pharmaceuticals, gene therapies, mRNA vaccines designed for specific pathogens or specific patients, custom probiotics, custom nutritional plans. The healthcare bucket is the largest by current revenue and the most regulated. The

leaders in 2026 are the established pharmaceutical firms with significant AI investments, plus a small number of new entrants who have raised significant capital. The Indian opportunity in this bucket is significant given the country's pharmaceutical industry, but realising it requires regulatory clarity that, in 2026, is still partial.

Materials. Custom catalysts for industrial processes, custom polymers for specific applications, custom alloys for aerospace and defence, custom semiconducting materials. The materials bucket is less regulated and more directly commercially competitive. The leaders are a mix of established materials firms with AI capabilities and new specialised firms. The Indian opportunity here is largely undeveloped, despite India's strong materials science research base.

Food and agriculture. Custom seeds for specific climates, custom fertilisers for specific soils, custom feeds for specific livestock, custom flavours and ingredients for the food industry. The food and agriculture bucket has lower margins and higher volume than the others, with significant social and policy considerations. The Indian opportunity is enormous given the country's agricultural sector, but progress is slowed by complex political dynamics around agricultural reform.

◆ ◆ ◆

IV. The competitive dynamics

The competitive dynamics of molecular services resemble those of software more than those of traditional pharmaceuticals or materials. The barrier to entry has collapsed. The development cycles are short. The most valuable assets are the AI models, the datasets, the synthesis capacity, and the regulatory expertise — all of which can be assembled by relatively small teams. The result is that the next decade will see a proliferation of small, specialised firms operating in molecular services, alongside the larger

established firms whose advantages are scale and regulatory capability.

The Indian opportunity in molecular services is, in my view, larger than the Indian opportunity in foundation-model AI, for two reasons. First, the cost economics favour countries with strong scientific talent at moderate cost, which describes India well. Second, the applications — particularly in healthcare and agriculture — are particularly suited to Indian conditions and Indian needs. The leaders in this category, by 2032, will include a meaningful number of Indian firms, if the regulatory and capital environments support their development.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in healthcare, materials, or agriculture, commission an AI-assisted design pilot in your most critical product category this year.**

   The pilot will surface what your firm can and cannot do at the molecular level today, and where the gaps are.

2. **Identify one personalisation possibility in your current product line that becomes economic at the new molecular cost levels.**

   The personalisation that was uneconomic two years ago is, in many cases, economic now. The firm that ships first will set the category.

3. **Engage with the regulatory development of personalised medicine, materials, and agriculture in India.**

   The regulatory environment is developing. The firms that engage now will help shape it favourably for legitimate practice.

CHAPTER TWENTY-FOUR

24

The Provable Truth Layer

Verification becomes a structural value.

—

"Trust is the only currency that is not subject to inflation — except when it is."

— A line attributed to several economists, none of whom can confidently be credited

◆ ◆ ◆

The chief executive of one of India's better-known news organisations called me, in mid-2024, to discuss what she described as an existential challenge. Her organisation had, over a long period, built a reputation for trustworthy reporting. The reputation had been valuable — it had supported the organisation's commercial business, justified premium subscription pricing, and given its journalists the access that produced the next round of trustworthy reporting. The reputation was now, she felt, under genuine threat from the rise of synthetic media. AI-generated images, AI-generated voices, AI-generated video, AI-generated text — all indistinguishable, to the casual reader, from the genuine article. The organisation's trusted reporting was now in a marketplace alongside a flood of synthetic content, much of it from sources that were either malicious or careless about verification. Her readers could not, in many cases, tell the difference. She was

not sure her organisation could either, in some cases, without significant additional verification work.

She wanted to know what to do. The conversation we had, and the work that has followed it, has become one of the most important strands of my own thinking. This chapter is the distilled version of that thinking. The category I will describe — the provable truth layer — is, in my view, one of the most important commercial categories of the next decade and one of the most important civic infrastructures of the digital era.

◆ ◆ ◆

I. The problem the category solves

The problem is that, by 2026, the marginal cost of producing convincing synthetic content has fallen to near zero, while the cognitive cost of verifying any specific piece of content has remained roughly the same. The result is an asymmetry that did not exist before: malicious or careless actors can produce a thousand pieces of synthetic content for the cost that a verifier can spend examining one. The asymmetry favours the producer of the synthetic content, and it scales.

The implication for the information ecosystem is significant and well-documented. Public discourse becomes harder to conduct on shared facts. News organisations face declining trust. Individuals become increasingly unable to verify, in any specific case, whether the content they are seeing is genuine. The cumulative effect is a kind of epistemic erosion that is, in 2026, well underway and that, without deliberate counter-measures, will worsen.

The provable truth layer is the category of infrastructure designed to address this erosion. It is the cryptographic, social, and institutional infrastructure that allows specific pieces of content to be verified as authentic — not merely asserted as authentic by a publisher, but verifiable by the reader through means that do not require the reader to trust the publisher. The category includes

content provenance standards, decentralised identity systems, cryptographic signing infrastructure, third-party verification services, and the institutional practices that produce trustworthy content in the first place.

◆ ◆ ◆

II. The technical infrastructure

The technical infrastructure of the provable truth layer has four components.

Provenance metadata. Content carries cryptographic metadata — generated at capture time by the camera, recording device, or composition tool — that records the chain of custody from creation to publication. The C2PA standard, finalised in 2023 and now widely adopted, is the industry's primary specification. Content with C2PA metadata can be verified as having originated from a specific device at a specific time, with any subsequent edits documented.

Decentralised identity. The publisher who signs the content has a verifiable identity that does not depend on a single platform. The decentralised identifier (DID) standards allow any party to publish content under an identity that any other party can verify cryptographically. This is the same technical layer that supports self-sovereign identity, applied to publishing.

Verification services. Third-party services examine specific pieces of content and produce verification reports — yes, this image was captured by this camera at this time, no, this video shows signs of synthesis, this audio matches the voice signature of this speaker. The verification services operate as trusted institutions in the same way that auditors and certifying bodies do in other industries.

Reputation infrastructure. The aggregate reputation of publishers, sources, and verification services is itself maintained as a public infrastructure. Readers can check, before trusting a specific piece of content, the reputation of its publisher and its verifiers. This

infrastructure is, in 2026, partially developed; it will mature significantly in the next several years.

[Figure 24.1 — to be supplied by illustrator. THE FOUR-LAYER PROVABLE TRUTH STACK. A vertical pyramid: provenance metadata, decentralised identity, verification servi…]

◆ ◆ ◆

III. The institutional layer

The technical infrastructure is necessary but not sufficient. The provable truth layer also requires institutional practices — practices that media organisations, governments, financial institutions, healthcare providers, educational institutions, and other producers of authoritative content adopt deliberately. The practices include, at minimum: signing all published content cryptographically, publishing the verification keys publicly, allowing independent third-party verification, accepting accountability for verified errors, and maintaining transparent records of how content is produced.

The institutions that adopt these practices early will, in time, become the trusted nodes of the information ecosystem. The institutions that do not will, in time, become indistinguishable from the flood of unverified synthetic content. The differentiation will be visible to readers, to advertisers, to regulators, and to capital markets. The trust currency, accumulated through deliberate practice, will become a structural commercial advantage.

◆ ◆ ◆

IV. The Indian opportunity, again

India has, in 2026, one of the more advanced regulatory and infrastructural positions on the provable truth layer of any major country. The deepfake regulations, the verifiable credentials work, the digital public infrastructure that supports decentralised identity, and the emerging governance around AI-generated

content are all ahead of where most jurisdictions are. The country also has, in its democratic public discourse and its institutional press, the mature institutional layer that the technical infrastructure requires.

The opportunity is to build the commercial infrastructure of the provable truth layer — the verification services, the publisher tools, the consumer-facing applications — to global standards from an Indian base. The market is global. The technical work can be done from anywhere. The institutional credibility of Indian firms in this space, given the country's regulatory posture, is significant. I expect a meaningful number of the global leaders in the provable truth layer, by 2032, to be Indian companies.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If your firm publishes content of any kind — news, marketing, financial reporting, scientific results — adopt content provenance standards this year.**

   The adoption is technically straightforward. The differentiation it produces, and the reader trust it builds, is significant.

2. **Audit your firm's exposure to synthetic content as an attack surface.**

   Most firms have not done this audit. The exposure is larger than they expect. The mitigation is feasible if begun early.

3. **Engage with the institutional development of the provable truth layer in your industry.**

   Whether you are in media, financial services, healthcare, government, or education, the institutional layer is being built now. The firms that participate will shape the standards. The firms that do not will, eventually, comply with standards shaped by others.

CHAPTER TWENTY-FIVE

25

Adaptive Everything

The cross-cutting category that affects every product.

—

"Anything that can be made adaptive will be made adaptive. Anything that cannot be made adaptive will lose."

— A working hypothesis I have come to hold

◆ ◆ ◆

The chief product officer of a consumer electronics company I have advised told me, sometime in 2024, that her team had begun to feel, in their internal product reviews, an emerging phenomenon they did not yet have a vocabulary for. The phenomenon was that products in their roadmap that had been planned as static — a fixed feature set, the same for every user, the same on day one and day three hundred — increasingly felt, in review, *broken*. The team had not changed its standards. The world had. The team's reviewers were, increasingly, judging static products against a benchmark of adaptive products. A product that did not learn its user, did not adjust to her preferences, did not improve over the course of the relationship — such a product no longer felt like a finished product. It felt like a half-built one.

The phenomenon she was describing is the cross-cutting category I have come to call adaptive everything. The previous five new categories are specific industries or technical infrastructures.

Adaptive everything is the *general principle* that, by 2030, every product in every category will be expected to adapt to its user, and products that do not will be perceived as broken. The category is therefore not a vertical to enter; it is a horizontal expectation that affects every vertical.

◆ ◆ ◆

I. The principle

The principle is simple: by 2030, the default expectation for any product the customer interacts with will be that the product learns her, adjusts to her, and improves over the course of the relationship. The static product — the product that is the same on day one and day one thousand — will be the exception, and will require justification rather than being the default.

This principle has been visible in software for a decade. The good consumer software of 2026 — the streaming service, the email client, the navigation app, the photo organiser — adapts to its user. The user does not configure it; it configures itself around her. By 2030, this expectation will have spread to physical products, to services, to spaces, to financial instruments, to vehicles, to consumer goods. The car will learn the driver. The chair will learn the sitter. The kitchen appliance will learn the cook. The financial product will learn the household. The retail brand will learn the customer. The non-adaptive version of each of these will feel, by 2030, the way a non-touchscreen phone feels in 2026.

The principle has implications for every product organisation. The products planned today on static-product economics will, by 2030, feel obsolete. The products planned on adaptive-product economics will be perceived as the new default. The transition is not optional, and it does not have a clear single moment of inflection; it is a gradual shift in expectation that occurs in different segments of the market at different times. The leader's responsibility is to understand the curve and to position the firm's

product roadmap such that the firm is shipping adaptive products before its segment of the market has come to expect them.

◆ ◆ ◆

II. Three levels of adaptation

Adaptation, in product terms, comes in three levels of depth.

Surface adaptation is what most consumer software did in the 2010s. The product remembers the user's basic preferences — language, time zone, recently used items — and adjusts accordingly. The adaptation is shallow but useful, and it is, in 2026, a baseline that customers expect even of the simplest products.

Behavioural adaptation is what the better consumer software does in 2026. The product learns the user's patterns — when she uses what, in what sequence, with what intent — and adjusts itself proactively. The product's interface is not static; it surfaces what the user is likely to need, hides what she is likely to ignore, and modifies its behaviour as her patterns change. The adaptation is meaningful and is increasingly the standard for high-quality digital products.

Contextual adaptation is the next frontier. The product not only learns the user but adapts to her current context — her physical location, her emotional state, her recent activities, her current goals, the time of day, the social setting, the device she is using. The product is, in effect, situationally aware. By 2030, contextual adaptation will be the standard for the best products in any category. By 2032, it will be the default expectation for most products.

[Figure 25.1 — to be supplied by illustrator. THE THREE LEVELS OF ADAPTATION. A vertical pyramid: surface adaptation (bottom), behavioural adaptation, contextual adap…]

◆ ◆ ◆

III. What this asks of product organisations

Building adaptive products requires a different product organisation from the one most firms have. I want to name three differences.

Data architecture as a first-class concern. Adaptive products require data — about the user, about her use, about the contexts she uses the product in. The architecture for collecting, storing, processing, and protecting this data is not an afterthought; it is the foundation on which the adaptation is built. Most product organisations in 2026 still treat data architecture as an infrastructure concern handled by a separate team, rather than as a core product concern handled by the product team itself. This will change.

Continuous deployment as the default. Adaptive products improve continuously, which means they are, in some sense, never finished. The product organisation that builds them must be capable of shipping improvements weekly or more frequently. This is the operating cadence of the better consumer software firms in 2026. It will be the default cadence for product firms in every category by 2030. Firms that operate on annual product release cycles will, by 2030, be visibly behind.

Trust and privacy as core product features. Adaptive products know the user, which means the user must trust the product. The trust is not built by terms-of-service language; it is built by transparent data practices, genuine user control, and visible respect for privacy. The product team must include this as a core feature, not as a compliance checklist item. Products that fail at this will, increasingly, be rejected by users and regulators alike.

◆ ◆ ◆

IV. The compounding implications

Adaptive products, by their nature, compound. The longer the user remains in the relationship, the more the product knows about her,

the more valuable the product becomes to her, the higher the switching cost. This is the same compounding I described in the experience economy framing earlier in the book; adaptive everything is the operational mechanism by which experience compounds at the product level.

The implication for competitive dynamics is significant. The first-mover advantage in adaptive products is large. The firm that builds the adaptive version of a product first, and earns the user's trust, accumulates user-specific value that competitors cannot replicate without spending the same time. The firm that ships the adaptive version second is competing not against an adaptive product equivalent to its own, but against an adaptive product that has had two years longer to learn its users than its own product has had.

This is why the leader's responsibility, in 2026, is not to wait for the curve to be obvious. The curve is obvious enough now to act on. The leaders who act now will accumulate the first-mover advantages. The leaders who wait will be competing, by 2030, against products that have learned their customers for years longer than their own products have had the chance to.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your top three products against the three levels of adaptation.**

   Most products score at surface adaptation, with aspirations at behavioural. The audit makes the gap between current and future visible.

2. **Identify one product in your portfolio that should be re-platformed as contextually adaptive within twenty-four months.**

   The re-platform is large. It is also necessary. Begin it.

3. **Re-architect your data, deployment, and trust practices around the requirements of adaptive products.**

   The architecture work is unglamorous. The architectural work is the foundation. Without it, the adaptive product layer cannot be built.

PART SIX

VI

Designing for What's Beginning



The interfaces and design principles for a post-screen world.

THE WORLD THAT'S BEGINNING

The interfaces and design principles
for a post-screen world.

CHAPTER TWENTY-SIX

26

From Screens to Spatial to Neural

Forty years of interface evolution, and what the next ten will look like.

—

"We become what we behold. We shape our tools,
and thereafter our tools shape us."
— Marshall McLuhan, again, deservedly

◆ ◆ ◆

I bought my first personal computer in the mid-1990s, when I was a young engineer recently out of college. The interface was a keyboard and a small monochrome screen running a command line. To do anything useful, I had to know the names of the commands. The computer did not anticipate me. The computer did not adapt to me. The computer did exactly what I asked, in the exact syntax it required, or it did nothing. I spent, I think, the first six months of my relationship with that computer in a state of quiet hostility, until I learned enough commands to make it cooperative. The hostility, in retrospect, was rational. The interface was hostile.

I am writing these words on a computer thirty years later. The interface is dramatically different. I have spoken several paragraphs of this chapter aloud, and the computer has transcribed them. I have selected sections to revise by gesturing at them. The AI assistant has, on its own initiative, suggested two factual checks and one stylistic improvement, both of which I have accepted. The

computer is anticipating my needs, adapting to my style, and intervening only when it has reason to believe its intervention will be helpful. The interface is not exactly invisible — I am still aware of typing and gesturing — but it is, compared to the command line of 1995, almost invisible.

This chapter is about the trajectory between those two moments, and about where the trajectory leads next.

◆ ◆ ◆

I. Four eras of human-computer interface

The trajectory, since the personal computer became a mass-market product around 1980, has gone through four major eras. Each lasted roughly fifteen years. Each was, when it began, less efficient than the era that preceded it; each, by the time it ended, had become so much better that the previous era looked obsolete in retrospect.

The desktop era — roughly 1985 to 2007. The interface was the desktop metaphor: windows, icons, menus, pointers. The dominant device was the personal computer at a desk. The interaction was deliberate; the user sat down at the computer to do specific things and walked away from it when she was done. The desktop era produced enormous productivity gains, particularly in knowledge work, and ended only when a more accessible mode of interaction arrived.

The mobile era — roughly 2007 to 2022. The interface became the touchscreen, with applications as the dominant unit of interaction. The dominant device was the smartphone, carried continuously, used in short bursts throughout the day. The interaction was casual; the user did not sit down to do things, she pulled out her phone, did the thing, and put it away. The mobile era multiplied the population of people who could use computers — by some

estimates, by a factor of three — and produced the platform economics that defined the 2010s.

The spatial era — roughly 2022 to 2032. The interface is becoming three-dimensional, with the computer no longer a screen but a presence in the user's physical space. The dominant devices are, increasingly, headsets, glasses, and ambient computing in the environment itself. The interaction is contextual; the computer is aware of where the user is, what she is looking at, what she is doing, and surfaces interaction at the appropriate moment. The spatial era is, in 2026, in its early years. By 2032, it will be the dominant interface paradigm for a meaningful portion of professional and consumer use.

The neural era — roughly 2030 onwards. The interface is, increasingly, direct between the user's nervous system and the computer. The early forms — non-invasive brain-computer interfaces for medical and accessibility applications — are already shipping. The mature forms, in which the user can direct the computer through thought and receive information through neural feedback, are still distant but increasingly visible on the technical horizon. The neural era will, in my forecast, become significant in the 2030s and dominant by the 2040s.

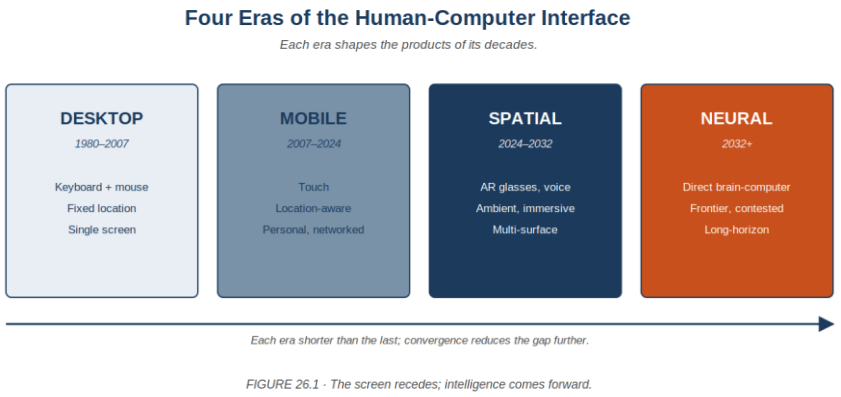


Figure 26.1

◆ ◆ ◆

II. What the spatial era actually looks like

The spatial era, which is the present-tense interface concern for most product organisations in 2026, has three operational characteristics worth naming.

The interface is contextual. The computer is aware of the user's location, her recent actions, her likely goals, and her surroundings, and surfaces interaction appropriately. The application paradigm of the mobile era — discrete apps, each launched explicitly by the user — is dissolving in favour of capabilities that surface contextually. The user does not "open the maps app"; the directions appear when she gets in the car. The user does not "open the recipe app"; the recipe appears when she stands in front of the kitchen counter.

The interface is multi-modal. The user interacts through voice, gesture, gaze, touch, and context, often combining several modalities in a single interaction. The single dominant modality of the desktop era (keyboard and mouse) and the mobile era (touchscreen) is replaced by a fluid combination that adapts to what the user is doing.

The interface is environmental. The interaction is no longer mediated entirely through a single device the user carries. It is distributed across the environment — in the glasses she wears, the room she is in, the car she is driving, the public spaces she walks through. The computer is, in some sense, ambient, with specific affordances becoming visible at specific moments.

The implications for product design are significant. Designers trained for the desktop or mobile eras have to develop new instincts for contextual, multi-modal, environmental interfaces. Most of them have not had the time or the training. The shortage of qualified spatial designers is, in 2026, one of the binding constraints on the rate at which the spatial era is being built. The firms that have invested early in this talent — and the educational

institutions that have begun to produce it — are, in the medium term, structurally advantaged.

◆ ◆ ◆

III. The neural era, briefly

The neural era is far enough out that it does not require operational planning by most product organisations in 2026. It does, however, require strategic awareness, because the work that is happening in neural interfaces today will mature into mainstream products on a timeline that is not as long as most leaders assume.

The current state of the art, in 2026, includes invasive brain-computer interfaces approved for specific medical applications (paralysis, certain neurological conditions), non-invasive interfaces using EEG and functional near-infrared spectroscopy for accessibility and consumer applications, and emerging research interfaces using transcranial focused ultrasound and other modalities. The capabilities are limited compared to the mature visions, but they are real, and they are improving rapidly.

By 2030, I expect non-invasive consumer applications to begin reaching market for specific use cases — meditation aids, focus enhancement, accessibility tools. By 2035, more general-purpose neural interfaces are likely to be shipping in early consumer form. By 2040, the neural era will be sufficiently mature that the interfaces of the 2020s will look, in retrospect, like the command lines of the 1990s.

The ethical and governance implications of the neural era are significant, and I will treat them in Chapter Thirty-Four. They are also far enough out that the leader of an enterprise in 2026 has time to engage thoughtfully rather than reactively. I encourage that thoughtful engagement; the firms that begin the conversation now will help shape it well.

◆ ◆ ◆

IV. The strategic implication

The strategic implication for the leader in 2026 is twofold. First, every product roadmap should include a deliberate strategy for the spatial era — what the firm's products look like in spatial interfaces, what new capabilities those interfaces enable, what user behaviours they make possible. Most firms have not yet committed to such a strategy. The ones that do will be ready when their customers expect spatial-era products. The ones that do not will, by 2030, be visibly behind.

Second, every leadership team should establish at least a watch on the neural era — not a roadmap commitment, which would be premature, but a deliberate awareness of the technology's progress and its implications for the firm's category. The lead time on capability building, regulatory engagement, and ethical positioning in the neural era is long enough that beginning the watch now is appropriate even though active investment is not.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Commission a spatial-era strategy for your firm's products this year.**

   The strategy does not need to be exhaustive. It needs to identify which products will be re-platformed for spatial interfaces, on what timelines, with what capability investments. Without this strategy, your firm will arrive at the spatial era unprepared.

2. **Hire or partner with one designer who has serious spatial-era skills, even if you do not yet have a spatial product.**

   The talent shortage is severe. Acquiring the talent now is meaningfully cheaper than acquiring it in 2028.

3. **Begin a quarterly internal briefing on neural interface developments.**

   The briefing is a low-cost insurance against being surprised. The probability that the neural era arrives faster than expected is non-trivial, and the firms that have been watching will be ready when the curve steepens.

CHAPTER TWENTY-SEVEN

27

The Five Post-Screen Principles

How to design when the screen is no longer the dominant medium.

—

"The medium is the message."
— Marshall McLuhan, for the third time

◆ ◆ ◆

I sat in a design review at a firm I have advised for several years, in mid-2025, watching a senior designer present a beautifully-rendered mock-up of a new product the firm intended to ship in 2027. The mock-up was, in its way, exquisite. Every screen was carefully designed. The typography was elegant. The visual hierarchy was clear. The interactions were thoughtfully sequenced. The team had spent four months on the design, and the work showed.

The chief product officer, who had brought me in as an outside reviewer, asked me what I thought. I asked, in turn, where on the device the user would be looking when she was using this product. The room paused. The designer answered, reasonably, that the user would be looking at the screen — that was the point of the design. I asked whether the user would be sitting at her desk, walking, driving, in a meeting, in her home, on a train. The designer answered that the design assumed a quiet desk environment. I asked what fraction of the product's intended use

would actually occur in a quiet desk environment. The designer thought for a moment and, honestly, said: "Probably not most of it."

The conversation that followed became one of the more useful design conversations I have been part of. The team had designed, as designers had been trained for thirty years, for the screen. The product they were going to ship in 2027 would be used, increasingly, in contexts where the user was not looking at a screen — in the car on the way to work, while preparing dinner, while walking through a public space, while in conversation with someone, while attending a meeting in a spatial-AR environment. The screen-first design they had produced would, in many of those contexts, be poorly suited to the actual use.

This chapter is the framework I have come to use for designing in the post-screen era. There are, I believe, five principles. They are not novel; experienced designers have intuited them for some time. The systematic articulation, however, is something I have not seen elsewhere, and the framework has been useful in my own work and in the work of teams I have advised.

◆ ◆ ◆

I. Spatial is the new flat

The first principle is that the canvas of design has moved from two dimensions to three. The flat screen — the rectangle on which most digital design has been done since 1985 — is no longer the dominant surface. The dominant surface is increasingly three-dimensional space, with the user at the centre and the content distributed around her.

Designing for three dimensions requires different instincts than designing for two. Spatial relationships matter — what is in front of the user, what is behind, what is to her left, what is to her right. Depth matters. Distance matters. Orientation matters. Movement matters — how does the user move through the space, and how

does the space respond to her motion. The visual hierarchy that was established by font size, colour, and position on a flat page is now established by spatial location, distance, and dynamic behaviour.

Most designers in 2026 are not yet fluent in this new language. The fluency is acquired by deliberate practice, with the right tools, in the right environments. Firms that invest in this fluency now will, in five years, have design teams that produce dramatically better spatial products than firms that have not invested.

◆ ◆ ◆

II. Voice is the new keyboard

The second principle is that the dominant input modality is shifting from text typed on a keyboard or touchscreen to voice spoken naturally. The shift is uneven across contexts — voice will not entirely replace typed input, particularly for precision work — but for the bulk of consumer interactions and for an increasing fraction of professional ones, voice will be the default.

Voice as an input modality has its own design vocabulary. Conversations have rhythm. They have turn-taking. They have implicit context. They have ambiguity that has to be handled gracefully. Voice interfaces that simply transcribe speech and treat it as if it were typed text — which is what most voice interfaces of 2024 and 2025 did — are inadequate. Voice interfaces that genuinely participate in conversation, with appropriate timing, tone, context-awareness, and graceful repair of misunderstandings, are dramatically better.

For Indian designers, voice has an additional dimension: it is genuinely multi-lingual, often within a single utterance. The user who codes-switches between Hindi and English in a single sentence is the median Indian user, not the exception. Voice interfaces designed only for English, or only for grammatical Hindi, will feel foreign to most Indian users. Voice interfaces that

handle the natural code-switching of Indian speech will be perceived as Indian, in a way that flat-text interfaces never quite manage.

◆ ◆ ◆

III. Anticipation is the new navigation

The third principle is that the user no longer navigates explicitly through a menu structure to reach what she wants. The system anticipates what she will want and surfaces it. The user accepts or rejects the anticipation; the system learns from the acceptance or rejection.

This shift requires a different design philosophy. The conventional product design of the desktop and mobile eras assumed that the user knew what she wanted and would navigate to it. The post-screen design assumes that the user has a goal and that the system should figure out the most useful next step. The user is no longer the navigator; the user is the verifier, accepting or refining what the system proposes.

This has implications for what to expose in the interface and what to hide. The conventional desktop application exposed everything — every feature, every option, every setting — on the theory that the user might need any of them at any time. The post-screen interface exposes very little, on the theory that the user wants the system to guess correctly and that exposed options, when not needed, are friction. The shift requires designers to think harder about what the user actually wants in any given context, rather than offering her every possibility and letting her choose.

◆ ◆ ◆

IV. Ambient is the new app

The fourth principle is that the discrete application, launched explicitly by the user, is being replaced by capabilities that surface contextually in the environment. The user does not "open the

calendar app" to check her schedule; her schedule appears when relevant, on whatever surface is appropriate, in whatever modality fits the moment. The user does not "open the messaging app" to send a message; she speaks the message and it goes, regardless of which device she is on or which platform her recipient uses.

The unit of design is, increasingly, the capability rather than the application. A capability — schedule access, messaging, navigation, search, payment — is something the user can invoke from any context, through any modality, on any device. The capability is, in this sense, atomic; it is the building block that the system composes into context-appropriate interactions on the user's behalf.

Designing capabilities rather than applications is a different exercise. The capability designer must think about every possible context in which the capability might be invoked, every possible modality, every possible device. The capability must work consistently across all of them while adapting to each. This is harder than designing a single application interface, but it is the work the post-screen era requires.

◆ ◆ ◆

V. Trust is the new login

The fifth principle is that the binary login — the user is either authenticated or not — is being replaced by continuous, layered, context-appropriate trust. The system knows, through a combination of passive and active signals, who the user is and how confident it is in that knowledge. Different actions require different confidence levels. The user does not log in once at the start of the session; the system maintains a continuous understanding of her identity and adjusts what it allows accordingly.

This shift requires a different design philosophy. The conventional design assumed that authentication was a discrete moment — the user enters a password, and from that point on the system trusts

her until she logs out. The post-screen design treats trust as a continuum that varies by action — the system might allow the user to check her balance with passive authentication but require active authentication before transferring funds, biometric authentication for large transfers, and multi-factor authentication for unusual transactions.

The user experience of this shift, when designed well, is invisible. The user does not notice the authentication; it happens in the background, in proportion to the action. The user experience, when designed badly, is intrusive — constant prompts to verify, constant interruptions of flow, constant friction. The differentiation between good and bad design in this layer is significant.

[Figure 27.1 — to be supplied by illustrator. THE FIVE POST-SCREEN PRINCIPLES. A pentagon diagram with five labelled vertices: spatial, voice, anticipation, ambient,…]

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your current product against the five principles.**

   Most products score well on one, partially on two, and poorly on the rest. The audit makes the gaps visible.

2. **Commission a single redesign this year that takes one of the five principles seriously.**

   The redesign will teach your team more than any conference. The lessons will inform the next four redesigns.

3. **Begin training your designers, formally, in the post-screen design language.**

   Most design education in 2026 still teaches screen-first design. The training your firm provides will, for some years, be the primary source of post-screen design expertise on your team.

CHAPTER TWENTY-EIGHT

28

The Irreducibly Human

What machines cannot do, and why that matters more, not less.

—

"The most important questions of life are, for the most part, only problems of probability."

— *Pierre-Simon Laplace, with apologies for the digression*

◆ ◆ ◆

I was, at a conference in Bengaluru in 2024, asked by a journalist what I thought AI could not do. The question is asked of me, in some form, by every audience I have spoken to since. I have, over time, come to give a more careful answer than I used to. The careful answer is that the question is ill-formed. There are very few cognitive tasks that AI cannot, in some form, do; the more useful question is which tasks are most valuable when done by a human, with AI as augmentation, rather than by AI alone. The list of those tasks is, I have come to believe, longer than most leaders assume — and growing in importance, not declining, as AI absorbs the surrounding work.

This chapter is about that list. It is, in some sense, the counterweight to the previous chapters of the book, which have spent considerable effort describing what technology will do. This chapter describes what humans, in the technology-saturated

economy ahead, will continue to do — and what the leader must protect, develop, and reward.

◆ ◆ ◆

I. Five capabilities that remain durably human

The capabilities I list here are not, individually, new. They have been part of human work for as long as work has existed. What is new is that, as AI absorbs the structured cognitive work that used to fill the human's day, these capabilities rise in relative importance. The human worker who can do them well becomes, in 2030, more valuable than she was in 2020. The human worker who cannot becomes, in 2030, less valuable — because the structured cognitive work that used to be her contribution is being done by AI.

Empathy and authentic emotional connection. The capacity to genuinely understand another person's emotional state, to respond to it with the appropriate human warmth, and to build a relationship of trust over time is not a capability AI replicates. AI can simulate empathy — and, in many limited contexts, the simulation is good enough. But in the contexts that matter most — the consequential medical conversation, the bereavement, the celebration, the deep personal trust — humans want and need other humans. The economic value of authentic emotional connection in 2030 will, in my forecast, be higher than it has been in any previous decade.

Judgment in genuinely novel and ambiguous situations. AI is excellent at tasks that resemble what it has seen before. AI struggles with tasks that are genuinely novel, where the right answer depends on subtle contextual factors that are not in the training data. The senior leader's most consequential decisions — the strategic pivot, the response to an unprecedented crisis, the call on a new market — sit in this territory. The judgment that produces good outcomes in

genuinely novel situations remains a deeply human capability, and the humans who develop it will be disproportionately valuable.

Ethical and moral reasoning. AI can apply ethical frameworks that have been articulated for it. AI does not, in any meaningful sense, *develop* ethical positions; it does not feel the moral weight of its choices; it is not held accountable in the way that humans are held accountable. The ethical reasoning that society and organisations require — the weighing of competing values, the acceptance of moral responsibility for choices, the development of new ethical understandings as conditions change — is irreducibly human. The humans who can do this work well will increasingly be needed as AI takes on more consequential tasks.

Creative breakthroughs that genuinely surprise. AI is excellent at producing combinations and variations within the space of what it has seen. AI is not yet excellent at producing the genuinely surprising creative breakthrough — the work that opens a new space rather than exploring an existing one. The artists, scientists, designers, and entrepreneurs who produce such breakthroughs are, in my forecast, more valuable in 2030 than in 2020, because the cost of competent imitation has fallen and the rarity of genuine breakthrough has been thrown into relief.

Embodied physical expertise. The surgeon's hands, the artisan's craft, the athlete's performance, the mother's care of her child — these embodied capabilities are not replaced by AI in the foreseeable future. Robots can do many things, but the integration of physical expertise with situated judgment that human bodies achieve over years of practice remains beyond machine capability. The economic and cultural value of embodied human expertise is, in 2030, higher than it was in 2020.

◆ ◆ ◆

II. Why these matter more, not less

The intuitive expectation, when AI absorbs more cognitive work, is that human work in general becomes less valuable. The opposite is true for the irreducibly human capabilities I have listed. As AI absorbs the structured cognitive work that used to surround them, these capabilities are no longer one component of a larger job; they are the central contribution the human is making. The human's time is freed from the structured work to focus on the irreducibly human work. The economic value of that work, per hour, rises.

The medical example is instructive. The physician who, in 2020, spent forty percent of her day on history-taking, documentation, prescription writing, and other structured tasks now spends, in 2026, perhaps ten percent of her day on those tasks. The remaining ninety percent is available for the irreducibly human work — the diagnostic reasoning in ambiguous cases, the empathic connection with the patient, the ethical judgment about treatment options, the leadership of the care team. The physician's productivity has risen. So has the value of her work. The patient who interacts with her experiences a more present, more attentive, more thoughtful physician than she used to. The economic and human value created per physician-hour has multiplied.

The same pattern, I expect, will play out across every profession in which the structured work is increasingly absorbed by AI. The lawyer's irreducibly human work is the negotiation, the courtroom presence, the client counsel, the strategic judgment. The teacher's is the mentorship, the ethical guidance, the inspiration of students. The salesperson's is the trust-building, the contextual judgment about the customer, the human warmth that a transaction does not require but that a relationship demands. As AI absorbs more, the human's work concentrates into what AI cannot do, and the value of that work rises.

◆ ◆ ◆

III. What this asks of leaders

The leader's responsibility, in this transition, is to protect, develop, and reward the irreducibly human capabilities, deliberately. The protection requires keeping these capabilities visible in the firm's organisational structure, even as AI absorbs the structured work around them. The development requires investment in the training, mentorship, and lived experience that produces these capabilities — they are not produced by curriculum alone. The reward requires compensation structures, recognition systems, and career paths that genuinely value these capabilities, rather than treating them as soft skills to be tolerated alongside the technical competencies.

Most firms, in 2026, do not do this work well. The compensation structures still reward the structured cognitive work that AI is absorbing, and undervalue the irreducibly human work that is becoming the central contribution. The career paths still privilege the technical specialist over the contextual generalist. The training programmes still focus on the skills that AI is replacing rather than the skills that AI is making more valuable. The leader who reorients the firm — its compensation, its career paths, its training, its recognition — toward the irreducibly human capabilities will, by 2030, be running a firm that is more productive, more humane, and more valuable than competitors who have not made the reorientation.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **List the five irreducibly human capabilities most central to your firm's value creation.**

   The list will not exactly match mine. It will reflect your industry, your customers, your culture. The exercise of producing the list is itself useful.

2. **Audit your compensation, recognition, and career paths against that list.**

   Most firms reward the structured work AI is absorbing and undervalue the human work AI is elevating. The audit will surface this misalignment.

3. **Commission a redesign of one career path or compensation system to genuinely reward the irreducibly human capabilities your firm depends on.**

   The redesign is unglamorous. It is also, I suspect, the highest-leverage human-capital decision you will make this decade.

CHAPTER TWENTY-NINE

29

The Devices of Tomorrow

Glasses, earbuds, vehicles, ambient.
The screen recedes; intelligence comes forward.

—

"The most profound technologies are those that disappear.
They weave themselves into the fabric of everyday life until
they are indistinguishable from it."

— *Mark Weiser*

◆ ◆ ◆

I will close this part of the book by stepping back from the principles and looking at the devices through which the post-screen era will reach the user. The previous chapters have argued for the design language of the era; this one is about the hardware and infrastructure that will carry that language. The chapter is brief, because the device landscape is moving fast enough that detailed predictions a year out would be useful and detailed predictions five years out would be foolish. What I will offer instead is the broad shape of the device landscape, with three or four critical inflection points that the leader should watch.

◆ ◆ ◆

I. The four device categories of 2030

By 2030, in my forecast, the consumer device landscape will be organised around four primary categories, with the smartphone — dominant since 2007 — retreating into a secondary role.

AR glasses will be the dominant personal device for daily continuous use. They will have replaced, for many users, the smartphone as the primary always-with-me device. The glasses will be lightweight, all-day comfortable, unobtrusive in social settings, and capable of overlaying contextual information on the user's view of the world. The smartphone will not disappear, but it will become the secondary device — picked up for tasks the glasses cannot serve — rather than the primary one.

Earbuds will be the dominant audio device, but they will be more than audio devices. They will be continuous voice interfaces, biometric sensors, and ambient AI assistants. The earbuds will be in the user's ears for most of the day, providing continuous connection to her companion AI, ambient awareness of her environment, and selective audio enhancement of conversations and other sounds.

Vehicles will be, for many users, the third-most-used compute environment after the glasses and the home. The vehicle will be a continuous workspace, entertainment environment, and personal assistant. The autonomous-driving capability will, by 2030, be sufficient for the user to genuinely use the commute time for other things. The vehicle's compute, sensors, and screens will be a more substantial portion of the user's daily computing than her smartphone.

Ambient computing — sensors, displays, and interaction surfaces embedded in the environment — will be the fourth pillar. The home, the office, the retail space, the public infrastructure will all carry computing capacity that the user accesses without carrying a personal device. The interaction is, in most cases, mediated

through the user's glasses, earbuds, or voice, but the computing is distributed across the environment.

The smartphone, in this configuration, becomes one device among several rather than the centre. It will continue to be sold, used, and improved, but it will no longer be the dominant interface. The leader who plans her firm's product strategy around the smartphone as the primary interface in 2030 will be planning for a world that no longer exists by then.

◆ ◆ ◆

II. The infrastructure beneath

The four device categories are supported by infrastructure that has its own development trajectory. Three pieces are worth naming.

5G and 6G networks. The high-bandwidth, low-latency mobile networks that the spatial era requires are, in 2026, partially deployed. By 2030, they will be substantially complete in most urban areas, with 6G beginning to deploy in the most advanced markets. The network availability is the precondition for spatial AR at scale.

Edge compute. Much of the AI inference required by the post-screen interface needs to run on or near the user's device, for latency and privacy reasons. The edge compute infrastructure — small data centres distributed at the network edge, AI accelerators in consumer devices, federated computing models — is being built. The firms that have invested in this infrastructure will be advantaged in delivering the spatial era at scale.

Foundation models accessible via API. The intelligence underlying the post-screen interface is, in most cases, not built from scratch by the firm offering the interface. It is accessed via API from a foundation model provider. The economics of who builds these models and on what terms will, in significant part, determine which firms can build durable post-screen products and which cannot. India's

position in this layer is, in 2026, weaker than I would like it to be — the country has very few foundation model builders at the global frontier, and the firms building post-screen products in India are dependent, mostly, on American foundation models. This is, in my view, a strategic vulnerability that requires deliberate attention, and I will return to it in Chapter Thirty-Three.

◆ ◆ ◆

III. The strategic implications

The strategic implications of the device landscape are several.

The smartphone-first strategy is becoming obsolete. The firm that designs its products principally for the smartphone will, by 2030, be designing for a secondary device. The post-screen design has to be primary; the smartphone has to be one of the surfaces on which it lands.

The integration of multiple devices is the new design challenge. The user does not interact with one device at a time; she interacts with several simultaneously, each playing a role. Her glasses display the contextual information; her earbuds carry the audio; her vehicle hosts the work session; her home carries the ambient state. The product must work coherently across all of them. This is harder than designing for any single device, and most firms are not yet equipped for it.

The platform economics of the smartphone era will not, mechanically, transfer. The Apple-Google duopoly that has shaped mobile platform economics for the past fifteen years will not, in my forecast, simply transfer to the device landscape of 2030. The dominant platforms of the spatial era are still being decided. The firm that wants to control its own destiny in this transition must engage with the platform decisions now, not later, even if the platform endpoints are still being defined.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Re-baseline your firm's product strategy on the assumption that the smartphone is one device among four by 2030, not the dominant one.**

   The re-baselining will surface where your current strategy is dependent on smartphone-first assumptions and where it is not.

2. **Identify, for each product in your portfolio, which of the four device categories it should principally serve in 2030.**

   The answer is, for most products, a combination. The combination is the strategic question.

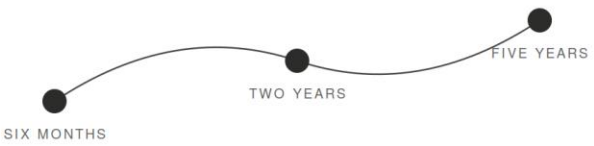
3. **Engage seriously with at least one platform decision affecting the spatial era this year.**

   The platforms are being decided now. Participation is more valuable now than it will be in 2028.

PART SEVEN

VII

The Leader's Playbook



The honest conversation about what this asks of leaders, what to do in the next six months, the next two years, and the next five — and the risks that must be designed for, not designed around.

THE WORLD THAT'S BEGINNING

The honest conversation about what this asks of leaders, what to do in the next six months, the next two years, and the next five — and the risks that must be designed for, not designed around.

CHAPTER THIRTY

30

The Honest Conversation About Work

What we owe each other in a decade when
the floor moves under millions of careers.

—

"The world is not yet finished with us. Nor we with it."
— Wendell Berry

◆ ◆ ◆

I want to begin this chapter where most leadership books end it, because the question I am about to ask is the question every other recommendation in this book ultimately rests on. The question is: what do we owe the people whose work, in this decade, will be reshaped or removed by the technologies we are deploying?

The question is uncomfortable. Most leadership conversations I am invited into, in 2026, treat it as a public-relations problem to be managed rather than a moral problem to be answered. The treatment is wrong. It produces speeches that read well in the annual report and decisions that read poorly in the lives they touch. It also, less obviously but more importantly, produces strategy that is not durable — because an organisation that has not honestly answered the question for itself will not have the moral authority, internally, to lead its people through the transition that is coming.

This chapter is the honest answer I have come to, after several years of being in the room when the question was being avoided. It is offered as a leader's answer rather than a policy treatise. The policy questions matter, and I will touch them briefly in the next chapter and again in the closing chapter on risks. But the leader's task, in 2026, is not principally a policy task. It is a moral task — the task of answering, for her own institution, what kind of employer it intends to be in a decade when the floor will move under millions of careers.

◆ ◆ ◆

I. What is actually happening to white-collar work

The factual situation, stripped of euphemism, is this. The kinds of cognitive work that have constituted the upper-middle of the global labour market for four decades — structured analysis, document drafting, routine professional services, mid-level management, customer support of moderate complexity, junior creative work — are, in 2026, in the early years of a transformation that will reduce the human headcount required to perform them by, in my honest forecast, somewhere between thirty and seventy percent over the coming decade. The range is wide because the variation across categories is wide. The direction of travel is uncontested.

This is not a forecast that requires AGI. It does not require any technology more advanced than what is already deployable in 2026 with deliberate engineering effort. It is the consequence of the agentic shift I described in Chapter Five, applied at scale across the categories of work most amenable to it.

The transformation is not uniform. Some categories will lose less than thirty percent and will largely re-shape rather than shrink. Other categories will lose more than seventy percent and will substantially disappear. The categories most exposed are, broadly, those whose work is structured, repeatable, transactional, and procedurally accountable — the work I named in Chapter Five. The

categories least exposed are those whose work involves genuine ambiguity, fiduciary accountability, contextual creativity, empathic presence, or physical embodiment in unpredictable environments. The middle, where most of the workforce sits, is the territory of redesign rather than replacement: the job survives, but its content changes substantially.

Three patterns are worth naming because they organise what is otherwise a bewildering landscape. The first is that *the work that disappears first is the work that was already being outsourced*, because outsourcing required the work to be structured for handoff, and the same structuring makes the work amenable to AI. Business process outsourcing, IT services delivery centres, shared service organisations — the categories India built much of its modern economy on — are among the most exposed. The second is that *the work that disappears most slowly is the work that requires physical presence in unpredictable conditions*, which means that the trades, skilled manual work, frontline care work, and craft work are, paradoxically, more durable than the office work that for decades commanded higher status and pay. The third is that *the work that grows is the work at both ends of the cognitive spectrum*: the highest-judgement strategic work expands because the agents amplify it, and the highest-empathy human work expands because the customer increasingly values the human moments more than ever before. The middle hollows.

This is what is happening. It is happening fast enough to be visible inside this decade and slow enough that the institutions that prepare deliberately can absorb most of it. The shape of the transition is, in significant part, in the hands of the leaders making decisions in 2026 and 2027.

◆ ◆ ◆

II. What the leader can choose

Within that situation, the leader has more agency than the macroeconomic framing suggests. The forces are large, but the

specific decisions made inside specific firms shape, materially, who absorbs the cost of the transition and who absorbs the gain.

Consider the two extreme postures a leader can take. The first posture: *capture the productivity gain entirely as headcount reduction*. The agents arrive. The firm reduces its workforce by the productivity multiplier. The freed margin flows to shareholders or is reinvested in expansion. The displaced workers absorb the cost. This is the posture most analyst models implicitly assume. It is also the posture that, in my view, no responsible leader should take, both because it is destructive of the social contract that ultimately licences the firm's operations and because it is, on a longer view, less commercially successful than the alternatives. A firm that has shed thirty percent of its workforce will, in five years, find that it has also shed thirty percent of its institutional memory, thirty percent of its customer relationships, and thirty percent of its ability to absorb the next transition. The cost is paid later but it is paid.

The second posture: *capture the productivity gain primarily as expansion of capacity, quality, and reach*. The agents arrive. The same workforce, augmented, can serve more customers, with higher quality, in more languages, at more hours, with more personalisation, with more attention to the moments that matter. The firm grows into the productivity gain rather than shrinking into it. The freed margin partially flows to shareholders, partially to customers as better service or lower prices, and partially to the workforce as expanded scope and elevated work. This is the posture that, in my view, most well-led firms in this decade will take, because it is both commercially superior over a five-to-ten-year horizon and morally defensible.

Most firms will, in practice, sit somewhere between these two extremes, and the position they take will be set, in significant part, by the leadership posture of the CEO and the board. The CEO who tells the board, in 2026, that the AI transformation is principally a margin-improvement story is taking the first posture by default.

The CEO who tells the board that the transformation is principally a growth-and-quality story, with margin improvement as a consequence, is taking the second posture by intention. The framing matters. It cascades through every subsequent decision.

◆ ◆ ◆

III. The four obligations of the responsible employer

I want to name, as concretely as I can, the four obligations I believe a responsible employer takes on in this decade. The list is not exhaustive. It is the floor.

The obligation of honest communication. The employee deserves to know, in honest terms, what the firm's plan is for the kind of work she does. If the work is being augmented and her role will expand, she should know. If the work is being substantially absorbed by AI and her role will shrink or change shape, she should also know. The firm that lets its employees discover the answer through layoffs has failed the obligation. The firm that communicates the trajectory openly, with enough lead time for the employee to act on the information, has met it. This is harder than it sounds. Most firms have not built the muscle of honest communication about disruptive technology change. They will need to build it, deliberately, in the coming year.

The obligation of meaningful reskilling. The employee whose work is changing deserves a real opportunity, funded by the firm, to acquire the new skills the firm now requires. The opportunity must be meaningful — paid time, quality programmes, a credible chance of reaching the new role — not the token reskilling allowance that most firms currently offer. The cost of meaningful reskilling per employee is, in my experience, between five and fifteen percent of one year's compensation. The cost of not doing it is the moral cost of discarding the employee and the strategic cost of losing the institutional memory she represented. Most firms currently underspend on reskilling by an order of magnitude. The leaders I most respect have, in 2026, begun to correct this.

The obligation of dignified transition for those who cannot be retained. Not every employee in a redundant role can be reskilled into a new one. Some roles disappear without replacement. Some employees, for entirely legitimate reasons, are not in a position to make the transition the firm needs. The responsible employer offers, in these cases, a transition that is dignified — adequate notice, generous severance, genuine career counselling, network access, retraining support, time to find the next thing. The employer who offers two weeks of severance to a fifteen-year employee whose role has been absorbed by AI has not honoured the relationship. The employer who offers six to twelve months of meaningful transition has. The cost is significant. The moral and reputational cost of not paying it is greater.

The obligation of voice in the transition. The employees whose work is being reshaped are the people who know that work most intimately. The firm that designs the transition without consulting them deeply will design it badly. The firm that includes them — through structured consultations, employee councils, technology committees with worker representation, pilots designed with workers as collaborators — will design it more wisely and execute it more smoothly. This is, in some firms, a cultural change as significant as the technology transition itself. It is also, I have come to believe, indispensable.

These four obligations are not heroic. They are, by historical standards, the floor of what an industrial-age employer was expected to provide its workforce during a major technology transition. The mid-twentieth-century American manufacturer, the mid-twentieth-century European industrial firm, the late-twentieth-century Japanese corporation — each had, with cultural variations, a version of these obligations. The obligations have weakened over the past forty years, in many parts of the world, under the pressures of shareholder primacy and short-termism. The convergence decade, in my view, is the moment at which these obligations need to be restored, not because of moral fashion, but

because the social licence of the firm requires them and because, on any honest accounting, they pay back over time.

◆ ◆ ◆

IV. The internal communication that has to happen

The practical question, for the leader who accepts the four obligations, is how to communicate them inside the firm in a way that is credible. Most firms in 2026 are conducting their AI conversations in two registers that do not connect: the executive register, in which AI is described as transformative, exciting, full of opportunity; and the worker register, in which AI is whispered about as threat, anxiety, the unspoken topic of every coffee-break conversation. The two registers do not meet. The disconnect is corrosive.

The leader's task is to bring them together — to have, openly, the conversation the workforce is already having quietly. The conversation has, in my experience of leading it inside client organisations, three components. First: an honest naming of what is happening, in plain language, without euphemism. This work will change. Some of you will see your roles expand. Some of you will see your roles change substantially. Some of you, in candor, will see your roles diminish or disappear. Here is what we know about which is which. Second: an honest naming of what the firm intends to do about it. We will invest in reskilling at this level for these categories. We will offer this kind of transition to employees whose roles cannot be retained. We will involve you in the design. We will not pretend the change is not happening. Third: a sustained programme of follow-through over years, not quarters, that demonstrates the commitments are real.

The leaders who have done this well — and I have, by now, watched several do it well — have found that the workforce response is not the resistance they had feared. The response is, more often, relief. The workforce had been carrying the anxiety alone. Sharing the conversation, sharing the plan, sharing the

agency in the transition, lightens the load on everyone. The firm that does this becomes, over time, a place where people want to work through the transition rather than a place they are trying to escape from before the transition reaches them. That is a competitive advantage in talent markets that will, in the coming years, be acutely tight for the categories of work that remain.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Hold an honest, all-hands conversation about AI and the future of work in your organisation within the next ninety days.**

   Speak in plain language. Name what is happening. Name what the firm intends to do. Take questions in public. The conversation is more important than the answers; the answers can be refined as the firm learns.

2. **Commit, in writing and in budget, to meaningful reskilling for any employee whose role is being substantially reshaped by AI.**

   Five to fifteen percent of one year's compensation per affected employee, deployed as paid time and credible programmes, is the order of magnitude. Less than this is theatre.

3. **Build an employee council on technology change, with real consultation rights, within the year.**

   The council changes the design quality of every subsequent technology decision. The cost is small. The dividend, over five years, is large.

CHAPTER THIRTY-ONE

31

The First Six Months

Foundation. Ten moves any leader can make this quarter.

—

"Vision without execution is hallucination."

— *attributed to Thomas Edison, more sharply than he probably said it*

◆ ◆ ◆

Every consulting framework I have ever seen for technology transformation makes the same mistake. It begins with the five-year vision, descends to the three-year roadmap, and then proposes a one-year operating plan. The mistake is structural. The vision is the easiest part. The operating plan is the hardest. By the time the leader has traversed the framework she has already lost the first six months of execution to slide preparation.

This chapter is the inverse. It begins with the next six months. It assumes the reader has, by this point in the book, internalised enough of the strategic argument to know roughly where she is going. It now offers the ten foundational moves that, in my experience working alongside dozens of leadership teams through this transition, separate the firms that compound from the firms that drift. The moves are not heroic. None of them requires capital expenditure that is unavailable in a normal operating budget. Most can be initiated by a CEO or a senior leader within a single quarter, with deliberate intent.

The reason to do them in the first six months rather than the second is mathematical. The capabilities and the market positions of 2030 will compound from the foundations laid in 2026 and 2027. A firm that begins the foundational moves in 2027 instead of 2026 will be, in 2030, somewhere between twelve and twenty-four months behind a firm that began earlier. In a decade in which capabilities double approximately every six months, twelve to twenty-four months is the difference between leading the category and being commoditised by it.

◆ ◆ ◆

I. The ten moves

I will take the moves in the order in which I recommend them be initiated, with the caveat that several can run in parallel and should. The order is not strictly sequential. It is approximately the order of precedence — the earlier moves enable the later ones.

Move 1: Build your AI fluency, personally, in the next sixty days. The single highest-leverage investment a CEO can make in the first six months is in her own working fluency with the technology. The reason is that AI strategy cannot be delegated to people with greater fluency than the CEO has, because the CEO will not be able to evaluate the proposals. Use the systems daily. Read the credible commentary weekly. Take a hands-on course over a long weekend. Build your own first agent. The investment is forty to sixty hours of personal time over the first quarter. The dividend is the ability to lead the rest of the moves with confidence.

Move 2: Map your business to the experience stack and the value migration. In one or two leadership offsites, work through where each of your major product lines currently sits — on the five-layer stack of Chapter Ten and on the value migration of Chapter Three. The output is not strategy. The output is shared understanding, captured on a single page, of where the firm's capabilities are

concentrated and where they are absent. The shared understanding is the foundation for every subsequent decision.

Move 3: Identify your three highest-impact AI use cases — and your three biggest risks. This is not a comprehensive scan. It is a focused identification of the three opportunities most likely to produce visible value in the first year and the three risks most likely to materialise as harm if unaddressed. Three and three. Each captured on a single page. Both lists assigned an executive owner. Both lists reviewed monthly.

Move 4: Build your data foundation — auditable, governed, integrated. Most AI initiatives that fail in the first eighteen months fail because the data foundation was not strong enough to support them. The first six months should include a deliberate audit of where the firm's data lives, who owns it, how it is governed, and what work is needed to make it usable for AI applications. The audit is not glamorous. It is non-negotiable. The firms that take this seriously in 2026 will be the firms that move fastest in 2027.

Move 5: Form a small, high-trust AI council reporting directly to the CEO. Three to seven people. A mix of business unit leaders, technology leaders, and outside expertise. The council's mandate is to drive the AI transition across the firm — not a technology committee that recommends, but a leadership body that decides. The council should meet weekly for the first six months and at least biweekly thereafter. The CEO should be in the room.

Move 6: Run two pilots — one customer-facing, one internal. The customer-facing pilot demonstrates to the market and to the workforce that the firm is serious. The internal pilot — typically agentic automation of an end-to-end internal process, as I described in Chapter Five — produces operational learning and visible productivity gain that funds the next round. Both pilots should be small enough to ship in ninety days and large enough to

matter to the people who execute them. Document both. Publish the results internally.

Move 7: Begin the reskilling commitment publicly. Announce, in writing, what the firm will fund in reskilling for which roles, on what timeline, with what guarantees. The announcement is itself an act of leadership. The workforce hears it. The market hears it. The competitors hear it. The firm is now committed to behaving in a particular way, and the commitment shapes every subsequent internal decision. Most firms have not yet made this announcement in any specific form. The firms that do, ahead of the wave, will be the firms talent flocks toward.

Move 8: Establish your AI governance and ethics charter. A short document — five to fifteen pages — that captures the firm's principles for AI use, its red lines, its accountability structures, its disclosure norms, its handling of customer data, its position on the cryptography migration. The document is not a public-relations exercise. It is the operating manual for the AI council and the basis for every subsequent technology decision. Most firms in 2026 do not have one. Most firms by 2027 will need to.

Move 9: Audit one cryptographic dependency for the post-quantum migration. I described this in Chapter Seven. The first six months is the right window to begin the inventory. Pick one critical system. Inventory the cryptographic dependencies. Identify the migration path. The first migration is the hardest. Subsequent migrations are projects rather than discoveries.

Move 10: Schedule the next quarterly review now, and set the agenda. The first six months is, in some ways, the easiest. The capabilities are visible. The team is energised. The wins are accessible. The discipline that compounds across the decade is the rhythm of review — every quarter, the same questions, the same accountability, the same willingness to update the plan as the technology changes. Set the rhythm now. Defend it later.

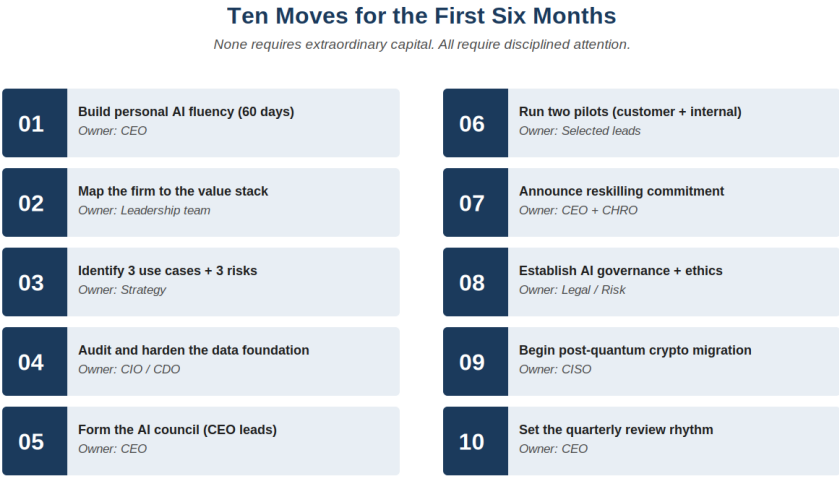


Figure 31.1

◆ ◆ ◆

II. What you do not do in the first six months

Equally important as the ten moves is the discipline of what *not* to do. Several actions tempt leaders in the first six months that, in my experience, almost universally consume capacity without producing returns.

Do not commit to an enterprise-wide AI platform. The technology is moving too fast. The platform you commit to in 2026 will, with high probability, not be the platform that wins your category in 2028. Use multiple foundation models through standard APIs. Maintain optionality. Decide the platform question in 2027 or 2028, when more is known.

Do not announce a sweeping reorganisation. The temptation to restructure the entire firm around AI is, in 2026, premature. The shape of the firm in 2030 will be different, but the shape is not yet legible enough to commit to. Make the small, foundational moves above. Allow the larger restructuring to emerge from the learning

the moves produce. The firms that re-organised early have, in my observation, mostly had to re-organise again.

Do not over-invest in flagship AI products at the expense of the foundation. The CEO's ego is naturally drawn to the flagship product. The firm's value is built in the foundation. Spend on the foundation in the first six months. The flagship can come in months nine through eighteen.

Do not delegate the leadership of the transition. The CEO who hands the AI agenda to a Chief AI Officer and disengages will, in twelve months, find that the Chief AI Officer has built a centre of excellence that the rest of the firm has not absorbed. AI strategy belongs in the CEO's office, with the CIO and the Chief AI Officer as senior partners, not as primary owners.

◆ ◆ ◆

III. The mood of the first six months

The first six months should feel, inside the firm, like a quiet acceleration rather than a noisy one. The leader who fills the period with town halls, big announcements, splashy launches, and external press is signalling activity rather than building capability. The leader who fills the period with disciplined foundational work — the ten moves above — and communicates the work openly to the workforce is building the substrate from which the noisy launches of 2027 and 2028 will become possible.

The firms that have done this well, by my observation, share a particular tone. Their leadership talks about AI confidently but without theatre. Their workforce understands what is happening because the leadership has been explicit. Their internal pilots are visible, and their results are shared. Their failures are also shared, without blame, as learning. Their cadence is steady. They are, six months in, quietly ahead of competitors who chose noise over discipline.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Within the next thirty days, schedule Move 1 — your personal sixty-day fluency programme — and tell your leadership team you have done so.**

   The signalling is half the value. The fluency is the other half.

2. **Within sixty days, identify the three executive owners of Moves 5, 6, and 8 (council, pilots, governance) and put them in place.**

   Without these three, the rest of the foundation does not hold.

3. **Within ninety days, hold the first all-hands conversation described in Chapter Thirty.**

   The conversation is itself one of the foundational moves. It cannot be postponed.

CHAPTER THIRTY-TWO

32

The Twelve to Twenty-Four Months

Re-architect. The structural changes that take a year to design and a year to install.

—

"We shape our buildings, and afterwards our buildings shape us."
— Winston Churchill

◆ ◆ ◆

The first six months are about foundation. The next eighteen — broadly the second half of year one and all of year two — are about re-architecture. This is the period during which the firm makes the structural changes that, by the end of year two, have visibly changed the shape of the company. The window is narrow. Eighteen months is barely enough to design these changes deliberately, build consensus across the leadership team, sequence the implementation, and absorb the disruption. A leadership team that delays starting these changes until year three is, in my experience, taking two more years than they have to spare.

I will name six structural changes that, in my honest view, every serious enterprise will need to make in this window. Each is hard. Each is the kind of change a firm can plausibly make once a decade. The case I want to argue is that this is the decade.

◆ ◆ ◆

I. Re-architect the data layer

Most firms in 2026 are operating data architectures that were designed for the world of business intelligence reporting — dashboards, monthly reviews, regulatory filings. The architecture was designed to produce reports for human consumption. It was not designed to feed AI systems that consume data continuously, bidirectionally, in real time, across all the firm's domains simultaneously.

The re-architecture is not glamorous. It involves data lakes, real-time pipelines, semantic layers, vector stores, governance models, and the painstaking work of translating decades of transactional data into formats AI systems can consume. The re-architecture is also non-negotiable. A firm whose data layer cannot serve its AI systems will find that its AI strategy is structurally limited regardless of how much it spends on the AI itself. The data layer is the floor; the AI is the ceiling. The ceiling can rise no higher than the floor allows.

The work, done deliberately, takes between twelve and eighteen months for a mid-sized firm and between two and three years for a large one. The work, done poorly or piecemeal, takes longer and produces less. The CIO and the Chief Data Officer should, by the end of year one, have a credible plan and an executive owner for every layer of the re-architecture. By the end of year two, the foundational layers should be operating.

◆ ◆ ◆

II. Re-architect the operating model around outcomes

The organisation chart Meera's HR director laid on her conference table in Chapter Fourteen will not, in 2030, be the organisation chart of the firm that retains its market position. The functional, hierarchical, role-based organisation that defined twentieth-century enterprise is, in the convergence decade, ill-suited to the

work that must be done. The replacement is the outcome-based, fluid-team, AI-augmented organisation I described in that chapter.

The transition does not happen in a quarter. It happens over eighteen to twenty-four months, in stages. Begin by composing a small number of outcome teams as pilots. Run them in parallel with the existing functional structure. Document what works and what does not. Expand the pilots. By the end of year two, the firm should be operating in a hybrid mode — significant parts of the work organised around outcomes, with the functional structure providing the disciplinary depth that the outcome teams draw from. The full transition continues into years three and four.

The leadership challenge in this re-architecture is significant. The functional leaders — the heads of marketing, finance, operations, technology — are accustomed to running their functions and reporting through the hierarchy. The outcome teams cut across this. The CEO's task is to redesign the senior leadership team's incentives and responsibilities so that the outcome teams have the authority they need without dismantling the functional capabilities the firm depends on. The redesign is harder than it sounds. Most firms get it wrong on the first attempt and have to revise it in year two.

◆ ◆ ◆

III. Re-architect the customer relationship as continuous

The customer relationship in most firms today is a sequence of discrete encounters: marketing exposure, lead capture, sales conversation, transaction, fulfilment, occasional support, repurchase. The encounters are stitched together by a CRM system that captures the data but does not, in most firms, animate the relationship. The customer experiences the firm as a series of separate touchpoints rather than as a continuous presence.

The re-architecture is to treat the customer relationship as a continuous, intelligent, adaptive presence in the customer's life — through the AI agents the customer increasingly uses, through the ambient interfaces in her environment, through the proactive value the firm can now deliver continuously rather than transactionally. This is operationally different from the old CRM. It requires the firm to have, for each customer, a continuously updated understanding of her state, her preferences, her current goals; to have, in production, the agentic capability to respond to that state; and to have the trust currency to be present in the customer's life without being intrusive.

The work, again, takes twelve to eighteen months to design and execute well. By the end of year two, the firm should be operating its top customer segments under the continuous-relationship model, with the rest of the customer base on a clear migration path.

◆ ◆ ◆

IV. Re-architect the talent model

The talent practices of most firms in 2026 — recruiting, onboarding, performance management, compensation, career paths, learning and development — were designed for the work that is being reshaped. They are increasingly miscalibrated. The recruiting filters select for credentials that no longer reliably predict capability. The performance management cycles measure outputs that increasingly are not the outputs the firm needs. The career paths assume role hierarchies that are dissolving.

The re-architecture is the move I described in Chapter Fourteen, taken seriously. Skills graphs replace CVs as the primary unit of professional identity. Performance management measures contribution to outcomes rather than execution of tasks. Compensation rewards the irreducibly human capabilities I described in Chapter Twenty-Eight. Career paths are reconceived as portfolios of capabilities the worker takes with her across many

assignments. Learning and development is reconceived as continuous rather than episodic.

The work belongs in the second year, not the first, because the data foundations of skills graphs and outcome-based performance need to be in place before the talent practices can be rebuilt on top of them. By the end of year two, the firm should have piloted the new talent practices on a meaningful subset of the workforce, with a credible plan to scale across the third and fourth years.

◆ ◆ ◆

V. Re-architect the supplier and partner relationships

The partner ecosystem most firms operate in is, in 2026, a set of bilateral commercial relationships managed by procurement and partner management functions. The partners are vendors, distributors, technology providers, agencies. The relationships are governed by contracts, SLAs, and quarterly reviews. The relationships were designed for the world in which the firm sold a product through a channel.

The re-architecture is toward what I think of as ecosystem orchestration. The firm operates as the conductor of a network of partners — some of whom contribute components, some of whom deliver the customer experience, some of whom hold the customer relationship, some of whom provide the AI infrastructure, some of whom validate trust. The orchestration is itself the firm's central capability. The firm does not try to do everything inside its own walls. It does try to ensure that the integrated experience the customer encounters is coherent, accountable, and compounds.

The implication is a change in how the firm relates to its partners. Bilateral procurement-style relationships give way to ecosystem-style alliances with shared data, shared customer relationships, shared accountability for outcomes. The partner contracts evolve. The partner management function evolves. The firm's posture in

the ecosystem evolves. By the end of year two, the firm should have re-platformed its three or four most strategic partner relationships under the new model.

◆ ◆ ◆

VI. Re-architect the board

The board of directors is, in most firms in 2026, the institutional body least equipped to govern the transition that is now underway. The board's typical composition — a mix of senior executives, financial expertise, industry experience, and a small number of independent voices — was assembled for the firm's previous era. It was not, in most cases, assembled with the convergence decade in mind.

The re-architecture is to refresh the board for the period ahead. This typically requires three or four new appointments over eighteen months: a director with deep AI expertise, a director with cybersecurity and privacy expertise, a director with experience of major technology transitions, and — for many firms — a director with deep customer experience expertise. The refresh is uncomfortable for the existing board. It is also non-negotiable for any firm whose board is going to give the executive team useful counsel through the transition.

The CEO who leaves the board composition unchanged through this decade is governed by a body that does not understand the period. The CEO who refreshes the board deliberately, in conversation with the chairman, is governed by a body that can hold the executive team to a higher standard and offer wiser counsel. The difference, by year three of the transition, is significant.

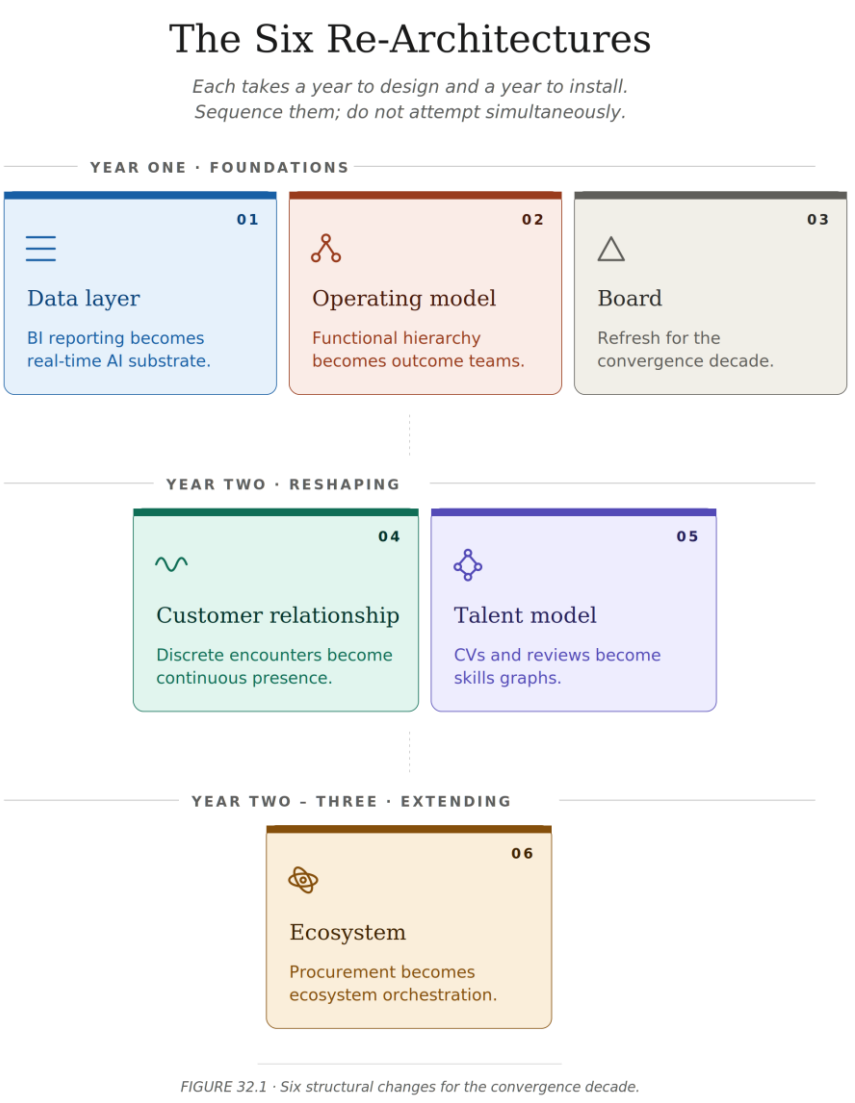


Figure 32.1

◆ ◆ ◆

VII. The discipline of sequencing

The six re-architectures cannot be done in parallel. Each is large. Each consumes leadership attention. The firms that try to attack all six simultaneously, in my observation, complete none well. The

firms that sequence them — typically beginning with data and operating model in year one, customer relationship and talent in year two, ecosystem and board in year three — complete the transition with manageable disruption.

The CEO's task is to hold this sequencing against the natural impulse to attack everything. The impulse is understandable; the gap between where the firm is and where it needs to be is alarming, and the leader's instinct is to close the gap fast. The discipline is to recognise that the gap closes faster through sequenced execution than through simultaneous chaos. The leader who can hold this discipline, against the noise, will deliver the transition. The leader who cannot will, by year three, be running a firm that has begun every transformation and finished none.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify, in your firm, which of the six re-architectures is most pressing — and commit to leading it personally.**

   The CEO who delegates the most pressing re-architecture is delegating the firm's future.

2. **Sequence the other five over the eighteen-month window with clear executive owners and quarterly milestones.**

   The sequencing is the strategy. The execution follows.

3. **Refresh your board this year with at least one new director who deeply understands the convergence.**

   The board refresh is the leverage point with the highest return on the smallest amount of CEO attention.

CHAPTER THIRTY-THREE

33

The Three to Five Year Bets

The strategic positions worth taking. The bets that compound.

—

"In the long run, we are all dead. But the long run is shorter than we think, and our successors live in it."

— paraphrased from Keynes

◆ ◆ ◆

If the first six months are about foundation and the next eighteen are about re-architecture, the three-to-five-year horizon is about strategic position. This is the horizon at which the bets the firm places now compound into the categories the firm will own in 2030 and beyond. Most firms place these bets implicitly — through a thousand small decisions whose cumulative direction is rarely stated — and end up, by 2030, with a strategic position they did not deliberately choose. The leaders who place these bets explicitly, in 2026 and 2027, will own clearer ground.

This chapter names six strategic bets that, in my view, deserve serious consideration from any enterprise of meaningful scale. Not every firm will take every bet. Each firm will choose, given its category, its capabilities, its capital, and its appetite, which subset to take. The point is not the specific list. The point is that the bets must be chosen, consciously, with the discipline of a portfolio rather than the accident of accumulated tactics.

◆ ◆ ◆

I. The bet on owning the customer's primary AI

The most consequential strategic bet of the next decade may be the bet on whose AI the customer principally interacts with. By 2030, every customer will have one or two AI companions she relies on for a substantial portion of her daily decisions — what to buy, where to go, who to consult, how to manage her money, when to see a doctor. The firm whose AI is in this position has access to the customer that no traditional channel can match. The firm whose AI is not in this position is, increasingly, on the wrong side of the customer relationship.

The bet has three credible postures. The first is to *build the primary AI ourselves* — appropriate only for firms with extraordinary scale and existing customer trust at platform level. Most firms cannot. The second is to *partner with the AI providers and ensure deep integration into their AI's recommendations* — the strategy that most firms will implicitly take, and the one in which the firm becomes a participant rather than the orchestrator of the relationship. The third is to *build a vertical AI for our category* that is sufficiently better than the general-purpose AI for the customer's needs in that category to win the customer's primary relationship within the vertical, even if not across her whole life. The third strategy is, in my view, the most underexplored and the most promising for firms of mid-tier scale.

The bet must be placed in 2026 or 2027. The vertical AI built in 2027 will, in 2029, have the customer relationships and the data that a competitor starting in 2029 cannot replicate. The window for placing the bet credibly closes within twenty-four months.

◆ ◆ ◆

II. The bet on infrastructure and capacity

For Indian firms in particular, but increasingly for firms in any country that does not have foundational AI capacity, there is a strategic bet to be placed on building the infrastructure on which the convergence decade will run. Compute capacity. Data centres. AI hardware. Foundation models trained on Indian data, in Indian languages, for Indian use cases. The capacity, in 2026, is concentrated almost entirely in two countries. The dependency this creates is structural and consequential.

The bet on building Indian infrastructure is not a bet for every firm. It is a bet for the firms with the scale, the capital, and the strategic imagination to participate in building the substrate of the country's convergence economy. The firms that have begun to take this bet — a small number of conglomerates, a small number of new ventures, a small number of public-private partnerships — are placing capital in 2026 that will, by 2030, have produced infrastructure no one else in India will have. The returns will be commensurate with the position.

The bet for smaller firms is to *participate as customers and partners of this infrastructure as it builds* — to use Indian foundation models where possible, to host on Indian compute infrastructure where possible, to contribute to the data ecosystem that the infrastructure depends on. The participation is itself a form of the bet. It signals the kind of country and ecosystem the firm wants to operate inside.

◆ ◆ ◆

III. The bet on programmable biology

Programmable biology is one of the four forces I treated briefly in Chapter Two, and it is, of all of them, the one most under-discussed in business circles. The reason is that it does not yet have a clear commercial template for non-pharma firms. The template will

emerge in this decade, and the firms that have begun to engage early will be the firms that define it.

The bet is to *develop a credible biology capability within the firm or a credible partnership with the right biology firms, before the category arrives*. For consumer firms, this means understanding longevity, personalised health, microbiome optimisation, precision nutrition, regenerative therapies. For industrial firms, it means understanding bio-manufacturing, biological materials, agricultural biology. For technology firms, it means understanding the convergence of AI and biology — the AI tools that are accelerating biology, and the biology applications of AI that are now becoming possible.

The bet is sized smaller than the previous two bets, because the commercial inflection is later. The bet is more strategic than tactical. The firm that has, by 2028, a credible biology capability will be the firm that participates in the wave when it arrives. The firm that does not will be the firm watching the wave from the shore.

◆ ◆ ◆

IV. The bet on energy and the climate transition

The energy transition is a five-trillion-dollar-per-year industrial reorganisation. It is also, for India in particular, a strategic moment of the same character as the IT services moment of the 1990s — a moment in which the country's combination of demand, talent, and policy posture create a window for national champions to emerge.

The bet, for firms of significant scale, is to *participate seriously in the energy and climate transition* — through investment, through procurement of green energy, through development of climate-adaptive infrastructure, through products that serve the transition. The participation is partly a moral obligation, partly a regulatory inevitability, and partly — for the firms that take it seriously enough — a strategic position. The Indian firm that builds, by 2030,

a credible position in distributed solar, in grid storage, in green hydrogen, in climate-adaptive agriculture, in sustainable urban infrastructure will be participating in the largest industrial wave of the next twenty years.

The bet, for smaller firms, is to *position the firm credibly within this transition* — to make procurement choices that align with the transition, to develop products that are compatible with it, to build the brand position of a firm that is on the right side of the energy decade. The bet is small in any single quarter. It compounds over a decade.

◆ ◆ ◆

V. The bet on Bharat-first design

The civilisational argument I made in Chapter Twelve has commercial implications beyond Indian markets. The brand or product designed deliberately from Indian civilisational substrate, with rigour and authenticity, has access to a creative resource that competitors do not. The bet is to *make Bharat-first design an explicit strategic capability of the firm* — to hire the humanities scholars, the regional design talents, the vernacular creators, the cultural advisors who can give the firm's products the depth that conventional design cannot.

The bet is small in capital. It is large in cumulative effect. The firms that have, by 2030, built genuine Bharat-first design capability will be the firms whose products carry a distinctive resonance — in India and, for the right products, in global markets seeking authenticity. The bet is, in some sense, a counter-bet to the homogenisation that AI-generated content might otherwise produce. The category that machine-generation cannot easily replicate is the category of genuinely culturally rooted creative work. Build that capability deliberately.

◆ ◆ ◆

VI. The bet on trust as institutional capital

The sixth bet is the most diffuse and, in my view, the most important. It is the bet that, in the synthetic-content era, the institutions that have accumulated trust will hold an asymmetric position that money cannot directly purchase. The bet is to *invest in the trust currency deliberately*, over years, through every operating decision the firm makes — the disclosures it offers, the data practices it adopts, the audits it commissions, the public commitments it makes, the failures it owns, the consistency it maintains across cycles.

The investment does not show up in any quarter's results. It shows up cumulatively, after several years, in the form of customer relationships that competitors cannot disrupt at any reasonable price, talent that joins because of the firm's reputation, regulators that grant the firm permission to operate where competitors are denied, partners that prefer the firm to alternatives. By 2030, the firms that have accumulated trust capital will, in many categories, hold the dominant positions. The firms that have not will, in those same categories, struggle to hold the customer relationships the trust would have anchored.

This is the bet least visible in any quarterly review and most consequential in any decade-long horizon. It is the bet I would most strongly urge the reader to take seriously, if she has not already, because it is the bet whose returns are real, whose returns compound, and whose returns are not available to a competitor who decides, three years from now, to imitate the strategy.

[Figure 33.1 — to be supplied by illustrator. SIX STRATEGIC BETS, ARRAYED BY HORIZON. A two-axis chart. Vertical axis: capital required (low to high). Horizontal axis…]

◆ ◆ ◆

VII. The bets you do not take

Not every firm should take every bet. Most firms cannot take more than two or three. The discipline of the strategic horizon is to choose deliberately. The leader's task is to identify the two or three bets that, given the firm's capabilities and capital, will compound most powerfully — and to walk away from the rest with intent rather than by default.

The non-bets are also positions. The firm that walks away from infrastructure, from biology, and from energy is taking the position that it will buy these capabilities from others. That position can be coherent. It is also a position that requires the firm to be exceptionally good at the bets it does take, because the bets it does not take will be its dependencies. The position must be examined with the same rigour as the bets themselves.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Choose two of the six bets — no more, no fewer — for explicit strategic commitment over the next three to five years.**

   The discipline of two forces a choice. The discipline of choice is the substance of strategy.

2. **For each bet you do not take, name explicitly the partner or supplier on whom the firm will depend.**

   The non-bet is a dependency. The dependency must be named.

3. **Begin the trust-capital investment regardless of the other bets you take.**

   It is the bet whose absence is most easily mistaken for prudence and whose absence is most costly over a decade.

CHAPTER THIRTY-FOUR

34

The Risks We Must Design For

Not the risks we hope to avoid. The risks we must design our institutions to absorb.

—

"Hope is not a strategy."

— *Vince Lombardi, by way of every consultant who has ever written a risk report*

◆ ◆ ◆

I want to close the playbook with the chapter that, in most leadership conversations I am invited into, gets postponed to the end of the agenda and then run out of time. The chapter on risk. I want to take it seriously, because the risks of the convergence decade are not the risks of previous technology transitions, and the institutions that prepare for the wrong risks are the institutions that will be surprised when the right ones materialise.

The risks I will name are not exhaustive. They are the seven that, in my honest view, deserve sustained leadership attention through this decade. Each is, to some degree, already visible in 2026. Each is, to some degree, avoidable through deliberate institutional design. The leader who designs for them — rather than hoping they will not arrive, or planning to react when they do — will be better positioned to absorb them when they materialise.

◆ ◆ ◆

I. The misinformation collapse

The first risk is that the trust currency I described in Chapter Nine collapses faster than the institutions that depend on it can adapt. Synthetic content production is, by 2026, capable of fooling the median consumer for the median piece of content. By 2028, it will be capable of fooling the trained observer for most pieces of content. The asymmetry between the cost of producing convincing synthetic content and the cost of detecting it is widening, and the widening favours the producer.

The institutional implication is that, without deliberate intervention, the public's ability to know what is true about anything mediated through digital channels will decline materially. Elections, public health, financial markets, social cohesion — each depends on a base level of shared epistemic ground that synthetic content erodes. The risk is not hypothetical. It is materialising now, in 2026, in every major democracy.

The design response, at the institutional level, is to build the provable-truth infrastructure I described in Chapter Twenty-Four — content provenance standards, cryptographic verification, accountable institutions for verified content distribution, regulatory frameworks that require disclosure of synthetic content. Each firm has a role to play. Media firms must adopt provenance standards. Financial firms must verify the authenticity of communications they receive. Educational firms must teach the new media literacy. Governments must regulate. None of these is sufficient alone. All of them together may be, just barely, sufficient.

◆ ◆ ◆

II. The skills displacement crisis

The second risk is that the displacement of work by AI happens faster than the reskilling and absorption capacity of the labour market can manage. The numbers I described in Chapter Thirty are large. If the displacement curve runs ahead of the reskilling curve

by even five years — entirely possible if firms behave badly — the consequence is an extraordinary spike in technological unemployment among middle-class knowledge workers. The political, social, and economic consequences of such a spike are, in my honest view, civilisationally serious.

The design response operates at three levels. At the firm level, the four obligations of Chapter Thirty — honest communication, meaningful reskilling, dignified transition, voice in the design — must be honoured. At the policy level, the unemployment insurance, retraining programmes, and labour market interventions need to be substantially upgraded for a transition of this scale. At the cultural level, the social meaning of work itself may need to be re-examined; the Western, in particular American, identification of meaning with paid employment may not survive a transition in which paid employment is structurally less abundant for some categories.

This is the risk India is most exposed to among the major economies, because of the size and the structure of the country's white-collar and service-economy workforce. It is also the risk for which India has the most policy room to maneuver, because the country's institutional capacity for large-scale public programmes — proven through the digital public infrastructure of the past decade — is, in my view, equal to the task if it is properly directed.

◆ ◆ ◆

III. The cybersecurity escalation

The third risk is that the AI capabilities that are transforming legitimate work are equally available to malicious actors. The AI agents that can compose persuasive text, write functional code, conduct convincing voice conversations, and execute multi-step plans can be — and are being — turned to fraud, social engineering, sophisticated cyberattacks, deepfake-based identity theft, and automated misinformation campaigns. The asymmetry between

attacker and defender, in some categories, has shifted toward the attacker.

The design response is in three parts. *Cryptographic infrastructure* must be modernised — both to harden against quantum threats, as I described in Chapter Seven, and to enable the verification of identity and content that the synthetic-content era requires. *Defensive AI* must be deployed alongside the offensive AI; the same capabilities that enable attack can, with deliberate engineering, enable detection and response at the speed and scale required. *Institutional readiness* must be built — incident response capabilities, mutual aid networks among firms, public-private cooperation on threat intelligence, regulatory frameworks for accountability.

For Indian firms specifically, there is an additional concern: much of the country's economic infrastructure — payments, identity, public services — has been built on digital rails that, while extraordinary in their reach and elegance, have not yet been hardened against the threat profile of the late 2020s. The hardening is a national-scale project, not a firm-level one, but it is the project that protects the country's digital advantages from the cybersecurity threats that increasingly target them.

◆ ◆ ◆

IV. The concentration risk

The fourth risk is that the convergence decade produces a concentration of economic and informational power, in a small number of firms and a small number of countries, that is structurally unhealthy for the rest of the global economy. The dynamics of foundation models, compute infrastructure, and platform-level AI naturally tend toward concentration, because the marginal returns to scale are immense and the barriers to entry are increasingly capital-intensive. Without deliberate intervention, the concentration will deepen through this decade.

The design response, at the firm level, is to maintain optionality — to use multiple foundation models, to host across multiple compute providers, to build core competencies that do not depend on any single platform. At the country level, the response is the infrastructure-and-capacity bet I described in the previous chapter, taken seriously by enough firms and enough countries that the concentration is meaningfully diluted. At the multilateral level, the response is the harder work of international agreements, antitrust frameworks, and standards bodies that prevent the new platforms from acquiring the kind of unchecked market power that the previous generation of platforms briefly enjoyed.

The risk is not theoretical. The concentration is occurring now. The question is whether it produces, by 2030, a digital order that is broadly accessible and competitive, or a digital order in which most economic activity flows through infrastructure controlled by a small number of firms in two countries. The answer depends on what the rest of the world does in the next twenty-four months.

◆ ◆ ◆

V. The accountability gap

The fifth risk is the gap between the systems that act and the institutions that are accountable for what those systems do. As agentic AI takes more autonomous action — the Mumbai investment firm scenario from Chapter Five, but at scale, across categories — the question of who is responsible when things go wrong becomes structurally harder. The agent acted. The agent was deployed by the firm. The firm relies on a foundation model from another firm. The model was trained on data the model provider does not entirely control. The agent's specific action was not, in the strict sense, foreseen by any human. Who is accountable?

The design response is to be explicit, in the firm's own governance, about accountability before the action is taken rather than after. Every agentic deployment must have a named human owner. Every category of decision the agent can take must be specified in

advance, with clear delegation boundaries. Every failure mode must have a defined response and a defined responsible party. The accountability chart must be as clear as the org chart. Most firms in 2026 do not have one. Most firms by 2027 will need one.

At the regulatory level, the question of liability for AI actions is being worked through in major jurisdictions, with significant variation across countries. Indian firms operating internationally must track the variation. The firm whose products are used in Europe is governed by one set of rules; the same firm in India by another; the same firm in the United States by a third. The complexity is real. The simplest defensive posture is to operate to the strictest standard across jurisdictions, even where local rules permit less. The reputational cost of an accountability failure is global; the regulatory cost is local. The defence must be planned for the global cost.

◆ ◆ ◆

VI. The geopolitical fragmentation

The sixth risk is the fragmentation of the global digital order into competing technological blocs, with reduced interoperability, divergent regulatory regimes, and increasing pressure on firms to choose sides. The fragmentation is already underway. By 2030, in my forecast, the global firm will operate across three or four substantially incompatible technology stacks, each with its own foundation models, its own data residency requirements, its own privacy and AI rules, its own approved hardware sources.

The design response, for the firm, is to plan for fragmentation rather than to hope it does not arrive. Multi-stack architectures must be designed in. Data residency must be structurally compliant. Localisation must be planned. The cost is real — fragmentation will materially increase the operating cost of global enterprises — but the cost of being unprepared is greater. The firms that have begun the multi-stack design now will, in 2030, be

operating routinely across the fragmented order. The firms that have not will be in expensive remediation.

For Indian firms, the fragmentation is, in some respects, an opportunity. India is one of the few large economies with the scale, the diplomatic posture, and the technical capacity to operate across multiple blocs. Indian firms with global ambitions can — with care — position themselves as the bridge between blocs that other firms cannot easily span. The position requires deliberate strategy and significant institutional discipline. It is not automatic. It is, however, available.

◆ ◆ ◆

VII. The deepest risk: existential

The seventh risk is the one that strategic literature has the hardest time treating, because it borders on the speculative and threatens to undermine the credibility of every other risk on the list. I will name it briefly and honestly. The risk is that the development of progressively more capable AI systems leads, in some plausible scenarios, to outcomes that are not merely disruptive but catastrophic at the level of the species — through accident, through misuse, through misalignment of advanced systems with human interests, or through the political instability that any of these could produce.

I do not believe these scenarios are the most likely outcomes of the next decade. I do believe they are sufficiently possible, and the consequences sufficiently severe, that institutions of every scale should treat the question with the seriousness it deserves. At the firm level, this means engaging with the safety practices of the AI providers one uses, supporting the safety research community, refusing to develop or deploy AI in ways that would unreasonably increase catastrophic risk. At the policy level, it means supporting the regulatory and international coordination that the safety community has been calling for. At the cultural level, it means refusing the false dichotomy that says we must choose between

technological progress and safety; the leaders who have, historically, been the most careful about safety are also the leaders who have, historically, produced the most durable progress.

This is, I recognise, the most uncomfortable risk to name in a leadership book. I name it because the alternative — pretending the risk is not real, or passing the risk to actors with less capacity to address it — is, in my honest view, an abdication of leadership. The institutions that take the question seriously are part of the answer. The institutions that do not are, however inadvertently, part of the risk.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Conduct a risk audit against the seven categories above within the next ninety days, with executive ownership of each.**

   The audit will surface the risks that are visible to the firm and the gaps where the firm has been unconsciously hoping rather than designing.

2. **For each risk for which the firm is not yet ready, name the specific institutional capability that must be built, and the timeline.**

   The capability building belongs in the year-two re-architecture programme of Chapter Thirty-Two.

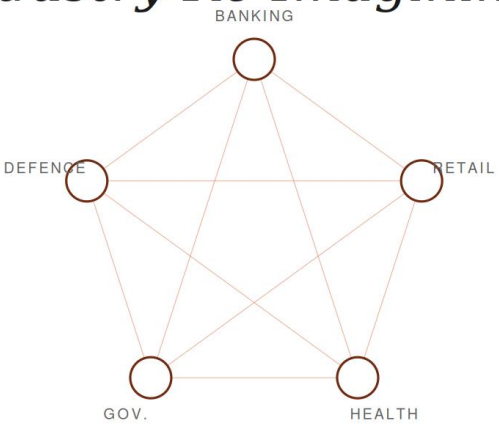
3. **Designate one director on the board with explicit responsibility for risk in the convergence decade.**

   The risk director's job is to ensure these questions are asked at every board meeting, even when the executive team would rather move past them.

PART EIGHT

VIII

Industry Re-Imaginings



Five sectors examined at the level a CEO in that sector needs to see. The convergence does not arrive uniformly. Each industry has its own clock, its own tipping points, and its own playbook.

THE WORLD THAT'S BEGINNING

Five sectors examined at the level a CEO in that sector needs to see. The convergence does not arrive uniformly. Each industry has its own clock, its own tipping points, and its own playbook.

CHAPTER THIRTY-FIVE

35

Banking and Financial Services Re-imagined

From the destination bank to the embedded one.

—

"Banking is necessary. Banks are not."

— Bill Gates, 1994, more right thirty years later than when he said it

◆ ◆ ◆

The chairman of one of India's largest private banks walked me through his board's strategic deck in early 2026 with what I found, retrospectively, to be remarkable candor. He turned to a slide that showed the bank's customer demographic by age cohort. The under-thirty cohort had declined by fourteen percent over the previous three years. The under-twenty-five cohort had declined by twenty-two percent. The bank's product had not deteriorated. The bank's interest rates had not become uncompetitive. The bank's branches had not closed. The bank had simply become, for a generation of Indians, *not where banking happens anymore*. Banking, for them, was happening inside three or four apps that did not have the word "bank" in their names — a payments app, a wealth-management app, a small-business app, an insurance app — each of which had absorbed a piece of what the bank used to do, while

the bank continued to operate as though the customer would still come to its branch for relationship management.

He asked me, with the directness of someone who had run out of comfortable answers, what I would do if I were running the bank. I told him, more or less, what this chapter contains.

◆ ◆ ◆

I. The destination is dying. The relationship is not.

The twentieth-century bank was a destination. The customer came to it — first physically to the branch, later digitally to the website or the app — to perform a banking task. The relationship was anchored in this destination model. Marketing brought the customer to the destination. Service handled the customer at the destination. Compliance audited the customer's behaviour at the destination. The destination was the bank's home turf, and the bank's competitive posture was to be the best destination in its category.

The convergence decade is dissolving the destination. Customers no longer go to the bank for banking. Banking comes to them, embedded inside the apps and interfaces and life moments where the financial decision is naturally being made. The customer making a purchase does not go to the bank for credit; the credit is offered, contextually, at the point of purchase, by whoever is fastest to recognise the moment. The customer planning a holiday does not go to the bank for foreign exchange; the foreign exchange is built into the booking. The customer running a small business does not go to the bank for working capital; the working capital is calibrated, automatically, against her real-time cash flow by whoever sees the cash flow most clearly.

The bank that continues to operate as a destination will, in five to seven years, find itself in the position of a railway station in 1995 — still operating, still patronised by some, structurally peripheral

to where the actual movement is happening. The bank that re-platforms as the embedded relationship — present invisibly inside every other interaction in the customer's financial life — will own the next decade.

Destination vs Embedded Banking

Two postures. Pick one. Cease pretending you are doing both.

Dimension	DESTINATION BANK	EMBEDDED BANK
Customer journey	Customer comes to the bank	Bank reaches the customer at the moment of decision
Primary surface	Branch, app, website	Embedded inside other apps, devices, places
Revenue model	Spread on deposits, fee on transactions	Take rate on contextual financial moments
Competitive moat	Trust, scale, regulation	API velocity, data, partner network
Customer relationship	Owned by the bank	Shared across the embedded ecosystem
Brand role	The destination's name	The trust signal inside other brands

FIGURE 35.1 · The destination dies; the relationship does not.

Figure 35.1

◆ ◆ ◆

II. Four moves the bank must make

The transition is not subtle. It requires four substantial moves, each of which most banks in 2026 have begun to discuss but few have begun to execute with the necessary commitment.

The first move is the agentic re-architecture of customer-facing operations. Most banking work that involves the customer — onboarding, KYC, credit assessment, customer service, dispute resolution, advisory — is amenable to substantial agentic automation. The banks that have done this well have absorbed eighty percent of the routine workload into agents and elevated their human relationship managers to deal with the moments that genuinely require human judgment and presence. The cost-to-serve has fallen

dramatically. The customer experience has, in many cases, improved.

The second move is the re-platforming of the bank as an API-first institution. The embedded-finance future requires the bank to be present, programmatically, inside other companies' products. The bank that cannot expose its capabilities cleanly through APIs cannot participate in the embedded economy. This requires deep technology investment, deep regulatory engagement, and a posture toward partnerships that most banks are culturally unaccustomed to. The banks that succeed are the ones whose CEOs understand that the bank's product is no longer principally consumed at the bank's own surface.

The third move is the radical re-imagination of the branch. Most retail banks in India still operate networks of hundreds or thousands of branches, sized for an era in which the branch was the primary channel. The branch is now a tertiary channel. The capital trapped in branch infrastructure is, in most banks, the single largest unproductive asset on the balance sheet. The honest move is not to gradually optimise the branch network but to fundamentally re-conceive what a branch is for. In my recommendations to clients, the branch of 2030 is a small number of high-trust advisory locations — fewer, larger, deliberately designed for the moments that genuinely require human presence (significant life-stage decisions, complex small-business advisory, high-trust wealth conversations). The other ninety percent of the network either closes or repurposes.

The fourth move is the re-engineering of the bank's data and trust architecture for the synthetic-content era. The bank is, in 2026, one of the institutions most exposed to synthetic-content risk. Voice deepfakes attempt to authorise transfers. Identity-cloning attacks target high-net-worth customers. AI-generated documents fool routine credit checks. The bank that has not, by 2027, hardened its

identity verification, its document authentication, and its customer-communication protocols against the synthetic-content threat will absorb losses that the bank that has hardened them will not.

◆ ◆ ◆

III. Where Indian banks have an asymmetric advantage

Indian banks have, in this transition, an asymmetric advantage that most of them have not yet capitalised on. The country's digital public infrastructure — UPI, Aadhaar, Account Aggregator, ONDC, the GST network, DigiLocker — provides the foundational layer of identity, payment, consent, and data exchange that the embedded-finance future requires. No other large economy has built this layer. The bank that integrates deeply with the public infrastructure has access to capabilities its global counterparts cannot easily acquire.

The advantage is not automatic. It requires the bank to treat the public infrastructure as a first-class platform on which to build, rather than as a regulatory inconvenience to comply with. The banks that have made this shift — and there are a small number — are now able to launch products in weeks that would have taken quarters to launch on legacy infrastructure. The compounding effect of this speed advantage is significant.

The Indian bank with the deliberate strategic posture to embed itself, via the public infrastructure, into every transaction in the customer's life — payments, credit, savings, insurance, investment, tax, identity — is, in my honest forecast, the bank most likely to define the category in 2032. The bank that defends the destination model is, in the same forecast, the bank most likely to be the regulated incumbent that the embedded competitors steadily disintermediate.

◆ ◆ ◆

IV. The wealth and small-business frontiers

Two specific frontiers within Indian financial services deserve particular attention, because they are large, fast-moving, and represent the largest opportunity for new entrants and the largest existential exposure for incumbents.

The wealth frontier. India's affluent middle class is producing the largest cohort of new financial-service customers in human history. The cohort is digitally native, distrustful of legacy financial brands, and willing to pay for genuinely good advisory experiences delivered through digital interfaces. The conventional Indian bank's wealth offering — built around relationship managers, product pushes, and quarterly portfolio reviews — is precisely the offering this cohort does not want. The opportunity is for an AI-augmented wealth experience that combines the warmth of a personal advisor with the scale and consistency that only digital delivery allows. Several new entrants are in the early stages of building this. The incumbents have, in most cases, not yet recognised the threat.

The small-business frontier. India's small and medium businesses — the country's tens of millions of MSMEs — are the largest under-served financial services market in the world. The conventional banking apparatus has struggled to lend to them, serve them, and integrate them, because the unit economics did not work and the data did not flow. The combination of agentic AI, the public-infrastructure data layer, and embedded finance changes this. A small-business owner in tier-three urban India, in 2026, can already access in-context financial services that did not exist three years ago. By 2030, the financial life of the Indian MSME will be substantially digital, deeply embedded in the operating tools she uses, and far better-served than it has ever been. The bank that captures this segment will own a generation of growth. The bank that does not will watch it captured by others.

◆ ◆ ◆

V. The trust premium

One consideration runs through every move I have described. The bank's central asset, even as the destination dissolves, is trust. The customer trusts the bank with her money. She trusts it with her data. She trusts it with her financial decisions in moments of vulnerability. This trust, accumulated over decades in many cases, is the bank's most durable competitive position.

The convergence decade puts this trust under stress in two directions. The synthetic-content threats erode it from outside. The agentic transformations, if mishandled, erode it from inside — when the customer feels she is no longer talking to her bank but to an algorithm with the bank's name on it. The bank's strategic task is to deepen and visibly defend the trust currency through the transition, which means deliberate transparency about how AI is used, deliberate human presence at the moments where trust is most acutely tested, and deliberate accountability for the agentic systems the bank deploys.

The banks that do this well will, in 2030, hold the high ground that no fintech entrant can easily contest. The banks that do not will find that their accumulated trust evaporates faster than their incumbent advantages can be defended.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit, this quarter, where your bank's customers are actually conducting their financial lives — and how much of that activity touches your brand.**

   The audit is uncomfortable. It is also the foundation for every subsequent decision.

2. **Commit to API-first re-platforming as a strategic priority, not a technology project.**

   The CEO who delegates this to the CIO is delegating the bank's future.

3. **Choose, this year, between the destination posture and the embedded posture. Cease pretending you are doing both.**

   Both is the comfortable position. Both is also the position that, in most banks, produces neither outcome cleanly.

CHAPTER THIRTY-SIX

36

Retail and Consumer Re-imagined

The store recedes; the relationship continues.

—

"Retail is not about the store. Retail is about getting the right thing to the right person at the right moment."

— *A retail executive, paraphrased, who would prefer not to be named*

◆ ◆ ◆

I sat with the founder of a mid-sized Indian apparel retailer, in a coffee shop in Bangalore, as she walked me through the metrics that had brought her to the conversation. Same-store sales, declining for six consecutive quarters. Online conversion, plateaued. Customer acquisition cost, rising at twelve percent per year. The product was not the issue, she said. The product had won design awards. The store experience was not the issue, in any obvious sense. The customers who came to the stores liked them. The issue was that fewer and fewer of the right customers were coming to either the stores or the website. Her brand had, somehow, become invisible in the precise moments when the customer was deciding what to wear next.

She asked me, with the same directness as the bank chairman in the previous chapter, what was actually happening to retail. The answer is the subject of this chapter, and it generalises beyond apparel.

◆ ◆ ◆

I. Discovery is migrating away from the retailer

The twentieth-century retail journey, simplified, had three stages: awareness, consideration, purchase. The retailer's marketing produced awareness. The retailer's catalogue or store produced consideration. The retailer's checkout closed the purchase. The retailer controlled, or at least heavily influenced, all three stages. Marketing budgets, store design, channel strategy, and merchandising were all the retailer's tools to operate within stages it could meaningfully shape.

The convergence decade is migrating discovery away from the retailer entirely. The customer's awareness of products in 2026 is increasingly produced not by retailer-controlled channels but by AI companions, social-platform recommendations, peer-network signals, creator endorsements, and contextual triggers in the customer's daily life. Consideration is happening inside aggregator interfaces — sometimes vertical aggregators (a meta-search across all apparel options matching her preferences), sometimes general AI assistants (the companion that suggests, when she mentions a friend's wedding, three outfits matching her style and budget). By the time the customer reaches the retailer's surface, the discovery and consideration stages have largely concluded. The retailer's role has been compressed to fulfilment.

The implication is that the retailer's traditional toolkit — brand marketing, store experience, merchandising — has been rendered, in significant part, structurally less powerful. The retailer is no longer in the room when the decision is being made. She is downstream of where the decision is being made, and increasingly, the customer encounters her brand only when the decision has, more or less, been made.

This is not a future scenario. It is the operating environment of mid-market retail in 2026. The founder I sat with in Bangalore was experiencing it without having yet named it.

◆ ◆ ◆

II. Three responses retailers can take

There are three credible postures the retailer can take in response. Each is a substantial commitment. None is easy.

The first posture is to become the AI surface in the customer's category. The retailer builds, owns, and operates the AI experience through which her customer makes category decisions. She is the companion the customer turns to when she needs an outfit, a meal, a gift, a home improvement. This requires building a vertical AI of genuine quality — better than the general-purpose AI in the category — and maintaining deep enough customer relationships that the customer continues to use it. The strategy is expensive and demanding, but it captures the discovery layer that other retailers have ceded.

The second posture is to embed the brand inside the AI surfaces the customer uses anyway. The retailer does not build her own AI. She invests, deeply, in being the brand the existing AI assistants surface, recommend, and route to. This requires very high product quality, very strong brand signals, very clean data feeds to the AI surfaces, and significant attention to the new SEO equivalent — the set of practices that determine which brands are surfaced by AI assistants and which are not. The strategy is less capital-intensive than the first but requires constant attention to a moving target.

The third posture is to differentiate radically on what AI cannot replicate — physical presence, human warmth, community, craft, and the irreducibly local. The retailer leans into the things the AI assistant cannot deliver: the in-person fitting, the staff who know the customer by name, the third place where the customer gathers with people who share her aesthetic, the curated community around the brand. The store becomes less a transaction venue and more a community venue. The economics shift from transaction-per-square-foot to relationship-per-customer. This strategy is, in many ways, the most distinctive and the most defensible, but it requires

the retailer to genuinely commit to the human-presence dimension rather than treating it as decoration.

Most retailers, in practice, will combine elements of all three. The discipline is to commit to a primary posture and execute it deliberately, rather than to dabble across all three and execute none well.

◆ ◆ ◆

III. The new unit economics

The unit economics of retail in the convergence decade are different from the unit economics that built the great retail brands of the twentieth century. Several shifts deserve naming.

Customer acquisition cost is rising, but customer lifetime value is rising faster. The customer who can be acquired through AI-mediated channels and held in a continuous relationship has a lifetime value substantially higher than the customer of the past. The retailer who optimises for first-purchase margin will lose to the retailer who optimises for lifetime relationship value, even when the first-purchase margin is lower.

Inventory is becoming more dynamically managed. Predictive AI, applied to demand sensing, is collapsing the inventory mismatches that used to characterise apparel, electronics, and home goods. The retailer that has integrated AI-driven inventory across her supply chain operates with a fraction of the working capital that legacy retailers carry, with fewer markdowns, with less waste.

The store's role is becoming experiential rather than transactional. The store that is purely a transaction venue is losing share to digital. The store that is an experience venue, and where the transaction is incidental, is holding share or growing. This means the store's design, its staffing, its programming, and its measurement must shift accordingly. Same-store sales is no longer the only relevant

metric; same-store engagement, time-spent, conversion-to-membership, and recurring-customer rate matter more.

Returns are becoming a first-class strategic concern. In categories where AI-mediated discovery is producing higher-volume, lower-conviction purchases, the return rate is rising. The retailer that has not engineered return logistics and return economics for this rising rate is bleeding margin invisibly.

◆ ◆ ◆

IV. The vernacular and aspirational frontiers

Two specific opportunities in Indian retail deserve attention because they are large, fast-moving, and currently under-served.

The vernacular frontier. Most retail in India has been built for the English-speaking, urban, upper-middle-class consumer. The vernacular consumer — Hindi, Tamil, Telugu, Bengali, Marathi, Punjabi — has been served by older or less ambitious brands, by unorganised retail, or by hyperlocal players. The combination of AI-driven personalisation in vernacular languages, the digital public infrastructure for payment and delivery, and the mobile-first habits of this consumer has, in 2026, made the vernacular retail frontier accessible at scale for the first time. The retailer who builds genuinely for this consumer — in her language, with her aesthetic preferences, with her economics, with her trust frameworks — will own a market that is larger than the English-speaking urban one and currently barely contested.

The aspirational-luxury frontier. India's affluent middle class is, simultaneously, building the world's largest aspirational-luxury consumer base. The luxury houses of Europe, Japan, and the United States have begun to recognise this and to invest in India. They are, in most cases, applying their global playbooks. The retailer who understands what aspirational luxury actually means in the Indian cultural register — which is quite different from what

it means in the European or American register — will out-compete the global houses at their own category. The civilisational substrate I argued for in Chapter Twelve applies sharply here. The luxury category is, in essence, a meaning-layer category, and the meanings that resonate in India are not the meanings the European houses were built around.

◆ ◆ ◆

V. The retailer's harder question

The harder question for any retailer in 2026 is whether her business is fundamentally a brand business or a logistics business. The two have different economics, different competitive moats, and different end-states.

The brand business builds a meaning-layer relationship with customers and uses logistics as a cost line. The brand wins by being what the customer wants to be associated with. The logistics business builds an operational moat around fulfilment, inventory, and unit cost, and uses brand as a marketing input. The logistics business wins by being able to deliver more reliably, more cheaply, more flexibly than competitors.

In the past, most retailers operated as a hybrid of both. The convergence decade, in my view, will force a choice. The brand businesses that compete on logistics will lose to specialist logistics players. The logistics businesses that compete on brand will lose to specialist brand players. The retailer who chooses, deliberately, which business she is in — and who invests accordingly — will, by 2030, have a defensible position. The retailer who continues to be both will, by 2030, be neither.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Decide whether you are principally a brand business or a logistics business — and invest accordingly.**

   The decision is uncomfortable. It is also the foundation of every subsequent retail strategy.

2. **Choose your AI posture — surface, embedded, or human-differentiated — and execute it.**

   The dabbling posture is the most expensive and the least successful.

3. **Re-baseline your store metrics for the era in which the store is principally an experience venue, not a transaction venue.**

   The metrics you use determine the strategy you execute. Update the metrics first.

CHAPTER THIRTY-SEVEN

37

Healthcare and Pharma Re-imagined

The patient becomes the platform.

—

"The doctor of the future will give no medicine, but will instruct his patient in the care of the human frame, in diet, and in the cause and prevention of disease."

— *Thomas Edison, who anticipated, perhaps, more than he knew*

◆ ◆ ◆

Chapter Sixteen treated the patient experience side of healthcare. This chapter takes the operator and pharmaceutical view — the institutional restructuring, the pharma economics, and the policy interfaces that the convergence decade is reshaping. The two chapters together complete the healthcare picture as I see it in 2026.

The conversation that prompted this chapter was with the chief executive of a pharmaceutical company that had spent the previous decade building a respected portfolio of generic drugs, and that was now confronting a strategic question its board had not been built to answer. The question was whether the company should remain a generic-drug company in the era when AI-augmented biology was making personalised therapeutics commercially feasible at scale, or whether it should re-platform

fundamentally as a precision-therapeutics company. The two paths required different capabilities, different capital structures, different regulatory engagements, and different timelines to commercial return. The CEO told me, with the candor I have come to expect from the leaders who eventually take the harder paths, that he honestly did not know which choice would be right. The point of our conversation was to think it through carefully.

◆ ◆ ◆

I. The patient as platform

The central architectural shift in healthcare, beyond the continuous-care shift I treated in Chapter Sixteen, is the emergence of the patient herself as the platform around which the system reorganises. In the previous era, the platforms of healthcare were institutions — the hospital, the insurance company, the pharma company, the regulator. The patient was an object on which these institutions operated. In the emerging era, the patient — armed with continuous biometric data, genomic information, AI-augmented self-awareness, and digital agency — is becoming the platform on which the institutions operate.

The shift has consequences that are not yet, in 2026, fully visible. The patient is increasingly able to choose which provider to consult based on real outcome data rather than reputation. She is increasingly able to authorise specific data flows for specific purposes, with cryptographic precision, rather than yielding undifferentiated consent at intake. She is increasingly able to participate in clinical research as an individual, contributing her data on terms she has set, in exchange for value she finds adequate. She is increasingly able to monitor her own treatment efficacy and to switch providers when efficacy is not delivered.

The institutions that will succeed in this configuration are the ones that have re-organised around the patient as the platform — that have built their operating systems on the assumption that the

patient owns her data, her relationships, and her decisions, and that the institution's role is to deliver value to her on terms she sets. The institutions that continue to operate as though the patient is the object of their operations will, in five to seven years, find themselves disintermediated by competitors that recognised the architectural shift earlier.

◆ ◆ ◆

II. The pharma re-architecture

The pharmaceutical industry is, in 2026, in the early innings of an architectural shift that is, in my view, of the same magnitude as the shift to biologics in the 1990s. The shift has three drivers and three commercial implications.

The drivers are AI-accelerated drug discovery, the maturation of precision and personalised therapeutics, and the increasing role of real-world evidence — collected from continuously monitored patients — in regulatory and commercial decision-making. The first compresses the front end of drug discovery from years to months for many categories. The second moves the centre of gravity from one-size-fits-all blockbusters to therapeutics calibrated to genotype, microbiome, and disease subtype. The third changes how drugs are approved, how they are reimbursed, and how their performance is measured after launch.

The commercial implications are several. *The blockbuster model is under pressure*, because the markets that supported single-molecule billion-dollar drugs are fragmenting into smaller, more tightly defined patient populations. *The traditional sales model is under pressure*, because the prescriber's decisions are increasingly informed by real-world evidence rather than by sales-force-mediated information. *The R&D economics are shifting*, because AI-augmented discovery and trial design can substantially reduce the cost per successful drug, but only for the firms that have built the capability. *The CMC and manufacturing economics are shifting*,

because personalised therapeutics require manufacturing at smaller scales with greater complexity, which advantages firms that have built the manufacturing capability and disadvantages those that have not.

For Indian pharma specifically, the strategic question is whether to remain principally a generic and biosimilars producer — competing on cost, scale, and execution — or to invest in becoming a precision-therapeutics producer for the Indian and emerging-market populations that the global majors have under-served. The two paths are not mutually exclusive, but the capital intensity of the second is significant, and the firms that try to do both without committing capital adequately to the second will, in my forecast, find themselves competitive in neither.

◆ ◆ ◆

III. The hospital as conductor, not factory

The twentieth-century large hospital was, organisationally, a factory. Patients came in, got processed through diagnostic and therapeutic workflows, and went out. The hospital's competitive position rested on the capital intensity of its equipment, the breadth of its specialties, the throughput of its operating theatres, and the brand of its name on the building.

The hospital that wins in the convergence decade is, organisationally, a conductor. It coordinates a network of diagnostic, therapeutic, and continuous-care capabilities, some of which it owns and many of which it does not, on behalf of a patient whose relationship it owns over years and decades. The hospital's competitive position rests on the depth of its patient relationships, the quality of its orchestration, the integration of its data, and the trust in which it is held by patients and referrers.

The transition is significant. The capital structure changes. The metrics change. The leadership skills change. The relationship with

insurers changes. The relationship with employers — who are increasingly buying health for their workforces directly rather than through intermediated insurance — changes. The hospitals that have begun this transition in 2026 will, by 2030, be the regional category leaders. The hospitals that continue to operate as factories will, by 2030, be commodity providers absorbed into larger orchestrators or in financial distress.

◆ ◆ ◆

IV. The insurance and payment re-think

The payment and insurance architectures of healthcare are, in most of the world, mid-twentieth-century institutions. They were built around episodic care delivered through professional silos, with payment tied to service volumes rather than outcomes. The architecture is increasingly miscalibrated.

Three shifts are now underway. *Payment is shifting from fee-for-service to fee-for-outcome*, slowly in some markets and rapidly in others, as the data infrastructure that allows outcome measurement becomes available. *Insurance is shifting from claims processing to continuous-health partnerships*, as the technologies of continuous care make it possible for the insurer to be a daily presence in the customer's health rather than an occasional adjudicator of bills. *Employer-led health buying is rising*, as employers — particularly in the US but increasingly elsewhere — recognise that their workforces' health is a strategic asset they can manage directly more effectively than through intermediated insurance.

For the Indian context, the relevant questions are different. India's insurance penetration is low; out-of-pocket spending is high; the public health system is variable; and the digital health stack — Ayushman Bharat Digital Mission, ABDM-linked records, the Health ID — is at an earlier stage than the financial public infrastructure. The country's healthcare payment architecture will, in this decade, be substantially built rather than substantially

reformed. The shape it takes is being decided, in part, in 2026 and 2027. Indian healthcare leaders, public and private, have an unusual opportunity to influence the architecture toward outcome-based, continuous, AI-augmented care — and a corresponding obligation to participate.

◆ ◆ ◆

V. The trust and ethics frontier

Healthcare is, of all the categories I treat in this book, the one in which the trust and ethics frontiers are most acutely tested. The patient is vulnerable. Her data is intimate. The decisions are consequential. The technologies are novel. The temptation to optimise for institutional efficiency at the expense of patient welfare is constant.

Three commitments must, in my view, anchor any responsible healthcare institution's strategy through this decade. *Transparency about AI*, because the patient deserves to know when an AI is participating in her care, what it is doing, and what its limitations are. *Patient ownership of data*, because the data is hers, and the institution's right to use it derives from her informed consent for specific purposes. *Outcome accountability*, because the institution that profits from a patient's care is morally and commercially obliged to be accountable for the outcome that care produces.

These commitments will, in some cases, slow the deployment of efficiency-optimising technology. They will, in every case, build the trust currency that is the durable competitive moat of healthcare in this era. The institutions that take them seriously will own their patients' lifetime relationships. The institutions that do not will find that, when the inevitable failure occurs, they have built none of the trust capital required to absorb it.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Whether you operate a hospital, a pharma firm, or an insurer, decide whether you are organisationally a factory or a conductor — and re-platform accordingly.**

   The factory model is the one most institutions are currently optimised for. The conductor model is the one this decade rewards.

2. **Build the data and trust architecture before you build the AI products.**

   The institutions that lead with AI products before they lead with trust will, on the first failure, lose both.

3. **If you are an Indian healthcare leader, participate actively in the policy and infrastructure decisions of the next twenty-four months.**

   The architecture is being decided. Participation now is more valuable than reaction later.

CHAPTER THIRTY-EIGHT

38

Government and Public Service Re-imagined

The state as platform; the citizen as participant.

—

"The state is the great fictitious entity by which everyone seeks to live at the expense of everyone else."

— *Frédéric Bastiat, who would, if he were alive, find India's digital public infrastructure an interesting partial counterargument*

◆ ◆ ◆

I have, over the past five years, spent substantial time with leaders of state governments, central ministries, public-sector utilities, and municipal corporations, advising on the deployment of AI and immersive-technology capabilities in public services. The conversations have been instructive in several directions, not least about how the convergence decade is reaching the public sector at a different pace and through different pressures than it is reaching private enterprise.

Public-service institutions in India, in my observation, are at a moment of asymmetric opportunity. The country has built, over the past fifteen years, a digital public infrastructure that is the most advanced of any large economy. The infrastructure is genuinely transformative. It has not yet, however, been fully matched by the

institutional re-platforming that would let the public sector deliver, on top of that infrastructure, the experience-quality citizens deserve. The gap between infrastructure and experience is the territory of this chapter.

◆ ◆ ◆

I. The infrastructure is built. The experience is not.

The Indian digital public infrastructure stack — Aadhaar for identity, UPI for payments, Account Aggregator for data flow, ONDC for commerce, DigiLocker for documents, the digital-health stack, the agriculture stack, the education stack — is, by 2026, approximately complete in its foundational layers. Citizens can be uniquely identified, can transact instantly, can authorise data flows precisely, can hold credentials digitally. The substrate is in place.

The experience layer is, in most public services, not in place. The citizen who has, on her phone, the most advanced payment infrastructure in the world still queues for hours at the regional transport office to renew her driving licence. The farmer who has, in his pocket, the digital identity that any system can recognise still cannot get his crop insurance settled without three site visits and four physical signatures. The patient whose health record sits in the digital health locker still presents physical reports at the hospital because the hospital's intake system has not integrated with the locker.

The gap is not technological. The substrate is technologically sufficient. The gap is institutional. It is the gap between a public agency designed for the world that is ending and a public agency redesigned for the world that is beginning. Closing the gap is the work of the next decade for Indian public services. The leaders who close it will deliver, in my view, the largest single improvement in citizen experience that the country has produced in half a century.

◆ ◆ ◆

II. Four moves the public sector must make

The moves are structurally similar to the moves I described for private enterprise, but the levers and the constraints are different.

The first move is the agentic re-architecture of citizen-facing operations. Most public services involve substantial work on the citizen's behalf — application processing, eligibility determination, scheme integration, status updates, dispute resolution, advisory. Each of these is amenable to substantial agentic automation. The agencies that have begun this work — a handful of state governments, a small number of central ministries — are seeing dramatic improvements in throughput, cost, and citizen satisfaction. The agencies that have not are still measuring their performance in lakhs-of-files-cleared-per-week.

The second move is the unification of citizen experience across siloed agencies. From the citizen's perspective, the public sector is one entity. From the public sector's internal perspective, it is hundreds of agencies operating in their own silos with their own data, their own systems, and their own fiefdoms. The technology now exists to abstract over these silos at the citizen-experience layer. The citizen interacts with one consistent interface — call it MyGov in its mature form, or whatever the equivalent is at the state level — and the interface routes her requests appropriately to the responsible agency. The technology is not the constraint. The institutional politics of agency boundaries is. The leaders who can navigate this politics will deliver experience improvements at a pace that is, in some cases, transformative.

The third move is the modernisation of the public-sector workforce for the agentic era. The Indian state employs tens of millions of people. The work many of them do is, in 2026, the kind of structured, procedural work that AI absorbs first. The challenge for public-sector leadership is to redesign the workforce — not to reduce its size, which is politically and socially difficult, but to elevate its

work, retrain it for higher-value functions, and absorb the agentic productivity gain as expanded service rather than headcount reduction. The Indian state is, in this respect, in a similar position to the responsible private employer of Chapter Thirty, but with stronger constraints and weaker tools. The leaders who navigate this carefully will, by 2030, have transformed the public-sector workforce into a more capable, more motivated, more respected institution than it currently is. The leaders who do not will preside over a public sector that is increasingly miscalibrated against the citizen's expectations.

The fourth move is the re-engineering of the public sector's data and security architecture for the threat profile of the late 2020s. The cybersecurity, post-quantum, and synthetic-content concerns I treated in Chapters Seven and Thirty-Four apply with particular urgency to the public sector, which is the custodian of the country's most sensitive data and the operator of its most critical infrastructure. The hardening required is significant. The institutional capacity to do it is, in many parts of the public sector, currently insufficient. The investment required is substantial but bounded. The cost of not making it is the systemic risk of a major cybersecurity incident at a moment when the public's confidence in digital systems is foundational to the country's economic functioning.

◆ ◆ ◆

III. The municipal layer

A specific frontier deserves particular attention. India is urbanising. By 2030, the country will have at least seven cities with populations greater than ten million, and approximately seventy-five cities with populations greater than one million. The urban experience that the country's growing middle class encounters daily — traffic, public services, infrastructure quality, environmental conditions — is, in 2026, well below the quality that

comparable urban populations in other middle-income countries encounter. The gap between citizen expectation and municipal delivery is widening.

The convergence technologies — sensor networks, AI-driven traffic management, predictive maintenance, citizen-feedback platforms, AR-based public information, digital municipal services — are now mature enough to substantially close this gap. A handful of Indian cities have begun to deploy elements of this stack. The cities that take it seriously over the next five years will, by 2030, be substantially more functional than their peers, will attract more talent and capital, and will become the templates that the rest of the country adopts. The cities that do not will face widening citizen frustration that is, in my view, a serious source of political risk for the leaders who failed to act.

The opportunity for state and national leadership is to identify, fund, and showcase a small number of city-scale transformations that demonstrate what is possible. The cost of these transformations is meaningful but not prohibitive. The political and economic return on a few well-executed municipal showcases is substantial. The municipal layer is, in my view, the highest-leverage public-sector investment available in the decade ahead.

◆ ◆ ◆

IV. The digital sovereignty question

A strategic question, distinct from the operating questions above, will increasingly press on Indian public-sector leadership through this decade. The question is the country's posture on digital sovereignty — the degree to which India's data, AI capabilities, and digital infrastructure are owned and operated within India versus dependent on foreign providers and platforms.

The question is not new. India has, over the past decade, made deliberate choices that distinguish its digital posture from that of

most other large economies. The public-infrastructure stack is sovereign, open, and Indian-built. The data localisation requirements have been progressively strengthened. The country has resisted the platform-monopoly outcomes that characterised the US and Chinese digital orders in the previous decade.

The questions for the decade ahead are sharper. How much of the country's foundation-model capability should be Indian? How much of the country's compute infrastructure should be sited in India? On what terms should foreign AI providers be permitted to operate in India? How should the country's strategic autonomy be balanced against the legitimate interests of cross-border collaboration and foreign investment? These are not technical questions. They are strategic questions with substantial economic and geopolitical consequences. The answers India settles on in 2026 to 2028 will shape the country's digital position for the rest of the century.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you lead a public-sector agency, identify one citizen-experience pain point this quarter and commit to closing it within the year using the convergence technologies.**

   The credibility of the broader transformation is built on visible early wins.

2. **If you advise on policy, contribute to the digital sovereignty conversation actively rather than waiting for it to be resolved.**

   The architecture is being decided. Participation is more valuable now than later.

3. **If you operate in the private sector with public-sector customers or stakeholders, build deliberate institutional relationships with the public-sector leaders driving the transformation.**

   The transformation is a public-private collaboration. The firms that participate as serious partners will operate inside the new architecture; the firms that approach the public sector as a transactional customer will operate outside it.

CHAPTER THIRTY-NINE

39

Defence and Strategic Sectors Re-imagined

Where the stakes are highest and the pace of change is fastest.

—

"In war, the moral is to the physical as three to one."

— Napoleon Bonaparte

◆ ◆ ◆

The chapter that follows is the most reluctant in this book. I write it because the convergence decade is reshaping defence and strategic sectors more profoundly than it is reshaping any commercial sector, and because the leaders of these institutions — civilian and military — are operating with less commercial-strategy literature than the leaders of commercial enterprises. I write it carefully because the questions involved are consequential at the level of human life and national security, and the temptation to offer glib commentary in this domain is one I have tried to resist.

I will limit my treatment to four observations that I believe are accurate, useful, and not exclusive to any specific national interest.

◆ ◆ ◆

I. The convergence is reshaping strategic capability

The technologies I have treated throughout this book — AI, agentic systems, spatial computing, quantum, biology, advanced materials, energy — are, simultaneously, reshaping the strategic capabilities that defence institutions must understand, possess, and counter. AI-augmented intelligence and decision support, autonomous systems across all domains, immersive training environments, quantum communications and cryptography, cyber capabilities of unprecedented sophistication, and information-environment operations of greater scale and subtlety than any previous era — each is now an active dimension of strategic competition.

The implication for defence institutions is that the technology decisions of this decade will shape strategic capability for half a century afterward. The decisions are being made now, in capital allocations, doctrinal evolutions, and organisational redesigns. The institutions that make these decisions well will, in 2050, have positioned their nations advantageously. The institutions that defer them will be reacting to circumstances chosen by others.

For India specifically, the strategic environment of the next two decades will be defined, in my honest reading, by the country's ability to combine indigenous defence-technology development, strategic partnerships with aligned countries, and a doctrinal posture that fits its civilisational and geographic position. The country's defence institutions have begun this work. They have not, in my view, yet fully matched the urgency the environment requires.

◆ ◆ ◆

II. The civilian-military boundary is dissolving

A structural feature of the convergence decade is that the technologies most consequential for defence are also, almost without exception, the technologies most consequential for the

civilian economy. AI, biology, energy, communications, advanced manufacturing — each is dual-use in a way that the strategic technologies of the previous era were not.

The implication is that the civilian-military boundary, which characterised much of the twentieth-century defence-industrial relationship, is dissolving. The cutting edge of strategic capability is, in many areas, being developed by civilian firms whose primary markets are commercial. Defence institutions that continue to operate as though their suppliers are principally defence contractors are increasingly mismatched against an environment in which their suppliers are principally commercial firms with defence applications.

For defence leadership, this requires a different mode of engagement with the technology base — a mode that treats civilian firms as primary partners rather than as occasional suppliers, that operates at the speed of commercial innovation rather than at the speed of traditional procurement, and that builds institutional pathways for talent to move between civilian and defence work without prohibitive friction. The countries that have done this well — the United States in particular, Israel notably — have produced disproportionate strategic capability. The countries that have not are finding their defence capability lagging the global frontier despite substantial spending.

◆ ◆ ◆

III. Information and influence are first-order strategic domains

The information environment has, in this decade, become a strategic domain of the same first-order character as land, sea, air, space, and cyber. The synthetic-content technologies I described in Chapter Twenty-Four are being deployed, in 2026, by sophisticated state and non-state actors to shape narratives, influence elections, divide societies, and erode the trust on which democratic governance depends. The asymmetry between the cost of producing influence operations and the cost of detecting and

countering them favours the producer. The defensive infrastructure required to counter these operations is, in most countries, underdeveloped.

For defence institutions, this is uncomfortable territory because the boundaries between defence, intelligence, civilian governance, and free expression are not sharply drawn and should not be. The institutional response cannot be primarily a military response without crossing important political lines. It must be a whole-of-society response that includes media institutions, educational institutions, technology firms, and civil society — coordinated, resourced, and led with deliberate care.

The countries that develop this whole-of-society capability over the next five years will be substantially more resilient to the information operations of the late 2020s. The countries that do not will be substantially more vulnerable. The choice is being made, by action or inaction, in 2026 and 2027.

◆ ◆ ◆

IV. The ethical frontier requires explicit attention

The technologies of this decade are introducing strategic capabilities — autonomous weapons, AI-augmented decision support at strategic levels, genetic technologies with potential weaponisation, quantum capabilities with strategic implications — that the existing frameworks of international law and military ethics are not fully prepared for. The choices being made, in 2026 and 2027, by major powers about how to deploy these capabilities will set norms that will be extraordinarily difficult to revise later.

I will not, in this book, take strong positions on specific defence-technology choices, because the questions are properly the responsibility of legitimate national institutions rather than of commentators. I will, however, register one observation that I believe applies broadly. The ethical frontiers of these technologies cannot be deferred to crisis moments. They must be addressed

deliberately, transparently to the extent national security permits, and with explicit institutional mechanisms for review and accountability. The countries that build these mechanisms will, when crises arise, have the institutional capacity to navigate them with their values intact. The countries that defer the work will, when crises arise, find themselves making consequential ethical choices under pressure with inadequate institutional support.

This is, in some sense, the same observation I have made throughout the book about the trust currency, applied to the strategic domain. The institutions that take ethics seriously in advance will, when tested, hold. The institutions that defer the work will, when tested, fail.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in defence or a strategic sector, accelerate the integration of civilian innovation into your institutional capability.**

   The traditional procurement model is increasingly mismatched against the speed of the technology environment.

2. **If you operate in the civilian technology sector, engage carefully and constructively with defence and strategic-sector counterparts.**

   The boundary between civilian and strategic is dissolving. The terms on which the dissolution occurs depend on the engagement of leaders on both sides.

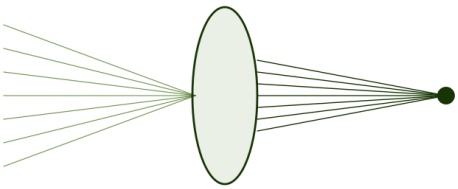
3. **If you advise on national policy, contribute actively to the ethical-framework conversations of this decade.**

   The frameworks are being set. Quiet contribution from thoughtful voices is more valuable now than louder contribution later.

PART NINE

IX

The Investor's Lens



How capital allocation must change, and what the convergence asks of those who direct it.

THE WORLD THAT'S BEGINNING

How capital allocation must change, and
what the convergence asks of those who direct it.

CHAPTER FORTY

40

How Capital Allocation Changes

The patterns of value creation, and value capture, in the next decade.

—

"Compound interest is the eighth wonder of the world. He who understands it, earns it. He who doesn't, pays it."

— quoted earlier in this book; worth repeating in the chapter that most depends on it

◆ ◆ ◆

I have spent significant time, over the past three years, in conversations with investors — venture, private equity, public-market, sovereign — who are attempting to translate the broad outlines of the convergence decade into specific capital-allocation decisions. The conversations have a recurring shape. The investor sees the technology shift clearly. The investor sees that the valuations of incumbent businesses do not yet, in most categories, reflect the disruption to come. The investor is uncertain about which businesses will be the winners, on what timeline, with what margins, in what configurations. The result is, in many cases, paralysis or its opposite — frantic activity that is more about being seen to act than about acting wisely.

This chapter is not investment advice in the regulated sense. It is a leader's view of how the patterns of capital allocation will shift in

the convergence decade, written for the readers of this book who deploy capital — directly or indirectly through the firms they run.

◆ ◆ ◆

I. Value migrates up the stack

The primary pattern, which underlies most of the others, is that value migrates from the layers at which it concentrated in previous decades to layers above them. The firms that built fortunes in the past forty years built them at the application layer of the personal-computing revolution, at the platform layer of the internet revolution, and at the data layer of the mobile-and-cloud revolution. The fortunes of the next twenty years will be built, in my forecast, at three layers that are still being defined.

The first is *the AI substrate layer* — foundation models, AI hardware, AI infrastructure, the agents that will mediate most knowledge work. The capital intensity of this layer is significant; the winners will be few and large. Investors with the scale and the appetite to participate at this layer will participate in returns that are, on a probability-weighted basis, comparable to the platform returns of the previous era.

The second is *the experience-orchestration layer* — the firms that combine AI, sensors, design, trust, and category-specific expertise to deliver the integrated experiences I have described throughout this book. This is a more diversified layer with many more potential winners than the substrate layer. The capital intensity per firm is moderate. The returns, for the firms that build durable customer relationships and that compound on the trust currency, are substantial. This is the layer where the most interesting venture capital will, in my view, be deployed in the next decade.

The third is *the trust and verification layer* — the institutions that anchor authenticity, identity, content provenance, and the verification infrastructure that the synthetic-content era requires.

This is a smaller layer in terms of revenue but disproportionate in terms of strategic importance. The firms that occupy it well will have outsized influence on every other layer.

The investor's pattern recognition for the next decade is, in essence, to identify these three layers in their specific category, to evaluate which layer the investment thesis depends on, and to size the allocation according to the layer's economic logic. The investor who continues to apply previous-era pattern recognition — looking for application-layer winners, for example — will, in my forecast, struggle to match the returns the new patterns offer.

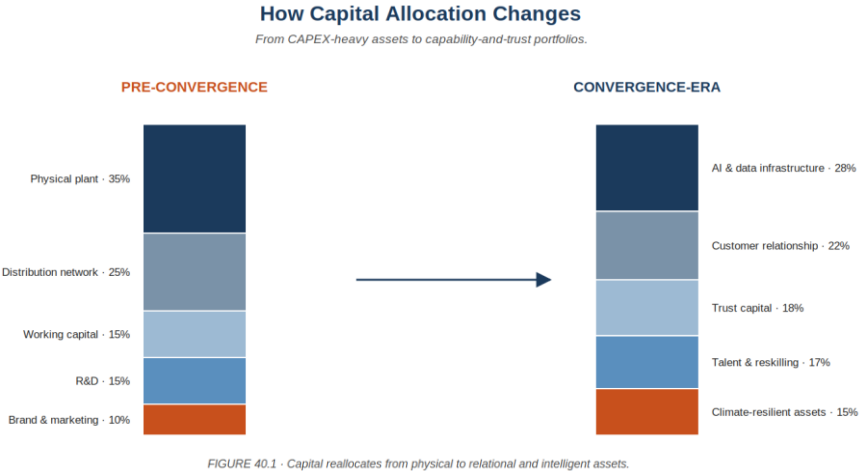


Figure 40.1

◆ ◆ ◆

II. Compounding versus accumulation

A second pattern, distinct from the first, is the difference between businesses that compound and businesses that accumulate. The distinction matters because the convergence decade rewards compounding far more than it rewards accumulation, and the two often look similar at the level of revenue growth.

A *compounding* business is one whose competitive position strengthens with every additional customer, every additional period of operation, every additional unit of data, every additional brand interaction. The customer relationships deepen. The data advantages widen. The trust capital accumulates. The cost of customer acquisition declines as the brand becomes recognised. The competitor's task to take a customer becomes harder year by year.

An *accumulating* business is one whose revenue grows without the underlying competitive position strengthening proportionally. The customer relationships are transactional. The data is held but not deepened. The brand is recognised but not preferred. Customer acquisition cost remains constant or rises. The competitor's task to take a customer is, in five years, no harder than it was in year one.

In the previous era, the distinction between these two patterns was meaningful but not decisive. Many accumulating businesses produced good investor returns over a decade. In the convergence decade, my honest forecast is that the distinction will be decisive. The compounding businesses will produce extraordinary returns. The accumulating businesses will increasingly be exposed to competitors that have built compounding mechanics, and their margins will, over a decade, be substantially eroded.

The investor's task is to distinguish the two with discipline. The diagnostic is not revenue growth but the trajectory of unit economics over time. If the customer acquisition cost is falling, the customer lifetime value is rising, the brand-preference signals are strengthening, and the competitive moats are deepening, the business is compounding. If the unit economics are flat or eroding, the business is accumulating. The market often confuses the two for periods of years. The discipline of distinguishing them is, in my view, the single most important investment skill of this decade.

◆ ◆ ◆

III. The role of capital duration

A third pattern is the increased importance of capital duration in the kinds of investments that will produce the largest returns. The compounding mechanics I described above operate over years, not quarters. The trust capital required to occupy the experience-orchestration layer accumulates over many cycles. The infrastructure investments required to participate in the AI substrate layer return capital over decades, not months.

The convergence decade therefore rewards capital that can stay deployed for long periods, that is willing to forego near-term liquidity, and that is structurally aligned with multi-year compounding rather than quarterly performance. This favours certain kinds of investors — sovereign capital, family offices, evergreen funds, well-built private-equity vehicles — over others — quarterly-performance-driven public funds, short-duration hedge funds, retail-driven flows. The long-duration capital pools that engage seriously with the convergence decade will, in my forecast, capture a disproportionate share of its returns.

For India specifically, this is an opportunity. Indian sovereign capital, family offices, and the country's growing institutional capital base are positioned to participate in long-duration, high-quality investments that the convergence requires. The institutional architecture for deploying this capital is, in 2026, less developed than it should be. The work of the next five years in Indian financial markets is, in significant part, to build the institutional vehicles that allow long-duration domestic capital to participate in the country's convergence-era opportunities. The countries that have built this institutional architecture have produced their own technology champions. The country that does not will fund its convergence economy with foreign capital, and the value will accrue accordingly.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Re-baseline your investment thesis on the three-layer migration of value.**

   The previous-era patterns will not, mechanically, work in this era.

2. **Distinguish, with discipline, between compounding and accumulating businesses in your portfolio.**

   The distinction is the single most important diagnostic of this decade.

3. **Match the duration of your capital to the duration of the compounding mechanic you are investing in.**

   The duration mismatch is, in my view, the most common cause of investor underperformance in transitions of this character.

CHAPTER FORTY-ONE

41

Valuation in the Convergence Era

The metrics that worked are not the metrics that will.

—

"Price is what you pay. Value is what you get."
— Benjamin Graham, by way of Warren Buffett

◆ ◆ ◆

A senior partner at one of India's largest private-equity firms told me, with uncomfortable honesty, that his firm's valuation models had become, over the previous twenty-four months, less and less predictive of the realised returns of his portfolio. The models continued to produce outputs. The outputs continued to be diligent, defensible, and well-documented. But the outputs were correlating more and more weakly with what the businesses, three or five years later, actually achieved. He had begun to suspect that the underlying frameworks — discounted cash flow, comparable trading multiples, precedent transactions — were, in their conventional application, miscalibrated for the businesses he was now evaluating. He asked me, in effect, what I thought was missing.

The honest answer is that the conventional valuation frameworks were built for businesses with characteristics that the businesses of the convergence decade do not, in most cases, share. The frameworks assume linear revenue growth, stable margins,

identifiable comparables, and capital structures whose cost can be mapped to risk through familiar formulas. The businesses of this decade often have non-linear revenue trajectories, margins that reshape as the business compounds, no clean comparables (because the categories are new), and capital structures whose risk-return profiles do not map cleanly to historical reference points. The frameworks have not become wrong, in any absolute sense. They have become, in many cases, insufficient on their own.

This chapter is about what the senior partner, and others like him, should now be doing differently. It is offered as observation, not as a complete methodology. The complete methodology is the work of the next decade for the valuation profession.

◆ ◆ ◆

I. The metrics that have become more predictive

Five metrics that were peripheral in conventional valuation have, in my observation, become substantially more predictive of the returns of convergence-era businesses. I will name them briefly.

Customer lifetime value to customer acquisition cost ratio over time. The static ratio has always mattered. The trajectory over time matters more. A business whose ratio is improving year over year is compounding. A business whose ratio is deteriorating, even from a high base, is in trouble. Most valuation models capture the static ratio at moment of evaluation. Few capture the trajectory rigorously.

Trust-currency indicators. Customer-recommended-without-incentive rate. Customer-defended-in-public rate. Net promoter trajectory across years rather than at a snapshot. Brand consideration in unprompted recall. These indicators, which most valuation models treat as soft, have become, in my view, substantially more predictive of long-term value creation than

many of the financial indicators that valuation models treat as hard.

Data and AI capability accumulation. The depth, breadth, and uniqueness of the proprietary data the business has accumulated, and the AI capability built on top of it. This is, in many businesses, becoming the primary source of competitive moat. Most valuation models do not capture it well, partly because data is hard to value cleanly and partly because the conventions for doing so have not yet matured.

Talent depth in the irreducibly human capabilities. The list of capabilities I described in Chapter Twenty-Eight — judgment under ambiguity, fiduciary trust, contextual creativity, empathic presence, cultural fluency. The businesses with deep talent in these capabilities will, in my forecast, outperform businesses without them, regardless of the AI tools they deploy. Most valuation models do not measure this depth.

Optionality on the platform-level shifts. The degree to which the business has positioned itself to benefit from, rather than be disrupted by, the major platform shifts of the decade — the shift to agentic AI, the shift to spatial computing, the shift to embedded finance, and others. A business whose strategy is well-aligned with these shifts has option value that conventional valuation models often do not capture.

◆ ◆ ◆

II. The metrics that have become less predictive

Conversely, three categories of metrics that were central to conventional valuation have, in my observation, become substantially less predictive.

Comparable-company multiples. The convergence decade is producing categories that have no clean comparables. The businesses that look like comparables on the surface are often built

on different economic logic underneath. Applying historical multiples to genuinely new categories produces valuations that are, in many cases, dramatically wrong in either direction.

Headcount-based productivity metrics. Revenue per employee, profit per employee, output per FTE — these metrics have been central to valuation in many sectors. They are increasingly miscalibrated as agentic AI absorbs the work of headcount, in ways that vary substantially across firms. A firm that has aggressively adopted agentic AI will look, on these metrics, dramatically more productive than its competitors. Whether the productivity is sustainable, value-creating, and durable is a different question that the metrics do not capture.

Capital-intensity comparisons. The capital intensity of building convergence-era businesses varies dramatically across categories and time. A foundation model company has capital intensity of an order of magnitude higher than an experience-orchestration company. Comparing them on capital-intensity-adjusted metrics produces misleading rankings. Each layer requires its own valuation framework.

◆ ◆ ◆

III. Toward a working framework

The framework I have begun to use, in conversations with investors, has four steps. It is rough. It is improvable. It is, in my view, more predictive than the conventional frameworks alone.

Step 1: Identify the layer. Is the business at the substrate layer, the experience-orchestration layer, or the trust-and-verification layer? Each layer has different unit economics, different capital intensity, different time-to-scale, and different competitive dynamics. The valuation must start from the layer.

Step 2: Diagnose the compounding mechanic. What is the specific mechanism by which the business's competitive position

strengthens over time? Is it data accumulation? Trust capital? Network effects? Switching costs? Brand preference? The strength and durability of the compounding mechanic is, in my view, the single most important valuation input.

Step 3: Test the alignment with platform shifts. Is the business positioned to benefit from the major platform shifts of the decade, or is it positioned to resist them? The alignment is binary in many cases, and the valuation implications are substantial.

Step 4: Apply the conventional frameworks as a check, not as a primary input. The DCF, the comparable multiples, the precedent transactions — these remain useful as sanity checks. They should not, in my view, be the primary inputs in convergence-era valuation. The primary inputs are the layer, the compounding mechanic, and the alignment with platform shifts.

The framework is incomplete. The conventional valuation profession will, over the next five years, develop more rigorous versions. The point of the framework is not to replace conventional approaches but to supplement them with the considerations that conventional approaches systematically under-weight.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you deploy capital, integrate the new metrics — trust currency, data accumulation, talent depth in the irreducibly human, platform-shift alignment — into your evaluation framework.**

   The conventional framework alone will, in this decade, lead to systematic mispricing.

2. **Treat comparable-company analysis with substantial scepticism.**

   Many of this decade's most successful businesses will have no clean comparables, and the comparables that exist will mislead.

3. **Build the in-house capability to evaluate these dimensions, rather than outsourcing it to traditional consulting frameworks.**

   The frameworks have not yet matured. The capability is, for now, a function of the team's pattern recognition rather than of an external methodology.

CHAPTER FORTY-TWO

42

What to Fund, What to Avoid

A pattern catalogue, with the caveats serious investors require.

—

"Bet rarely, but bet big."
— A maxim attributed, with disagreement, to several investors who would all reject the simplification

◆ ◆ ◆

I will, in this chapter, offer a pattern catalogue of the kinds of businesses I would, in 2026, look at favourably and the kinds I would look at sceptically. This is not, of course, a recommendation for any specific investment. It is a pattern catalogue offered for reflection by readers who deploy capital. The patterns are sufficient, not necessary. There will be exceptions in both directions.

◆ ◆ ◆

I. Patterns I would look at favourably

Five patterns, in my honest view, deserve sustained investor attention through the convergence decade.

The vertical-AI category-leader pattern. A firm that has built genuinely deep AI capability for a specific industry vertical — health, law, education, finance, real estate, agriculture — and that has

accumulated proprietary data and customer relationships that broader-purpose AI cannot easily replicate. The defensibility comes from the vertical depth. The compounding comes from the data. The trust capital is built through specialised understanding of the domain. The category leaders in vertical AI will, in my forecast, be among the highest-return investments of the decade.

The ecosystem-orchestrator pattern. A firm that has positioned itself as the conductor of an ecosystem in its category, rather than as a participant trying to do everything in-house. The orchestrator captures the customer relationship, the data flow, the trust currency, and the integration value, while accepting that many components of the offering are produced by partners. The orchestrator pattern is harder to build than the vertical leader, and it requires unusual leadership skill in partner ecosystems, but it produces some of the most durable competitive moats this era will generate.

The trust-anchor pattern. A firm that has positioned itself as a trust anchor — for content authenticity, for identity verification, for AI provenance, for ethics in a specific domain — at a time when trust currency is becoming structurally more valuable. The category is small in revenue but large in influence, and the firms that establish the standards in 2026 to 2028 will hold positions that are extraordinarily difficult to displace.

The frugal-innovation-with-AI pattern. A firm that has applied AI to deliver, at substantially lower cost, services that have historically been available only to wealthy customers. AI-augmented financial advisory at the cost of an algorithm. AI-augmented medical advisory at the cost of a teleconsultation. AI-augmented legal advisory at the cost of a routine subscription. The pattern is particularly powerful in India and emerging markets, where the addressable market is enormous and previously under-served. The firms that build this pattern well will produce returns that are, on

a probability-weighted basis, comparable to the global platforms of the previous decade.

The infrastructure-as-a-bet pattern. A firm that is building the physical or digital infrastructure on which the convergence economy will run — compute capacity, data centres, AI hardware, climate-adaptive infrastructure, defence-relevant capabilities. The capital intensity is high. The time-to-return is long. The strategic importance is substantial. For investors with the duration and scale to participate, the infrastructure layer is, in my forecast, where the most consequential national-level returns will be generated this decade.

◆ ◆ ◆

II. Patterns I would look at sceptically

Five patterns, in my honest view, deserve substantial investor scepticism, even when the businesses are growing rapidly and attracting positive market attention.

The undifferentiated-AI-wrapper pattern. A firm whose primary product is a thin wrapper around a foundation-model API, with no proprietary data, no significant domain depth, no defensible moat, and no compounding mechanic. The firms in this pattern are growing rapidly in 2026, often at attractive multiples. They are also, in my view, structurally exposed to the inevitable improvement of foundation models and the inevitable arrival of more capable competitors. The pattern is not categorically bad — some firms in it will, through skilful execution, build durable positions. The base rate of long-term success, however, is unfavourable.

The destination-business in a category that is dissolving its destination. The bank that is doubling down on branch experience. The retailer that is doubling down on store experience without re-thinking the relationship. The educational institution that is doubling down on the campus without re-thinking the learning. These are not bad

businesses in any absolute sense, but they are positioned against the structural direction of their category. The capital invested in them is, in many cases, capital invested against the convergence decade rather than with it.

The legacy-platform business with no path to AI integration. A firm whose business model depends on platform economics — marketplaces, software-as-a-service, social platforms — that has not credibly integrated AI into its core value proposition. The platform was a competitive advantage in the previous era; without AI integration, it becomes, in this era, a legacy that competitors who started AI-native can disrupt with surprising speed.

The capital-intensive business in a category that AI is making capital-light. Many businesses built fortunes on the capital intensity required to operate in their categories — the capital-intensive distribution networks, the capital-intensive manufacturing, the capital-intensive customer-acquisition machines. AI is, in many of these categories, making the capital requirements substantially smaller. The firms whose competitive moats depended on capital intensity that is now being eroded are exposed in ways that conventional analysis often misses.

The trust-deficit business. A firm that has, through previous behaviour, accumulated a trust deficit with its customers, regulators, employees, or society. In the previous era, trust deficits could often be repaired through marketing and time. In the synthetic-content era, where trust is structurally more valuable and the alternatives are more numerous, trust deficits are more punishing. The firms that have not invested in trust capital, and that carry an active deficit, will struggle disproportionately in this decade.

◆ ◆ ◆

III. The harder pattern: the early-stage exception

One pattern deserves separate treatment because it does not fit cleanly into the favourable or sceptical categories above. It is the pattern of the genuinely novel category — a business attempting something that no comparable business has attempted before. The vertical-AI leader pattern, the ecosystem orchestrator pattern, the trust anchor pattern — each requires the existence of comparable categories from which to draw the pattern. A genuinely novel category, by definition, has no comparable.

For these businesses, conventional pattern recognition fails. The business may be a generation-defining category creator. The business may be a fascinating idea that does not, ultimately, find product-market fit. The business may be both at different times. The investor's task in evaluating these businesses is, in significant part, to evaluate the founder, the depth of insight into the unmet need, the quality of the early customer signal, and the team's ability to navigate uncertainty. Conventional metrics provide weak signals.

The investors who do this well, over the long arc of their careers, are rare. The returns they produce, when they get it right, are extraordinary. The returns they produce, when they get it wrong, are total loss. The discipline required is the discipline of running a portfolio that can absorb the total losses while exposing the portfolio to the extraordinary upside. This discipline is, in my view, not as rare as the success of category-creating investments suggests; what is rare is the institutional capital structure that allows it to be applied without distortion from short-term performance pressure.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your current portfolio against the five favourable patterns and five sceptical patterns.**

   Most portfolios contain meaningful exposure to patterns the holder would not, in retrospect, have chosen to maintain.

2. **For each holding, identify the specific compounding mechanic — and rank holdings by the strength of the mechanic.**

   The ranking will, in my experience, surface several mid-quality holdings that should be reduced in favour of higher-quality compounding businesses.

3. **Establish, for your firm, an explicit posture toward category-creating investments, and the discipline to maintain it.**

   The posture is a strategic choice. The discipline is what separates the investors who succeed at it from the investors who attempt it without success.

CHAPTER FORTY-THREE

43

India's Capital Architecture

Building the institutional vehicles that domestic capital deserves.

—

"Capital is not just money. Capital is patience, conviction, and the willingness to be wrong for a while."

— *paraphrased from various Indian investors who would all claim it*

◆ ◆ ◆

I want to close Part IX with a chapter that is principally addressed to Indian readers, because the questions it treats are sharper here than in most large economies and the answers required are partly a matter of national institutional architecture rather than only of firm-level investment skill.

India in 2026 is, from a capital-architecture perspective, in a peculiar position. The country is producing world-class entrepreneurs, building world-class businesses, and generating extraordinary opportunities at the convergence of its demographic, infrastructural, and technological positions. At the same time, the institutional vehicles for deploying domestic Indian capital into these opportunities at scale and with appropriate duration are, in 2026, less developed than the opportunity merits. The result is that a substantial fraction of Indian convergence-era value is being financed by foreign capital, and a substantial fraction of the long-term returns will accrue accordingly.

This chapter is about that gap, and what closing it requires.

◆ ◆ ◆

I. The patient-capital deficit

The primary gap is in patient capital — capital deployed with a five-to-fifteen-year horizon, willing to forego near-term liquidity, structurally aligned with the multi-year compounding I described in Chapter Forty. The countries that have built this kind of capital architecture — the United States through endowments, family offices, and large pension funds; Singapore and the UAE through sovereign wealth funds; Japan through corporate cross-holding and strategic capital; Sweden through Wallenberg-style multi-generational vehicles — have produced, over decades, disproportionate technology and industrial champions.

India's institutional capital landscape is, in many respects, less mature than the size of the economy would suggest. The pension system, while growing, is small relative to the economy. The endowment culture is nascent. The family offices, though increasingly sophisticated, are individually fragmented and collectively under-coordinated. The sovereign wealth function is not at the scale of comparable economies. The result is that, when long-duration capital is required for genuinely transformative investments, the country must turn, more than it should, to foreign sources.

The work of the next decade is, in significant part, to build out this missing institutional architecture. This is not principally an entrepreneurial task. It is a financial-institutional task that requires sustained policy attention, careful regulatory design, and the deliberate development of the talent pools required to staff the institutions. The countries that have built these institutions over decades have not done so by accident. India is currently in the early stages of doing so. The pace and quality of the building over the

next ten years will substantially shape the country's economic position in 2050.

◆ ◆ ◆

II. The venture and growth equity gap

A specific gap within the patient-capital deficit is in venture and growth equity at the scale the convergence decade requires. India has produced, over the past decade, a credible venture-capital industry; the industry has funded a generation of Indian entrepreneurs to early-stage success. The industry has, however, been substantially dependent on foreign limited-partner capital, particularly from US institutional investors. The growth-equity layer above it has been even more foreign-dependent.

The implication is that, when an Indian convergence-era business reaches the scale at which it requires several hundred million or more dollars of growth capital, the capital is, in most cases, foreign. The valuations are set by foreign investors. The exit pathways are, in many cases, oriented toward foreign acquirers or foreign listings. The long-term value capture is partly outside the country.

The work to address this gap is partly the work of building domestic institutional capital pools, and partly the work of building the regulatory and market infrastructure that allows those pools to deploy efficiently into venture and growth equity. The recent years have seen progress — domestic pension and insurance pools have been progressively permitted to allocate to alternative assets; the regulatory structure for AIFs and InvITs has matured; the Indian listing market has deepened. The progress is real. The pace required to match the opportunity of the convergence decade is faster than the current trajectory.

◆ ◆ ◆

III. The public-markets question

A related question is whether India's public-markets architecture is well-suited to the convergence-era businesses the country is producing. The Indian public markets have, historically, been deep, liquid, and well-regarded for traditional industry sectors — financial services, consumer products, IT services, infrastructure. The markets have been less well-suited to early-stage technology businesses, businesses with unconventional cost structures, and businesses whose value depends on long-term compounding rather than near-term profitability.

Several of the most successful Indian technology businesses of the past decade chose to list abroad — in the United States or in Singapore — partly because the foreign markets were more receptive to their economic logic. Several of those businesses subsequently chose to relist or dual-list domestically as the Indian market matured. The trajectory is positive. The work to be done remains substantial.

For the convergence decade specifically, the public-markets architecture should evolve to support listings of vertical-AI companies, ecosystem orchestrators, infrastructure plays, and other categories that may not, at the moment of listing, have the financial profiles that Indian public-markets investors are accustomed to evaluating. The maturation of the Indian institutional investor base, the development of analyst coverage capable of evaluating these businesses, and the regulatory frameworks that permit appropriate listings are all part of the work.

◆ ◆ ◆

IV. The role of the responsible founder

One consideration runs through this chapter that is properly the responsibility of founders and management teams rather than of

capital institutions. The terms on which Indian convergence-era businesses raise capital, and the institutional partners they choose, will substantially shape who captures the long-term value of their success. The founder who optimises only for the highest valuation at the next financing round, and who chooses partners on transactional terms, will produce a value-distribution outcome that may not align with the founder's longer-term goals or with the country's interests.

The responsible founder, in my view, considers the long-term partnership characteristics of capital sources alongside the immediate financial terms. Domestic institutional partners who can stay the course over fifteen years; foreign partners who bring genuinely strategic value beyond capital; capital structures that keep the founder and the management team aligned with long-term value creation rather than near-term liquidity events. These choices are, in aggregate, more consequential than any individual valuation negotiation. The founders who have made these choices well are the founders whose businesses, at the end of the decade, have produced both extraordinary returns and a value distribution that strengthens rather than weakens the country's broader economic architecture.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you allocate capital institutionally, examine your current allocation to long-duration, India-focused investment vehicles — and increase it deliberately if it is below your considered target.**

   The institutional architecture is built by the cumulative decisions of the institutional investors who participate in it.

2. **If you are an Indian founder raising significant capital, consider the long-term partnership characteristics of capital sources alongside the immediate financial terms.**

   The choice you make at financing time substantially shapes the value-distribution outcomes a decade later.

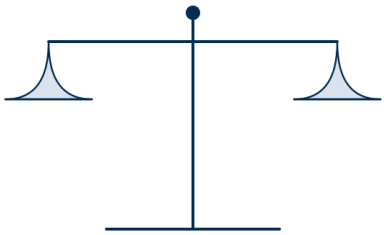
3. **If you advise on capital-markets policy, contribute to the institutional-architecture conversation actively.**

   The policy choices being made now will shape the country's capital architecture for a generation.

PART TEN

X

Policy and the Public Interest



The questions that markets cannot resolve alone, and that the next twenty-four months will substantially decide.

THE WORLD THAT'S BEGINNING

The questions that markets cannot resolve alone, and that the next twenty-four months will substantially decide.

CHAPTER FORTY-FOUR

44

AI Governance — From Principles to Practice

The hard work begins where the white papers end.

—

"Between the idea and the reality, between the motion and the act, falls the shadow."

— *T. S. Eliot, The Hollow Men*

◆ ◆ ◆

I have, over the past three years, read more AI-governance white papers, principle statements, voluntary commitments, and aspirational frameworks than I would have thought possible. Most of them are well-intentioned. Many are well-written. Few of them, in my honest view, contain the operational specificity required to actually govern AI in production. The gap between principles and practice — between the words "fair," "transparent," and "accountable" on a page and the engineering and institutional decisions those words require — is where the work of this decade actually lies.

This chapter is about that gap, and about what closing it requires from leaders, from governments, and from the multilateral institutions that are now attempting to coordinate across borders.

◆ ◆ ◆

I. The principles are converged. The practice is not.

The AI-governance conversations of the past five years have produced a remarkable convergence at the level of principle. The OECD AI Principles, the EU AI Act framing, the various US executive orders and standards, the UK AI Safety Institute, the Indian governance frameworks now under development, the corporate AI principles from the major firms — each emphasises, with minor variations in vocabulary, the same broad set of commitments. AI should be safe. AI should be transparent. AI should be fair. AI should be accountable. AI should respect privacy. AI should not discriminate. AI should be subject to human oversight where consequential.

The convergence at the level of principle is genuine and useful. The convergence at the level of practice is, in 2026, almost entirely absent. What "transparent" means in operational terms — what specifically must be disclosed to whom, in what form, with what frequency, with what verification — is not standardised across jurisdictions or even across firms within jurisdictions. What "accountable" means — who specifically is responsible when an AI system causes harm, in what proportions, to whom, with what remedies — is contested. What "fair" means in the context of statistical models that necessarily produce different outcomes for different groups is a deeply unresolved technical and philosophical question.

The result is that firms attempting to implement responsible AI in good faith are, in 2026, often making consequential decisions in legitimate uncertainty about what the law actually requires of them. Firms attempting to evade responsible AI in bad faith are often able to claim compliance with principles that are too vague to constrain them. Both situations are corrosive to the public trust that the AI transition requires.

◆ ◆ ◆

II. The four governance frontiers

Four specific governance frontiers will, in my view, substantially shape the AI environment of the next decade. Each is being negotiated, in different forums, with different participants, on different timelines.

The frontier of model evaluation and disclosure. What testing must a foundation model undergo before it is released? What capabilities and risks must its developers disclose? What independent evaluation must it submit to? The frontier is being shaped in 2026 by the AI Safety Institutes in the UK, US, and elsewhere; by the EU AI Act's evolving standards; by the voluntary commitments of major AI developers; and by an emerging set of standards organisations. The choices being made in this frontier will, by 2028, become the de facto compliance regime for foundation models globally.

The frontier of liability and accountability. When an AI system causes harm, who is responsible? The user who deployed it? The developer who built it? The provider whose foundation model underlies it? The platform on which it operates? The regulatory regime that permitted its deployment? The frontier is being contested in courts in the United States, in legislatures in Europe, and in regulatory frameworks in India and other jurisdictions. The clarity that emerges over the next three years will substantially shape the risk allocations under which AI is deployed for the rest of the decade.

The frontier of content authenticity and synthetic media. What disclosure obligations apply to AI-generated content? What technical infrastructure is required to make authenticity verifiable? What enforcement mechanisms operate against malicious synthetic content? The frontier is moving, in 2026, through industry-led provenance standards (the C2PA coalition), through legislative responses to election-related deepfakes, and through

evolving platform-level moderation policies. The infrastructure that emerges will be one of the most consequential public-interest decisions of this decade.

The frontier of AI in high-stakes decisions. Where can AI participate in high-stakes decisions affecting individuals — credit, employment, healthcare, criminal justice, education, immigration — and under what constraints? The frontier is the subject of the most active legislative attention globally, and the frameworks emerging vary substantially by jurisdiction. Firms operating internationally will, by 2030, navigate a complex patchwork of obligations that have profound implications for how AI is deployed in consumer-facing applications.

◆ ◆ ◆

III. The Indian opportunity and obligation

India is, in 2026, in a peculiarly strong position on AI governance. The country has the technical capacity to engage seriously. It has the digital public infrastructure that demonstrates how thoughtful governance can produce both innovation and citizen protection. It has a large enough domestic market that its choices will materially shape the products firms build for it. It has a diplomatic posture that allows engagement with both Western and other governance frameworks. It has, in the recent India AI Mission and various regulatory proposals, the institutional architecture beginning to take shape.

The country's opportunity is to develop a governance framework that is genuinely third-way — neither the prescriptive comprehensive regulation of the EU AI Act, nor the lighter-touch sector-by-sector approach of the United States, nor the state-controlled orientation of China, but a framework grounded in India's particular combination of pluralism, public-infrastructure thinking, and pragmatic capacity-building. The framework, well-executed, would influence global AI governance

disproportionately to India's share of AI development, because much of the developing world will look to India's framework as the template for their own.

The obligation is to actually develop the framework with the seriousness the opportunity requires. The temptation, in any large bureaucracy, is to defer the hard choices, produce a vague principles document, and let the practical questions be resolved by litigation or by foreign frameworks operating extraterritorially. This temptation must be resisted. The country's leaders in this domain — in MeitY, NITI Aayog, the AI Mission, the various sectoral regulators, and the legal community — are doing genuinely serious work. The work needs to accelerate.

◆ ◆ ◆

IV. What firms can do regardless

For firms operating in 2026, the appropriate posture is not to wait for governance clarity before acting. The clarity will arrive over years, not quarters; the firms that defer their internal governance work until then will be unprepared when the rules materialise.

The firm-level work is, in summary: develop and publish internal AI principles that are operationally specific rather than aspirational; build the governance infrastructure (review boards, technical evaluation processes, incident response capabilities) that can make those principles real; engage with the standards-setting processes externally; and operate to the strictest reasonable standard across jurisdictions. The cost of this work is meaningful. The cost of failing to do it — when the inevitable incident occurs and the firm has to scramble to demonstrate it acted responsibly — is greater.

The firms that have done this work well in 2025 and 2026 are now in a structurally advantaged position. They have the institutional capacity to deploy AI confidently. They have the public trust to

occupy categories where less prepared firms are excluded. They have the working relationships with regulators that allow them to participate in shaping the frameworks rather than only complying with them. The firms that have deferred the work are increasingly playing catch-up against a moving target.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If your firm has principles for AI use, audit them for operational specificity within the quarter.**

   Principles that cannot be operationalised are aspirational documents, not governance.

2. **Engage actively in at least one governance-standards-setting forum that affects your domain.**

   Participation is more valuable than reaction.

3. **If you advise governments, urge specificity over generality in AI governance frameworks.**

   The hard work is at the level of practice. The principles documents, however well-meaning, are insufficient.

CHAPTER FORTY-FIVE

45

Data Sovereignty and Digital Public Infrastructure

Whose data, on whose terms, hosted where, used how.

—

"He who controls the data controls the future."
— paraphrased from George Orwell, with the appropriate acknowledgement

◆ ◆ ◆

A senior policy adviser to one of India's larger states told me, with the unvarnished candor of a long career, that the question of who controls India's data was, in his view, the single most consequential strategic question facing the country in the next decade. He framed it more sharply than the diplomatic conversation would normally permit. "Our citizens generate the largest pool of behavioral, transactional, and biological data on Earth. The data is, today, principally training the AI capabilities of two foreign countries. The capabilities those AIs develop will, in turn, be sold back to us as services we depend on. If we do not change this trajectory in the next five years, the country will have built the world's most valuable data resource and given it away."

The framing is sharp. It is also, in its essence, accurate. This chapter is about the policy frontier of data sovereignty, the institution that

India has built to address it — the digital public infrastructure stack — and the work that remains.

◆ ◆ ◆

I. The asymmetric data flow

The flow of data between India and the rest of the world, in 2026, is asymmetric in ways that are not widely appreciated. Indian citizens, businesses, and institutions generate data on a scale that, on a per-capita basis, may be lower than in some advanced economies but in absolute volume rivals or exceeds most. This data — search behavior, communication patterns, transaction records, health information, learning interactions, location traces, biometric and genomic data — is, in many cases, processed and stored on infrastructure operated by foreign firms, used to train AI systems whose intellectual property is held by foreign parents, and monetised through services that are sold back into India.

The flow is not, in itself, scandalous. The arrangement has been mutually beneficial in significant respects. Indian users have received high-quality digital services at low or no cost. Indian firms have leveraged foreign infrastructure to scale rapidly. Indian researchers have benefited from access to global AI capabilities. The global AI improvements derived in part from Indian data have accrued, ultimately, to Indian users alongside others.

The asymmetry, however, has consequences that compound over time. The ownership of the AI capabilities derived from Indian data accrues abroad. The strategic value of the inferences drawn from that data accrues abroad. The negotiating leverage in commercial relationships accrues abroad. By 2030, on the current trajectory, the country's negotiating position with foreign technology providers will be substantially weaker than it should be, because the providers will hold more of the country's data, more of the AI capability built on it, and more of the institutional relationships that depend on both.

◆ ◆ ◆

II. Three layers of policy response

The policy response to this asymmetry operates, in my view, on three layers.

The first layer is data localisation — the requirement that certain categories of data be physically stored within India and processed under Indian jurisdiction. India has, over the past decade, progressively expanded localisation requirements, first for payment data, then for personal data more broadly, with current debates around health data, biometric data, and AI training data. The localisation requirements are not, in themselves, sufficient. They are necessary but not adequate. They ensure that the data is hosted within Indian jurisdictional reach; they do not, on their own, ensure that the value derived from the data accrues within India.

The second layer is data sovereignty more broadly — the framework of rights, obligations, and institutional arrangements that governs how data may be collected, processed, and shared across borders. The Digital Personal Data Protection Act, finalised in its current form, establishes much of this framework, with implementation rules now being developed. The Act, in its current shape, gives Indian citizens substantial rights over their personal data and gives the government substantial ability to regulate cross-border flows. The implementation will, in practice, be where the framework's strength is tested.

The third layer is digital public infrastructure — the deliberate construction of public-good platforms (identity, payment, data exchange, commerce, health, education, agriculture) that serve as alternatives to foreign-controlled platforms and that ensure key categories of data flow through architectures over which the country has substantial control. India has built more of this layer than any other large economy. The work continues; the maturation

of the platforms over the next five years is, in my view, the single most important institutional work the country has underway.

◆ ◆ ◆

III. The DPI moment, internationalised

The Indian digital public infrastructure has, in the past three years, attracted substantial international interest. India's G20 presidency in 2023 prominently featured DPI as a frame for global digital cooperation. The Modular Open Source Identity Platform (MOSIP), developed in India and now adopted in multiple African and Southeast Asian countries, demonstrates the exportability of the approach. UPI's interoperability arrangements with several international payment systems extend the model abroad. The OpenAI-style debates about whether comparable infrastructure should exist for AI capability — open foundation models accessible as a public good — draw, explicitly or implicitly, on the DPI template.

The implication is that India's DPI is not only domestic infrastructure but, increasingly, a model that the country is exporting and that influences global digital governance. The country's strategic position in the convergence decade depends, in part, on how this export is conducted — whether through formal multilateral arrangements, bilateral partnerships, open-source community models, or commercial deployments. Each option has different implications for the country's influence, the model's adoption, and the long-term value capture.

The work of the next five years includes, in my view, deliberate institutional development of how India's DPI is exported and on what terms. The country has the substance. It has not yet, in my honest reading, built out the institutional architecture to leverage the substance fully.

◆ ◆ ◆

IV. The harder questions

Two questions, beyond the framework above, deserve sustained policy attention.

The question of AI training data. Should foreign foundation-model developers be permitted to train on Indian data without restriction, with restriction, or with explicit consent and compensation? The economic implications of the answer are substantial. The technical implementation is complex. The international precedent is contested. The frameworks emerging in the EU and elsewhere are imperfect models. India's answer to this question, in 2026 to 2028, will substantially shape the country's AI economic position.

The question of cross-border AI services. Should foreign AI services be permitted to operate in India under what conditions? The current regime is permissive. The mature regime, in my forecast, will be more conditional, in ways that depend on the nature of the service, the data it processes, the alternatives available, and the strategic considerations involved. The regulatory development in this area is, in 2026, at an early stage. The choices made in the next twenty-four months will shape the operating conditions for AI in India for the rest of the decade.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate a firm whose business depends on Indian data, develop an explicit data-sovereignty posture and the operational architecture to support it.**

   The regulatory environment is consolidating. The firms that are prepared will absorb the changes; the firms that are not will scramble.

2. **If you advise on policy, contribute to the harder-question conversations actively.**

   The framework is being decided. The questions are not principally technical; they are strategic, and informed contribution from operating leaders is valuable.

3. **If you operate internationally, plan for the increasing fragmentation of global data regimes.**

   The fragmentation is structural. The firms that have planned for it will operate efficiently; the firms that have not will absorb compliance costs that compound.

CHAPTER FORTY-SIX

46

Labour, Reskilling, and the Social Contract

The transition that markets cannot manage on their own.

—

"The test of our progress is not whether we add more to the abundance of those who have much; it is whether we provide enough for those who have too little."

— Franklin D. Roosevelt

◆ ◆ ◆

I treated the firm-level obligations of the responsible employer in Chapter Thirty. This chapter takes the policy-level view of the same questions — what governments, multilateral institutions, and educational systems must do to manage a transition whose scale exceeds what any single firm can absorb on its own. I have argued throughout this book that the technology of the convergence decade will produce, alongside its considerable benefits, substantial disruption to white-collar and middle-skill employment. The disruption will not, on its current trajectory, be evenly distributed. The policy response, or its absence, will substantially shape who absorbs the costs and who absorbs the gains.

◆ ◆ ◆

I. The scale of the displacement

Estimates of the scale of AI-driven workforce displacement vary widely, in part because the methodology differs across studies and in part because the genuine uncertainty is large. The ranges, however, are sobering. McKinsey's 2023 generative-AI study estimated that occupations representing approximately twenty-five to thirty percent of current US working hours could be substantially automated by 2030. Goldman Sachs's parallel analysis suggested that approximately three hundred million full-time-equivalent positions globally could be exposed to automation. The OECD's 2023 employment outlook estimated that roughly twenty-seven percent of jobs in OECD countries are highly exposed to AI-driven automation. These estimates are not predictions of unemployment at those levels — most affected jobs will be reshaped rather than eliminated — but they indicate the magnitude of the workforce reorganisation underway.

For India specifically, the structural exposure is, in some respects, more serious than for advanced economies. The country's IT services industry, which directly or indirectly employs over five million people and generates substantial export earnings, is exposed to the agentic-AI shift in ways that the leadership of the industry has begun to acknowledge but has not yet fully metabolised. The business process outsourcing industry, employing millions more, is exposed even more sharply. The country's transition to the convergence economy, if managed poorly, could displace several million white-collar workers over a period in which the new categories of work expand more slowly. The social and political consequences of mismanaging this transition are, in my view, substantial.

◆ ◆ ◆

II. The four pillars of a serious policy response

A serious policy response to this transition rests, in my view, on four pillars. Each is necessary; none is sufficient on its own.

The first pillar is income support during transition. Workers whose roles are eliminated or substantially reshaped will, in many cases, require months or years to acquire new skills and find new employment. The income gap during this period must be bridged, either by sufficiently generous unemployment insurance, by direct income support, by employer-funded transition packages, or by some combination. India's current unemployment insurance and social safety net are, on most international comparisons, less developed than the transition will require. The expansion of these systems is one of the largest policy projects of the decade.

The second pillar is large-scale reskilling infrastructure. The reskilling required to move workers from displaced categories to growing categories is substantial, and substantially larger than current educational infrastructure can absorb. National-scale reskilling programmes, delivered through public-private partnerships, leveraging AI-augmented adaptive learning, and connected to actual labour market demand, must be built. India has, in the National Skill Development Corporation and various sectoral skill councils, the institutional architecture to do this work. The quality, scale, and connection-to-demand of the work being done currently is, in my honest reading, below what the transition will require.

The third pillar is labour market intermediation. The matching of displaced workers to new opportunities is a non-trivial problem, especially when the workers have been in roles for many years and the new opportunities have unfamiliar skill requirements. AI-augmented matching platforms, connected to the skills graphs I described in Chapter Fourteen, can substantially improve this matching. The deliberate development of this intermediation

infrastructure, with appropriate public investment and private participation, is the third pillar.

The fourth pillar is the social meaning of work. Beyond the practical pillars, there is a deeper question about how societies will give meaning, identity, and dignity to people in a period when the conventional employment relationship is reshaping. The Indian context is, in this respect, both more challenging and more resilient than many advanced economies. More challenging because the social safety net is thinner and the cultural identification of personal worth with employment is strong. More resilient because Indian families and communities have, historically, supported their members through labour-market transitions in ways that more atomised societies have not. The policy framework should support and strengthen these resilience mechanisms rather than undermining them.

◆ ◆ ◆

III. The reskilling that actually works

The reskilling industry, globally, has a poor track record. Most reskilling programmes produce weak outcomes — high attendance, low completion, weak labour-market translation. The gap between what reskilling programmes claim and what they actually deliver is, in my long observation of the field, substantial.

The reskilling that has produced genuine outcomes shares several characteristics. *It is connected to specific labour-market demand* — not training in abstraction, but training for specific employers who have committed to hire the graduates. *It is rigorous in admissions and assessments* — selecting the workers most likely to succeed and assessing their progress credibly. *It is sufficient in duration* — typically months to a year, not weeks, for genuine skill acquisition. *It is supported during transition* — with stipends, mentorship, and connection to employers, not abandoned once the formal

programme ends. *It is iterated* — programme outcomes measured rigorously and the curriculum revised based on what worked.

Few public reskilling programmes meet all five criteria. The work to build programmes that do is, in significant part, the work of the next decade for any country navigating this transition seriously. India has, in pockets, programmes that approach this standard. The expansion of these programmes to the scale the transition will require is, in my honest view, beyond what current institutional capacity can deliver. The institutional building required is substantial. It is also the necessary work.

◆ ◆ ◆

IV. The cooperative contribution from firms

The policy framework above will not be implemented without substantial cooperation from firms. The firms have the information about which roles are changing on what timeline. The firms have the resources to fund meaningful reskilling at the scale of their own workforces. The firms have the placement capacity to give reskilled workers credible new opportunities. The firms have, in effect, much of what the policy framework requires.

The question is whether they will contribute willingly, through some combination of moral obligation and enlightened self-interest, or whether they will do so only under regulatory pressure. My honest view is that the most thoughtful firms will choose the former path, because they will recognise that the social licence of their operations depends on it and because the talent they will need over the long arc of the transition will come, in significant part, from the workers they helped through the early stages of it. The less thoughtful firms will be moved by regulation, by liability concerns, or by reputational damage. Both paths produce, eventually, comparable contribution. The voluntary path is faster and produces better outcomes.

For Indian leaders specifically, I would argue, the voluntary path is also the right path culturally. The country's industrial leaders, at their best, have always understood their relationship with their workforces and their communities as more than a transactional one. The convergence decade is the moment to renew that understanding deliberately, at scale, through the institutional contributions that the policy framework will require.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you lead a firm of significant scale, commit to a reskilling investment proportionate to the workforce changes your AI strategy implies.**

   The proportionate level is, by my estimate, between five and fifteen percent of one year's compensation per affected worker. Most firms are currently below this level by orders of magnitude.

2. **If you advise on policy, support the four pillars actively — income support, reskilling, intermediation, and the social-meaning conversation.**

   Each pillar is necessary. None is sufficient.

3. **If you operate in education or skill development, audit your programmes against the five characteristics that produce genuine outcomes.**

   The programmes that lack these characteristics consume resources without producing the transitions the country needs.

CHAPTER FORTY-SEVEN

47

Education Policy for the Convergence Decade

The longest-term policy choice this decade requires.

—

"Education is the most powerful weapon
which you can use to change the world."
— Nelson Mandela

◆ ◆ ◆

I treated the institutional question of learning in Chapter Fifteen. This chapter takes the policy view — what governments must do, what they must stop doing, and what the public-private architecture of education in the convergence decade requires.

I begin with an observation that I find difficult to escape. The choices that countries make about education in this decade will, more than any other policy choice, determine their position in 2050. The decisions are, in many cases, unglamorous. They will not show up in any single election cycle. They will not produce headlines that give the politicians who make them visible credit. They will, however, compound over a generation in ways that other policy choices do not.

◆ ◆ ◆

I. The education policy of the previous era

Indian education policy of the previous era — broadly, 1960s through 2010s — accomplished a remarkable expansion of access. The country went from a population that was largely illiterate to a population in which the great majority of children were enrolled in primary education and increasingly in secondary education. The expansion was real, valuable, and historically consequential. It also, in many cases, came at the cost of quality. The schools were built; the teachers were hired; the children were enrolled; but the quality of what occurred inside the classroom remained, in too many cases, poor. The Annual Status of Education Reports (ASER) over two decades documented, with painful consistency, the gap between enrolment and learning outcomes.

The policy reforms of the past decade — the Right to Education Act, the National Education Policy 2020, the various skill-mission initiatives, the digital education programmes — have begun to address the quality gap. The work is genuine and important. It is also, in my view, working from a frame that the convergence decade is rapidly making obsolete. The frame assumes a particular shape of school, a particular shape of teacher-student relationship, a particular structure of curriculum and assessment, that the technologies of this decade are now substantially capable of changing. The reforms within the existing frame are necessary; they will not be sufficient.

Education for the Convergence Decade

Seven shifts that policy must enable, not impede.

Element	FROM (twentieth century)	TO (convergence decade)
Curriculum	Fixed, age-cohorted	Adaptive, AI-tutored, lifelong
Assessment	Annual exams, summative	Continuous, formative, AI-graded
Teacher role	Lecturer and grader	Coach, mentor, designer of experience
Skills measured	Memorisation, recall	Judgement, creativity, collaboration
Credential	Degree at end of cycle	Skills graph, continuously updated
Access	Geographic, economic gates	Universal, AI-democratised
Industry alignment	Loose, lagging	Tight, real-time

FIGURE 47.1 · Seven shifts. None optional in a convergence-era system.

Figure 47.1

◆ ◆ ◆

II. Five policy choices that will shape the next generation

Five specific policy choices will, in my view, substantially shape Indian education in the coming decade. Each is contested. Each has implementation challenges. Each is consequential.

The first is the integration of AI tutoring at scale into the public education system. AI tutors capable of personalised instruction in Indian languages have, by 2026, reached a level of quality at which their integration into public education would substantially improve learning outcomes. The integration is non-trivial — it requires devices, connectivity, teacher training, content development, and quality assurance — but the technical and cost barriers are now, for the first time, surmountable. The choice is whether to integrate AI tutoring deliberately, with quality and equity as primary considerations, or to allow it to enter the system unevenly through private channels that widen the urban-rural and rich-poor gaps. The deliberate integration is, in my honest view, one of the highest-leverage policy choices of the decade.

The second is the modernisation of the curriculum for the convergence-era workforce. The curriculum that students absorb today, particularly in upper-secondary and undergraduate education, is, in many cases, substantially miscalibrated against the work that will exist when they graduate. The AI-augmented workforce of 2030 requires different skills than the curriculum of 2026 emphasises. Critical thinking under uncertainty, working with AI systems, contextual judgment, multi-disciplinary integration, communication across languages and cultures — these are not fringe skills. They are the central skills of the convergence-era worker. The curriculum reform required is substantial. It is also urgent.

The third is the redesign of teacher preparation and continuous development. The teacher in the convergence-era classroom is not principally an instructor delivering content. She is a curator, a mentor, a designer of the student's experience, a guide through the AI-augmented learning environment. The teacher preparation programmes of 2026 are, in most cases, training teachers for the previous role, not the emerging one. The redesign of these programmes — and the continuous professional development of the teachers already in service — is an institutional undertaking of significant scale. The countries that do this well will produce dramatically better educational outcomes. The countries that do not will, over a generation, fall further behind.

The fourth is the reform of assessment and credentialing. The standardised, examination-based assessment that has dominated Indian education for decades is increasingly miscalibrated against what the convergence-era worker actually needs to demonstrate. Skills graphs, portfolio-based credentials, demonstrated capability assessments, and continuous formative evaluation should progressively replace one-time high-stakes examinations as the dominant credentialing mechanisms. The institutional reform is contested — examinations are deeply embedded in the country's educational and labour-market institutions — but the direction of

travel is, in my view, irreversible. The pace of the reform will substantially shape how well the country's workers compete globally.

The fifth is the integration of lifelong learning into the educational architecture. The convergence decade requires the country's adults to re-skill multiple times across their working lives. The educational architecture, as currently configured, is largely built for the under-twenty-five demographic. The integration of adult re-skilling into the educational architecture — through public funding, employer co-funding, individual learning accounts, and the institutional partnerships required to deliver — is the fifth choice. The countries that have begun this integration are positioned to absorb the workforce transition smoothly. The countries that have not will, in the late 2020s, face workforce transition challenges that overwhelm their existing institutions.

◆ ◆ ◆

III. The hardest political economy

Each of the five choices above involves political economy challenges that policy intent alone cannot resolve. The teacher unions, in many states, have legitimate interests in the existing institutional architecture. The examination boards have deeply embedded constituencies. The private education industry, which has grown to fill gaps in the public system, has its own interests. The parents — perhaps the most powerful constituency of all — have expectations shaped by the institutions they themselves grew up in.

The political economy challenges are not arguments against the reforms. They are the conditions within which the reforms must be designed, communicated, and implemented. The leaders who navigate this political economy successfully will, in my view, be remembered as the architects of the country's educational transformation. The leaders who avoid the political-economy

difficulty will preside over an education system that, by the end of this decade, will be substantially mismatched against the country's needs.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in education, identify which of the five policy choices most directly affects your work — and contribute to that conversation actively.**

   The reforms will be designed by those who participate in their design.

2. **If you employ workers in significant numbers, participate in the lifelong-learning architecture as a co-funder and partner.**

   The infrastructure will not be built without substantial private participation.

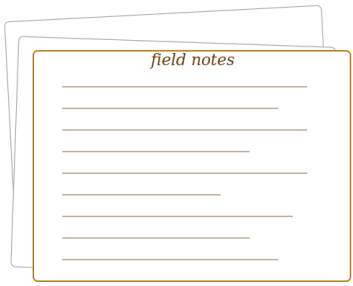
3. **If you advise on policy, advocate the integration of AI tutoring at scale, deliberately and with equity considerations as primary.**

   The integration will happen with or without policy attention. With policy attention, it will be more equitable. Without it, it will substantially widen existing gaps.

PART ELEVEN

XI

Case Studies from the Field



Five composite portraits of organisations navigating the convergence. The names are changed; the patterns are real.

THE WORLD THAT'S BEGINNING

Five composite portraits of organisations navigating the convergence. The names are changed; the patterns are real.

CHAPTER FORTY-EIGHT

48

A Real Estate Developer, Re-platformed

From selling square feet to operating environments.

◆ ◆ ◆

The case I want to walk in this chapter is, to be transparent, a composite. It draws on three engagements over the past four years with Indian real-estate developers — two in north India, one in the south — that, taken together, illustrate the migration from product-economics real estate to environment-as-platform real estate that the convergence decade is making possible. The composite developer I will call Vrishti Estates, after no real entity. The patterns are real. The dialogue is, where attributed, paraphrased from actual conversations conducted with permission to discuss the substance.

◆ ◆ ◆

I. The starting position

Vrishti Estates, in 2022, was a respected mid-tier developer with a twenty-year track record across approximately fifteen residential and three commercial projects in two cities. The company had built a reputation for engineering quality, on-time delivery, and aesthetic design that, in the relevant market segments, commanded a meaningful price premium. Revenue was approximately ₹1,500

crore. Profit margins, at the project level, were healthy. The company was not in distress. It was, by every conventional measure, a successful regional developer.

The chairman, however, was uncomfortable. He had spent the previous eighteen months observing two phenomena that conventional industry reporting did not capture. First, his under-thirty-five buyers were asking different questions than his over-fifty buyers had asked twenty years earlier. The questions were not about square footage, location, or price. The questions were about what the building's air quality would be in 2035, what the building's energy infrastructure would support, what the building's connectivity would allow, what the community of the building's residents would be like, and how the building would adapt as the family's needs changed over the next two decades. Second, the resale market for his completed projects, which had historically been a strong indicator of project quality, was beginning to differentiate by specifications that had not previously mattered — the building's environmental certification, its smart-home integrations, its connectivity infrastructure, its community programming.

He came to a conclusion that he had not previously articulated. His customers had stopped buying square feet. They had started buying integrated living environments, of which square feet were one component. His company was still selling square feet.

◆ ◆ ◆

II. The transformation

Over the next twenty-four months, working with my team and with several partner firms, Vrishti undertook a transformation across four dimensions.

The first was the redesign of the product itself. Future projects were re-architected from the ground up as sensor-rich, adaptive, integrated environments. Each apartment had air quality monitoring, energy

management systems, and ambient computing infrastructure built into the construction rather than added afterward. Common areas were designed for community formation rather than only as pass-through spaces. The buildings were designed to integrate with municipal infrastructure (waste, water, energy) more deeply than previous projects had. The capital cost increased by approximately seven percent. The price premium achieved on sale increased by approximately twenty-two percent.

The second was the construction of an operating layer above the development. Vrishti formed a separate operating company that would, on a long-term basis, operate the integrated environments after the units were sold. The operating company offered residents a recurring service — environment monitoring, predictive maintenance, community programming, AI-assisted resident services, integration with health and education partners. The recurring revenue per unit, after eighteen months of operation, averaged approximately ₹3,500 per month. Across the company's first 800 occupied units, this generated approximately ₹3.4 crore per month of recurring revenue, growing as additional projects came online.

The third was the formation of a partner ecosystem around the residents. Vrishti partnered with selected providers in healthcare (a continuous-care provider, integrated through the operating layer), education (an AI tutoring platform, accessible through the resident portal), retail (curated grocery and household services), and entertainment (a media and event programming partner). Each partnership generated revenue for Vrishti through commercial arrangements and improved the resident experience. The partner ecosystem, the chairman told me, was the single innovation that residents most consistently cited as a reason for staying.

The fourth was the redesign of the customer relationship. Vrishti shifted from a sales-and-handover model to a continuous relationship that began with the customer's first interaction and continued through

ownership, potential resale, and post-occupancy support. The CRM system, the data architecture, the customer-service protocols, and the brand communications all shifted to support this continuous relationship.

◆ ◆ ◆

III. The financial results

The financial results, four years after the transformation began, were striking. Revenue had grown from ₹1,500 crore to approximately ₹2,800 crore, with a substantially higher proportion of recurring revenue from the operating layer. Profit margins, at the consolidated level, had expanded from approximately fifteen percent to approximately twenty-three percent, driven by the higher-margin recurring services. Customer lifetime value, calculated across the unit sale, the recurring services, and the partner-ecosystem revenues, had increased by a factor of approximately 2.4. Customer acquisition cost, as a proportion of customer lifetime value, had declined from approximately five percent to approximately two percent, driven by referrals from existing residents.

More importantly, perhaps, the resale market for Vrishti units had outperformed the comparable market by a meaningful margin. Resale buyers were paying a premium for Vrishti units over comparable competitor units, in part because they were also buying into the operating-layer relationship. The brand had, in three years, become the regional category leader for the integrated-environment segment, a category that, four years earlier, had not existed in the local market.

◆ ◆ ◆

IV. What did not work

Not all of the transformation succeeded. Three specific elements either failed or fell short.

The branded retail experiment. Vrishti attempted to launch its own branded retail offering, primarily groceries and convenience services, sold to its residents through the operating layer. The unit economics, after eighteen months, did not support continuation. The retail capability was not Vrishti's competitive strength, and the partner-ecosystem alternative (sourcing from established retail partners) produced better outcomes at lower cost. The branded retail was discontinued.

The premium tier overreach. Vrishti attempted, in one project, to introduce a premium tier with substantially higher recurring service fees and a more elaborate amenity programme. The tier did not find a sufficient buyer base. The lesson was that the integrated-environment proposition was strong across the developer's traditional buyer segments but that pushing it into a premium tier required different brand positioning, different service capabilities, and different customer acquisition than Vrishti had built. The premium tier was repositioned to align with the rest of the portfolio.

The cross-city expansion. Vrishti attempted, ahead of plan, to expand the new model to a third city. The expansion required local partner relationships, local regulatory engagement, and local market understanding that the company had not built. The third-city projects, eighteen months in, were behind plan. The lesson was that the integrated-environment model is partly a function of local context — local partners, local culture, local infrastructure — and that the model does not transfer cleanly across markets without dedicated local development.

◆ ◆ ◆

V. The lessons

The lessons from the Vrishti case generalise, in my view, to most consumer-facing real estate in India and likely elsewhere.

The first lesson is that the migration from product to environment is real, and that the developers who lead it in their markets command substantially better economics than developers who continue to operate the conventional model. The second is that the migration is operationally substantial — it requires not just new product features but new operating capabilities, new partner relationships, new customer relationships, and new revenue models. The third is that not all elements of the migration succeed; the discipline is to attempt the changes deliberately, measure honestly, and continue the elements that work while discontinuing those that do not. The fourth is that the migration is, financially, accretive — when executed well, it produces both revenue growth and margin expansion. The fifth is that the migration is, strategically, defensive — the developers who do not undertake it will, over a decade, lose share to those who do.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in real estate, audit your customer journey from sale through occupancy through resale, and identify where you have ended your relationship prematurely.**

   The premature ending is, in most cases, where the recurring-revenue opportunity sits.

2. **Identify the operating capability you would need to build to extend your customer relationship beyond the sale.**

   The capability is the strategic question. The opportunity is enabled by it.

3. **Begin with one project as a deliberate test of the integrated-environment model.**

   The capability is built through doing, not through planning.

CHAPTER FORTY-NINE

49

A Hospital Chain Goes Continuous

From procedures to relationships.

◆ ◆ ◆

The case I want to walk in this chapter draws on engagements with two private hospital chains in India over the past three years. The composite I will call Aaroh Health, after no real entity. The journey is a real journey, and the patterns are widely applicable across Indian healthcare.

◆ ◆ ◆

I. The starting position

Aaroh Health, in 2023, operated approximately fourteen hospitals across four states, with an aggregate bed count of approximately 2,800 and annual revenue of approximately ₹1,800 crore. The chain was profitable, well-regarded clinically, and had completed two waves of expansion through the previous decade. The chairman, however, was reading the same indicators that the chairman of Vrishti Estates was reading, with comparable concern.

His indicators were different in substance but similar in pattern. His patients, particularly the under-fifty cohort, were increasingly approaching the hospital with prior diagnostic context that the hospital had not generated and, in many cases, did not have efficient ways to incorporate. His patients were increasingly

expecting follow-up engagement that the hospital was not structured to provide. His patients were increasingly comparing his hospital not just to other hospitals but to digitally native health platforms that offered different relationships for different stages of the patient journey. His revenue per bed-day was healthy. His revenue per patient relationship over time was, when he calculated it, weak. The hospital saw the patient at high-margin moments and lost her in between.

◆ ◆ ◆

II. The transformation

Over the next thirty months, Aaroh undertook a transformation across five dimensions.

The first was the integration of continuous-care offerings. Aaroh launched, initially as a pilot in two cities and subsequently across the chain, a continuous-care subscription that connected patients to wearable-based monitoring, AI-mediated health check-ins, primary-care advisory, and rapid escalation to the chain's specialists when needed. The subscription, priced at approximately ₹400 to ₹1,200 per month depending on the tier, was positioned as a complement to insurance rather than a replacement.

The second was the modernisation of the clinical workflow. The intake AI, ambient transcription, AI-assisted diagnosis support, agentic billing and follow-up scheduling — each was rolled out progressively across the chain. The clinical staff time per patient encounter shifted significantly toward direct patient interaction and away from documentation, in ways that the staff broadly welcomed. The patient throughput per physician increased by approximately forty percent over eighteen months, with patient satisfaction scores increasing rather than decreasing.

The third was the redesign of the patient relationship. Aaroh created, for each patient, a longitudinal health record that the patient owned,

accessed through a dedicated application, and could share selectively with providers inside and outside the Aaroh ecosystem. The record integrated clinical encounters, wearable data, lab results, prescriptions, and the patient's own annotations. The patient-experience implications were substantial. The strategic implications — the patient had a continuous relationship with Aaroh whether she was actively in clinical encounter or not — were even more substantial.

The fourth was the formation of partner relationships. Aaroh partnered with a wearable provider for hardware, with a longevity-clinic partner for a higher-end care offering, with several mental-health platforms for integrated mental-health support, and with a pharmacy chain for prescription fulfilment integrated into the patient app. Each partnership extended the patient's experience of Aaroh while keeping Aaroh's capital deployment focused on its core capabilities.

The fifth was the engagement with employer customers. Aaroh shifted significant marketing and sales attention from individual patients to employers, particularly larger enterprise employers seeking to manage their workforces' health more directly. The employer offering combined the continuous-care subscription, the clinical infrastructure, and the data architecture in ways that pure insurance products could not. By month thirty of the transformation, employer-channel revenue had grown to approximately twenty-eight percent of new revenue.

◆ ◆ ◆

III. The financial and clinical results

The financial results were positive but more modest, in the short term, than Vrishti's. Revenue grew from approximately ₹1,800 crore to approximately ₹2,400 crore over thirty months — solid but not spectacular. The recurring-subscription revenue, while growing rapidly, was still small in absolute terms. Operating

margin expanded by approximately three percentage points. The metrics that improved most dramatically were less financial — patient satisfaction, clinical outcomes on managed-condition cohorts, and patient retention across episodes.

The strategic position, however, had shifted substantially. Aaroh was, by the end of month thirty, positioned as a continuous-health provider rather than as a hospital chain. The valuation that the company commanded in a partial financing transaction reflected this — a multiple substantially higher than peer hospital chains, justified by the higher-quality recurring revenue, the deeper customer relationships, and the strategic optionality the platform created.

◆ ◆ ◆

IV. What worked, what did not

The elements that worked best were the clinical workflow modernisation, where the productivity and patient-experience gains were both real and significant, and the longitudinal record, where the patient response was uniformly positive and the strategic implications proved substantial. The elements that worked less well were the initial pricing of the continuous-care subscription, which was set too low and had to be revised upward; and the early branding of the platform, which conflated the hospital and the platform in ways that confused the market and required rework.

The lesson on pricing, in particular, is widely applicable. Continuous-care subscriptions, if priced as add-ons to existing services, are typically underpriced relative to the value they create. The honest pricing reflects the value of continuous health rather than the marginal cost of the digital service.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in healthcare, define your continuous-care offering this year, even in a minimal initial form.**

   The capability is built through operating, not through planning.

2. **Modernise the clinical workflow before launching the patient-facing experience.**

   The internal capabilities must be in place before the external promise is made.

3. **Engage the employer channel as a primary commercial route.**

   The employers' interest in workforce health is a substantial, growing, and currently under-served opportunity in Indian healthcare.

CHAPTER FIFTY

50

A Public Utility Becomes Predictive

From outage response to outage prevention.

◆ ◆ ◆

The case I want to walk in this chapter is, again, a composite, drawing on engagements with two state electricity utilities in different parts of India over the past four years. I will call the composite Suryadeep Power. The work it did, and the lessons it learned, illustrate what public utilities — electricity, water, gas, transport — can become in the convergence decade.

◆ ◆ ◆

I. The starting position

Suryadeep Power served approximately thirty-eight million consumers across an Indian state, with a generating, transmission, and distribution footprint of substantial scale. It was, by the standards of Indian state utilities, well-managed. Its performance metrics — aggregate technical and commercial losses, supply hours per consumer, billing efficiency — were better than the median state utility in the country, though substantially behind the best globally. The chairman, a career civil servant with deep technical capability, was, like the previous two leaders in this chapter, uncomfortable with what the metrics did not capture.

The metrics did not capture, he said, the citizen's experience of the utility. The citizen's experience was, in a word, reactive. When the supply failed, the utility responded. When the supply was variable, the utility coped. When the consumer queried a bill, the utility eventually answered. The relationship between utility and consumer was almost entirely defined by the utility's response to problems. There was, in his words, "no relationship in the moments when nothing was wrong, which is most of the time."

He believed that the convergence technologies could change this fundamentally. The work he led over the next forty-eight months tested that belief.

◆ ◆ ◆

II. The transformation

Five elements of the transformation deserve specific attention.

The first was the deployment of comprehensive grid sensing. Suryadeep, with central-government and external-financing support, deployed approximately 4.6 million smart meters across its consumer base, alongside transformer-level monitoring and feeder-level analytics. The data architecture was rebuilt to ingest, store, and analyse the resulting data flows in real time. The capital cost was substantial; the operational implications were transformative. The utility moved from monthly meter-reading cycles to fifteen-minute consumption visibility across its entire footprint.

The second was the deployment of predictive analytics for outages and equipment failures. AI systems, trained on the new data flows, learned to predict transformer failures, cable degradation, and demand spikes hours to days in advance. Maintenance shifted from scheduled and reactive to predictive. The number of unplanned outages declined by approximately sixty-three percent

over thirty months. The duration of outages, when they did occur, declined by approximately forty-one percent.

The third was the redesign of the consumer relationship. Consumers received, through a consumer-facing application, real-time consumption information, predictive billing, energy efficiency advisory, and proactive outage notifications. The application also enabled bill payment, complaint registration, and service requests. The combination shifted the consumer's experience from reactive to proactive. The number of consumer complaints declined by approximately fifty-one percent over thirty months. Consumer satisfaction scores improved substantially.

The fourth was the integration with rooftop solar and distributed energy. As Indian rooftop solar adoption expanded, Suryadeep moved early to integrate distributed energy into its grid management. Net-metering became efficient. Solar adoption among Suryadeep's consumers grew faster than in comparable utilities. The state's renewable-energy targets were met ahead of schedule.

The fifth was the workforce transformation. The transformation absorbed the human work of meter reading, routine maintenance scheduling, and complaint logging, while elevating field staff to higher-value functions — predictive maintenance execution, consumer-service relationships, and renewable-energy integration support. The headcount remained approximately constant; the work shifted substantially. The training programme to support this shift was, in my view, the single most important element of the transformation, and the one that other utilities have most often gotten wrong.

◆ ◆ ◆

III. The financial and operational results

The financial results, after four years, were significant. Aggregate technical and commercial losses had declined from approximately

twenty-four percent to approximately thirteen percent — an extraordinary improvement by Indian utility standards, generating annual savings in the hundreds of crores. Tariff increases were avoided over consecutive years. The state government, which had previously subsidised the utility heavily, saw its subsidy burden decline. The utility's credit metrics improved enough to enable a meaningful capital-markets transaction.

The operational results were perhaps more strategically important. The utility had moved from a position of perpetual catch-up to a position of operational leadership. Its peer utilities in other states were studying its work. Its CEO was speaking at international forums on grid modernisation. The state's industrial customers, who had previously considered captive generation seriously because of supply unreliability, were extending their commitment to grid supply.

◆ ◆ ◆

IV. The lessons

The lessons from the Suryadeep case generalise, in my view, to most public-utility infrastructure across India and likely across emerging economies.

The first lesson is that the technology to deliver dramatically improved utility performance now exists, is affordable at scale, and produces returns that justify the investment within reasonable horizons. The second is that the transformation is, fundamentally, an institutional transformation rather than a technology project. The technology is necessary; the institutional capability to absorb it productively is what determines outcomes. The third is that the workforce transformation is the most important element. Utilities that deploy the technology without redesigning the workforce produce uneven results; utilities that redesign the workforce alongside the technology produce transformative results. The fourth is that the consumer relationship is the most strategically

important element. The improved technical metrics matter; the improved consumer relationship matters more, because it shifts the political economy of the utility's operating environment and creates the foundation for further investment.

The Indian public sector has, in pockets, demonstrated the ability to execute transformations of this scale. The replication of these transformations across the country's many utilities is the single most important infrastructure project of the next decade for Indian public services.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you lead a public utility, commit to comprehensive sensing as a foundational investment.**

   The visibility is the precondition for every subsequent capability.

2. **Redesign the workforce alongside the technology, not after it.**

   The two transformations must move together; sequenced separately, neither succeeds.

3. **Treat the consumer relationship as a primary strategic objective, not as a secondary outcome.**

   The improved relationship is, ultimately, what permits the broader transformation to continue.

CHAPTER FIFTY-ONE

51

A Faith Institution Goes Digital

Where the technology meets the deepest human needs.

◆ ◆ ◆

I want to close the case-studies part with a case that may surprise some readers. Some of the most thoughtful applications of the convergence technologies I have encountered, over the past five years, have been at religious and faith institutions in India. The institutions are deeply rooted in millennia-old practices. They serve constituencies that range from highly traditional to highly modern. They operate at scales — millions of devotees, in some cases — that few commercial enterprises match. And they have, in several cases I have worked with, moved with thoughtfulness and seriousness to integrate technology into their work in ways that respect the tradition and serve the constituencies.

The composite case I will walk is, again, a composite. I will call the institution Sankalpa Mandir. It draws on engagements with three different faith institutions — two Hindu temples of significant scale, one inter-faith organisation — over the past four years.

◆ ◆ ◆

I. The starting position

Sankalpa Mandir, in 2022, was a major north Indian temple complex serving approximately fourteen million devotees per year

for darshan and many millions more for festivals, ceremonies, and services. The institution operated a substantial physical infrastructure, employed several hundred priests and staff, ran significant charitable activities, and managed a complex web of administrative, religious, cultural, and commercial functions.

The trustees, in early 2022, commissioned a strategic review that focused on a question that had been growing in their minds. The question was whether the institution was serving its devotees as well as it could in the changing context of how those devotees lived their lives. The under-forty devotees were increasingly distributed across geographies, increasingly digital in their daily practices, increasingly time-constrained, and increasingly seeking forms of connection with the institution that did not require physical presence at the temple complex. The institution's existing services, designed for a constituency that was largely local and present, were not, the trustees believed, fully meeting these constituencies.

◆ ◆ ◆

II. The transformation

The transformation that followed was carefully designed and respectfully executed. Five elements deserve attention.

The first was the construction of a rigorous virtual-darshan capability. The institution invested in immersive recording of its principal sanctum at high fidelity, with the religious authorities setting careful guidelines about what could and could not be recorded and shared. The resulting virtual-darshan, accessible through an application and through immersive devices, allowed devotees who could not be physically present to participate in darshan from anywhere. The reception, the institution reports, has been remarkable. Many devotees who had not been able to visit the temple complex for years have begun a continuous virtual-darshan practice. The daily traffic on the virtual-darshan channels is, by 2026, several times the foot traffic at the physical complex.

The second was the development of an integrated devotee relationship platform. Devotees who chose to register received a personalised relationship with the institution that connected their devotional practice to the institution's calendar, allowed them to participate in ceremonies and rituals at distance, supported their charitable contributions, and enabled communication with priests and counsellors associated with the institution. The platform handled, by 2026, several hundred million devotee interactions per year.

The third was the integration of AI-assisted spiritual counsel. This was, predictably, the most carefully developed element. The institution worked with senior religious authorities to design AI assistants that could answer routine questions about devotional practices, festival significance, ritual procedures, and interpretive teachings, while being clearly bounded against doctrinal positions that the religious authorities reserved to themselves. The AI assistants operated in multiple Indian languages, with cultural and linguistic register appropriate to each. They were positioned as supplements to, not replacements for, human spiritual guidance. The reception has been positive.

The fourth was the modernisation of the institution's operational infrastructure. The administrative, financial, and logistical functions of the temple complex were progressively modernised — donation management, ceremony scheduling, volunteer coordination, accommodation booking, food service — using AI-augmented operations. The result was a substantial improvement in the experience of devotees who did visit physically, alongside a substantial improvement in operational efficiency and accountability.

The fifth was the integration of immersive learning and cultural programming. The institution developed immersive educational experiences for devotees and visitors — walking through the institution's history, the significance of its rituals, the cultural and philosophical depth of its tradition — accessible through both on-

site and remote interfaces. The programming has, the institution reports, substantially deepened the engagement of younger devotees who had previously been more distant from the tradition.

◆ ◆ ◆

III. The reception and the controversy

The reception of the transformation has been, on the whole, positive. Devotee satisfaction has improved across most measures. The institution has expanded its reach without diluting its religious or cultural authenticity. The religious authorities involved have, with care and reflection, embraced the transformation as a contemporary expression of the tradition's longstanding posture toward continuous renewal.

The transformation has not been without controversy, however. Some traditionalist constituencies argued that the virtual-darshan diminished the sanctity of physical presence at the temple. Some commentators argued that the AI counsel was inappropriate in the spiritual domain regardless of how carefully it was bounded. Some critics argued that the modernisation reflected commercial more than spiritual considerations.

The institution's response to these critiques has, I think, been instructive. They engaged the critiques seriously rather than dismissing them. They acknowledged that the technology is one expression among several, and that those who preferred the traditional modes of practice would find their traditional modes preserved and supported. They committed to ongoing reflection about which technological capabilities should be deployed and which should not. The combination of seriousness, humility, and commitment, in my observation, defused most of the genuine concerns and improved the trust between the institution and its broader constituencies.

◆ ◆ ◆

IV. The lessons

The lessons from this case are perhaps less commercially specific than the previous three but, I think, more deeply applicable across the institutions that serve human meaning rather than human consumption.

The first lesson is that technology, thoughtfully integrated, can serve traditions of great depth without diminishing them. The institutions that have done this well have done it by leading with religious and cultural authority, with technology in service rather than in command.

The second lesson is that the constituencies these institutions serve are larger, more dispersed, and more diverse than the institutions' physical footprints alone can serve. The technology allows institutions to serve constituencies they could not previously reach, in ways that complement rather than replace their physical presence.

The third lesson is that the trust currency, in these institutions, is everything. The transformation succeeds when the institution's accumulated trust is preserved and deepened. It fails when the technology is deployed in ways that erode the trust. The discipline required to do the former rather than the latter is, I think, what most distinguishes the institutions that have done this work well from those that have not.

The fourth lesson is more general. The technologies of the convergence decade are not, fundamentally, about commerce. They are about how human beings engage with one another, with their institutions, with their traditions, with their meaning. Faith institutions that have understood this — and there are now many across the world — have used these technologies to serve their deepest purposes more powerfully than would otherwise have been possible. The lesson generalises beyond faith institutions to all institutions whose purpose is principally about human meaning rather than human consumption — educational institutions,

cultural institutions, civic institutions, family institutions. The technology is a servant of meaning, when we choose to use it that way. It is also a corrosive of meaning, when we do not.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **If you operate in any institution whose purpose is principally about meaning, examine how the convergence technologies could serve that purpose more deeply.**

   The technology is available. The institutional imagination required to use it well is rarer.

2. **Lead with the institutional authority, not with the technology.**

   The technologies that succeed in these contexts succeed because the institutional authorities embrace them as expressions of the tradition.

3. **Engage the critiques seriously.**

   The critics often see what the enthusiasts do not. The institutions that engage them well end up with stronger transformations than the institutions that dismiss them.

CHAPTER FIFTY-TWO

52

Patterns Across the Cases

What the four transformations share — and what they do not.

◆ ◆ ◆

Before I close Part XI, I want to step back and observe what the four cases share, and what they do not. The exercise is brief but, I think, useful for any reader contemplating a comparable transformation in her own institution.

◆ ◆ ◆

I. What they share

Five patterns, in my view, run across the four cases.

The first pattern is that each transformation began with an honest diagnosis of what the existing operating model was missing. In each case, the leader looked beyond conventional metrics to ask what the customer or constituent was actually experiencing, and recognised that the institution was structurally absent from the moments where the relationship now needed to live. This diagnosis is, I have come to believe, the single most important precondition for transformation. Institutions that have not made it do not transform; they merely deploy technology onto existing operating models.

The second pattern is that each transformation involved both a customer-experience layer and an operational-capability layer, deployed together

rather than sequentially. The customer-experience layer (the resident application, the patient platform, the consumer app, the devotee platform) made the transformation visible. The operational-capability layer (the operating company, the clinical workflow modernisation, the grid sensing, the temple operations) made the transformation real. Institutions that deploy only the customer-experience layer produce hollow transformations; institutions that deploy only the operational layer produce invisible transformations.

The third pattern is that each transformation extended rather than shrunk the institution's relationship with the people it served. In every case, the relationship became more continuous, more integrated, and more durable. This extension was, in financial terms, the source of most of the value created. It was also, in non-financial terms, what most fundamentally changed the institution's character.

The fourth pattern is that each transformation produced new partnership relationships, alongside the institution's existing capabilities. The institutions that attempted to do everything in-house produced thinner offerings; the institutions that built partner ecosystems produced richer ones. This is, I think, a structural feature of the convergence decade. The institutions that win are orchestrators rather than full-stack producers.

The fifth pattern is that each transformation required substantial workforce transformation alongside the technology deployment. The institutions that addressed the workforce transformation explicitly produced sustainable transformations. The institutions that addressed it as an afterthought produced fragile ones.



Figure 52.1

◆ ◆ ◆

II. What they do not share

The cases differ in several important respects, and the differences are themselves instructive.

They differ in the timeline of the transformation. Vrishti's transformation produced visible results in two years; Aaroh's took thirty months; Suryadeep's took four years. The timelines reflect the complexity of the institution and the depth of the transformation. Leaders should expect their transformation to take longer than they initially plan, and to value the long-arc execution discipline accordingly.

They differ in the financial trajectory of the transformation. Vrishti's transformation produced strong financial results within two years; Aaroh's transformation produced strong strategic positioning but more gradual financial results; Suryadeep's transformation produced strong operational and financial results over a longer arc. The financial trajectory depends on the unit economics of the new operating model, which vary substantially across categories.

They differ in the relationship to existing institutional authority. Vrishti's transformation extended the existing commercial authority; Aaroh's extended the clinical authority; Suryadeep's extended the operational authority; Sankalpa's extended the religious authority. The transformations succeed when they are clearly in service of, and led by, the existing institutional authority. They struggle when they appear to bypass or supersede that authority.

◆ ◆ ◆

III. The general implication

The general implication is that transformation in the convergence decade is, fundamentally, an institutional discipline rather than a technological one. The technology is available. The institutions capable of absorbing it productively are rarer. The leaders who can build that capability — through the diagnosis, the dual-layer deployment, the relationship extension, the partner ecosystem, and the workforce transformation — are rarer still.

This is, in my honest view, why the convergence decade will be a decade of substantial divergence in institutional outcomes. The technology will be approximately equally available to comparable institutions. The institutional capabilities to absorb it will not. The institutions that build the capabilities will pull substantially ahead. The institutions that do not will fall substantially behind. The gap, by 2030, will be visible to anyone looking at any category of institution.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Apply the five-pattern test to any transformation you are currently leading or contemplating.**

   The pattern test is a quick diagnostic of whether the transformation will succeed.

2. **If you are missing one or more of the five patterns, address the gap before continuing.**

   The transformations that ignore the gaps produce thin, fragile, often disappointing results.

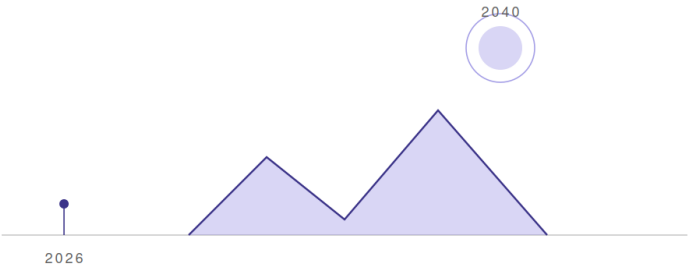
3. **Plan for a longer arc than you initially imagine.**

   Most transformations of this character take three to five years to produce their full results. The leaders who plan for that arc execute it. The leaders who plan for shorter arcs are disappointed and abandon prematurely.

PART TWELVE

XII

The Longer View



Past 2030. What the convergence opens up, what it asks of us, and what we should be building toward.

THE WORLD THAT'S BEGINNING

Past 2030. What the convergence opens up, what it asks of us, and what we should be building toward.

CHAPTER FIFTY-THREE

53

Beyond 2030 — What 2040 Looks Like

The next inflections are already visible to those who know where to look.

—

"The best way to predict the future is to invent it."
— Alan Kay, 1971

◆ ◆ ◆

I have, throughout this book, focused principally on the decade between 2026 and 2030 — the decade in which the convergence becomes real, the decade in which leaders and institutions must make their consequential decisions, the decade in which the foundations of the rest of the century will be laid. The focus has been deliberate. A book that drifts too far into the long future loses its operational usefulness for the leaders who must act now. A book that confines itself only to the near term, however, misses the strategic context that the longer view provides.

This chapter is the longer view. It asks what the world looks like in 2040, ten years past the horizon of most of this book. The forecast is, of course, less reliable at this distance. The patterns visible in 2026 will combine with developments that no one can confidently predict. But the structural shape of 2040 is, in my honest view, more

visible than most observers credit, because the technologies that will define it are largely present in some form in 2026, and the institutional decisions that will shape it are largely being made now.

◆ ◆ ◆

I. The five trajectories converged

The five forces I described in Chapter Two will, by 2040, have substantially converged into a single integrated technological substrate. AI will be ubiquitous and ambient, present in every interface and most physical environments. Spatial computing will have replaced screens for most consumer interactions. Quantum computing will be commercially mature in specific categories. Programmable biology will have produced the first comprehensive personalised medicine, the first credible longevity interventions, and substantial bio-manufacturing capacity. Energy and climate technology will have largely completed the transition off fossil fuels in advanced economies and made substantial progress in the rest of the world.

The integration of these forces produces capabilities that are difficult to describe in 2026 vocabulary because they require new categories. The ambient AI that knows us across years, contributes to our decisions continuously, mediates our relationships with institutions, and progressively becomes part of how we think — this is not "AI" in the sense the term carried in 2024. It is something new. The integrated digital-physical environments we inhabit — sensor-rich, predictive, adaptive, individualised — are not "buildings" or "cities" in the sense those terms carried before. They are something new. The services that draw on our biological specifics in real time to optimise our health, our nutrition, our cognitive performance, and our pharmaceutical prescriptions are not "healthcare" in the conventional sense. They are something new.

The leaders of 2040 will be operating in this new operating environment with a fluency that 2026 leaders are only beginning to develop. The institutions that will serve them will be institutions designed for that environment, not institutions retrofitted from previous eras. The next decade is, in significant part, the decade of designing and building those institutions.

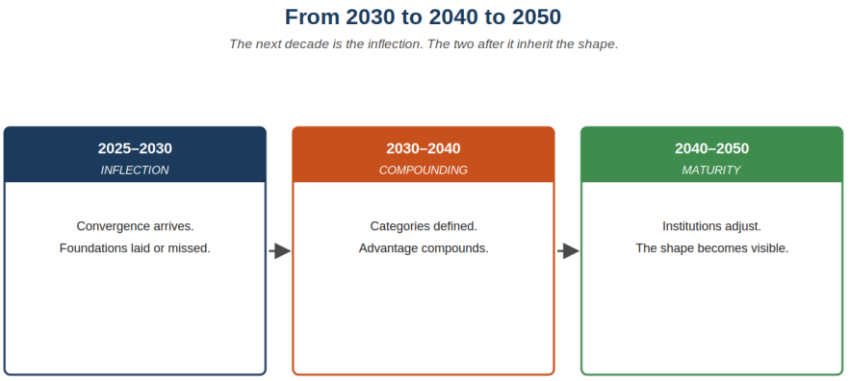


Figure 53.1

◆ ◆ ◆

II. The new arrivals

Three technological developments that are still in early form in 2026 will, in my honest forecast, be substantial features of the 2040 landscape.

The first is genuinely capable artificial general intelligence. Whether by the strict definition (matching human capability across essentially all cognitive tasks) or by a looser definition (matching human capability across most professional tasks), AI by 2040 will, in my view, have crossed a threshold that 2026 AI has not. The implications are profound and not fully predictable. Some categories of human work will become substantially more leveraged, with one human supervising what teams of humans previously did. Other categories will be substantially redefined.

Scientific research will accelerate in ways that compress decades of progress into years. The institutional, policy, and cultural responses to this development are, in 2026, only beginning to be sketched.

The second is the maturation of brain-computer interfaces. Direct neural interfaces — first for medical applications, then for consumer applications, then for professional applications — will have moved from research demonstrations to meaningful commercial deployment by 2040. The early consumer applications will, in my forecast, be principally for hands-free interaction with digital systems, with reading and writing capabilities developing later. The implications for accessibility, for productivity, and ultimately for the nature of human-computer interaction are substantial. The ethical and policy questions are, similarly, still in early development.

The third is the meaningful presence of intelligent humanoid and task-specific robots. The robots of 2040 will not be the science-fiction robots of 1980s imagination, but they will be genuinely capable physical agents in homes, workplaces, healthcare settings, and public spaces. The labor market implications, the social-relationship implications, and the institutional implications of meaningfully present robots are, again, substantial and only beginning to be addressed.

◆ ◆ ◆

III. The continuing tensions

Three tensions visible in 2026 will, in my forecast, remain prominent in 2040, with their resolution varying substantially across societies.

The tension between centralised AI and distributed AI. The path of AI development has favoured centralisation — large foundation models, run on substantial compute, controlled by a small number

of firms. The counter-development — open-source models, federated learning, on-device intelligence — has been steady but slower. By 2040, the balance between these two patterns will be different in different societies. Societies that have invested in the distributed pattern will have more diverse AI ecosystems. Societies that have not will have more concentrated ones. The implications for innovation, for competition, and for the political economy of technology are significant.

The tension between technological capability and human agency. The convergence technologies will, by 2040, enable many things that humans could not previously do. They will also enable many things that human institutions could not previously do to humans. The balance between expanded individual capability and expanded institutional capability is contested in 2026 and will remain contested in 2040. The societies that have built strong individual rights, strong privacy infrastructures, and strong institutional accountability will preserve human agency. The societies that have not will see human agency progressively reduced.

The tension between speed and wisdom. The technologies are arriving faster than most institutions are absorbing them well. The mismatch between technological capability and institutional capacity to deploy it wisely is a structural feature of this decade and will likely remain so. The societies that have built institutional capacity to absorb the technologies thoughtfully will compound advantages over time. The societies that have not will, periodically, generate failures that erode the trust on which broader adoption depends.

◆ ◆ ◆

IV. India in 2040

For India specifically, the 2040 picture, on current trajectories, is one of substantial transformation. The country will, by then, have

completed its demographic dividend window and be entering the next phase of its demographic arc. Its economy will have grown substantially, with the trajectory depending on the policy choices made in the coming decade. Its position in the global order will be more central than it is in 2026, both because of its economic scale and because of the convergence-era infrastructure it will have built.

The specific shape of Indian society in 2040 will depend on choices being made now. If the convergence technologies are deployed equitably, with strong public-infrastructure underpinnings, with serious attention to workforce transitions, and with the civilisational depth that the country is capable of bringing to design — the India of 2040 will be a substantial demonstration of what good technology, well-deployed, can accomplish for human flourishing at the scale of a continental civilisation. If the technologies are deployed inequitably, with weak institutional infrastructure, with neglected workforce transitions, and without the civilisational rooting — the India of 2040 will be a society of deep divisions, with those who participated in the convergence economy having pulled far ahead of those who did not.

The choice is not a matter of certainty. It is a matter of the cumulative effect of the institutional decisions being made by leaders, in 2026 and 2027, in the firms, the public agencies, the educational institutions, and the civil-society bodies of the country. The leaders who take these decisions seriously will, in 2040, be remembered as the architects of the country's transformation. The leaders who do not will be remembered, if at all, as having been present at the moment but absent from the work.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Maintain a explicit, written, ten-year strategic view that you revise annually.**

   The view is wrong, in detail, every year. The discipline of maintaining it produces strategic clarity that no shorter horizon delivers.

2. **Identify the institutional capabilities that the longer view requires and that your firm currently lacks.**

   The capability gaps are the work of the decade.

3. **Engage younger colleagues seriously in the longer-view conversations.**

   They will, in 2040, be running the institutions you are currently leading. Their pattern recognition will be more relevant to that environment than yours.

CHAPTER FIFTY-FOUR

54

The Civilisational Question

What this technology is for, and what it is in service of.

—

"We are stardust, we are golden, we are billion-year-old carbon,
and we've got to get ourselves back to the garden."
— Joni Mitchell, "Woodstock," 1969

◆ ◆ ◆

I want to take one chapter, near the close of this book, to address a question that runs underneath every other question in it but that this book has not yet addressed directly. The question is what all of this technology is for. The question is contested in 2026 in ways that I think it will be more contested in 2030. I want to offer my honest answer, with the caveat that the question deserves answers from many voices and that mine is one among them.

The conventional answer is that technology is for productivity. We build technology to do more with less, to produce more value per unit of input, to free human time for more interesting work. The answer is not wrong. It is not, in my view, sufficient.

◆ ◆ ◆

I. Productivity is not the destination

The purely productivity-oriented answer leaves unstated what the additional value, the freed time, and the more interesting work are

for. If the answer is merely "more of itself" — more productivity, freeing more time, producing more value, freeing more time, in an endless recursion — the answer becomes hollow. Most thoughtful commentators on technology, across the centuries, have understood this. The productivity gains of previous technological revolutions were means to ends — the elimination of the diseases that killed children, the abolition of the most degrading forms of labour, the expansion of education and political participation, the lengthening of the span of healthy life, the connection of distant peoples, the deepening of human knowledge and human art. These were the ends. Productivity was the means.

The convergence decade will produce productivity gains of unprecedented scale. The question is what ends they will be deployed toward. The institutions of 2026 are, in many cases, organising as if productivity were itself the end. This is, I think, a category mistake. It is a mistake that will, if not corrected, leave the convergence decade poorer in human terms than it could have been.

◆ ◆ ◆

II. The civilisational ends worth pursuing

Three civilisational ends, in my honest view, deserve to organise how we deploy these technologies. The list is not exhaustive. It is a starting point.

The first is the substantial expansion of healthy human flourishing. Not just longer lives but better lives — lives in which the years of meaningful capability extend further, in which the diseases that diminished previous generations are progressively eliminated, in which the cognitive and emotional well-being of human beings is supported by institutions and practices designed for that purpose. The convergence technologies make this end newly achievable, in ways that previous technological eras could only approach.

Whether they are deployed toward it depends on what we choose to build.

The second is the substantial expansion of human creative and intellectual range. The augmentation of human capability with AI tools, the access to immersive learning environments, the opportunities for collaborative work across distances, the access to the deepest accumulated human knowledge — each of these expands the range of what individual humans can know, create, and contribute. The expansion is, in principle, available across the social spectrum more broadly than any previous educational expansion. Whether it is delivered in practice depends on the institutional architectures we build around the technology.

The third is the strengthening of human community at every scale. Family, neighbourhood, profession, faith, citizenship, civilisation, species. The technologies of this decade can be used to strengthen these communities or to dissolve them. They have, in many specific deployments, been used to dissolve them — the social platforms that fragmented attention and amplified division, the algorithmic media that polarised civic discourse, the convenient services that eroded face-to-face relationships. The same technologies, deployed differently, can strengthen these communities — supporting families in caring for elderly relatives at distance, allowing neighbourhoods to coordinate resources, allowing professions to share knowledge across geographies, supporting faith institutions in their continuing work, deepening civic participation, supporting human-scale relationships across cultural lines. Which deployment dominates depends on the institutions and the leaders who choose what to build.

◆ ◆ ◆

III. The Indian contribution

India, in my honest view, has a particular contribution to make to the civilisational question, drawing on the depth I described in

Chapter Twelve. The country has thought, longer than most, about what human flourishing actually consists of, what role technology should play in service of it, and what civilisational architectures support it. The traditions are diverse — Hindu, Buddhist, Jain, Sikh, Islamic, Christian, secular humanist, philosophical-rationalist — but the cumulative engagement with the question is unusually rich.

The country's contribution to the convergence decade is not principally, in my view, in the foundation models, the compute infrastructure, or the device platforms — though it should participate seriously in all of these. The country's distinctive contribution is in the *institutional and design imagination* that comes from a civilisation that has thought about human meaning as a first-class category. The institutions that India builds in this decade — the educational systems, the healthcare systems, the public-service systems, the workplaces, the cultural institutions, the faith institutions — can be expressions of this civilisational imagination, deployed with the convergence technologies as the means.

If India does this well, the country will, in 2040, be not only an economically significant participant in the convergence decade but a moral and design leader of it. The leaders who do this work will be doing so not as a national project narrowly defined but as a contribution to the broader human civilisation that India has long understood itself as participating in.

◆ ◆ ◆

IV. The leader's responsibility

The civilisational question, in the end, is not abstract. It is the question that each leader running each institution must, however imperfectly, answer for her institution. What is this institution for? What ends does its work serve? What contribution to human flourishing does it make? What community does it strengthen? What kind of world is it building?

The questions are uncomfortable in a commercial context that has, in recent decades, often answered them only in financial terms. They are also, I have come to believe, the questions that distinguish the institutions that will be looked back on in 2050 as having been worth building from those that will be remembered, if at all, as having been merely successful.

The leader's responsibility is to answer these questions with care, to build her institution accordingly, and to defend the answers against the perpetual pressure to substitute the means for the ends. The pressure is constant. The defense must be constant. The institutions whose leaders defend the ends produce the lasting work. The institutions whose leaders surrender to the means produce, eventually, work that no one is glad to have built.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Articulate, in writing, what your institution is for in civilisational terms — beyond the financial.**

   The articulation is the work. The articulation is also, in my experience, what most leaders have never done explicitly.

2. **Audit the institution's actual decisions over the past year against that articulation.**

   The audit will surface the gap between what the institution says it is for and what it actually does. The gap is the work.

3. **Build the institutional habits — the meetings, the metrics, the communications — that keep the civilisational ends visible alongside the operational metrics.**

   The institutions that lose sight of the ends drift. The institutions that maintain the visibility of the ends continue to do meaningful work over decades.

CHAPTER FIFTY-FIVE

55

A Builder's Generational Responsibility

Closing reflection: what I am asking of the reader, and of myself.

—

"Forgive yourself for not knowing what you
didn't know before you learned it."
— paraphrased from Maya Angelou, who said something close

◆ ◆ ◆

I want to close the body of this book with a personal reflection, because the book has been, in many places, more impersonal than I usually like to be. I have written, throughout, with the leader's framing — the leader's diagnosis, the leader's options, the leader's decisions. The framing has been, I think, useful for the reader who carries leadership responsibility. It has not, however, captured the personal dimension of what the convergence decade asks of the people who lead through it.

The personal dimension is, in my honest view, more demanding than the strategic dimension. The strategic dimension can be analysed, planned, and executed with discipline. The personal dimension cannot be analysed, planned, or executed in the same way. It must be lived. The leader who navigates this decade well

will, on the personal dimension, be required to do work that the strategic literature does not address.

◆ ◆ ◆

I. What this decade asks of the leader, personally

Three specific demands, in my observation, weigh most heavily on the leaders I work with.

The first is the demand of perpetual learning. The leader of 2026 must, simply, learn faster than she has at any previous point in her career. The technology shifts every quarter. The competitive landscape shifts every year. The customer expectations shift continuously. The leader who is not learning continuously is, within months, operating from obsolete pattern recognition. The discipline of continuous learning is, for many senior leaders, uncomfortable. The seniority that previously rewarded accumulated expertise now requires the renewal of expertise on a cadence that the senior leader has not previously experienced.

The second is the demand of equanimity in uncertainty. The leader of 2026 makes consequential decisions on incomplete information at a higher rate than she has at any previous point in her career. The information is inadequate not because the leader has been negligent but because the underlying questions are genuinely uncertain. The leader who waits for the information to be adequate is, in this decade, the leader who acts too late. The leader who acts at appropriate speed must, simultaneously, develop equanimity about the consequent fact that some of her decisions will be wrong. The equanimity is, for many leaders, harder than the decision-making.

The third is the demand of continued moral clarity. The leader of 2026 confronts choices that her predecessors did not face — what to deploy, what not to deploy, whom to displace, whom to elevate, whom to invest in, whom to abandon. The choices have moral

content, often substantial moral content, often in directions where the conventional commercial frameworks provide weak guidance. The leader who resolves these choices well does so by drawing on something more than commercial logic — a moral framework, a civilisational tradition, a personal sense of who she intends to be. The work of maintaining this moral clarity, against the constant pressure of commercial expediency, is, in my experience, the most demanding work of the decade.

◆ ◆ ◆

II. What I am asking of myself

I want to be transparent about what I am asking of myself in writing this book and in the work I do alongside it.

I am asking myself to engage with the technologies of this decade not as a vendor of them but as a builder for whom the deployment of the technologies is in service of something — the institutions I help to build, the leaders I help to develop, the constituencies I help to serve. The temptation to vend, in this decade, is constant and often substantial. I have tried, with imperfect success, to resist it.

I am asking myself to engage with the workforce transition responsibly. I am asking my own firm, Plexus Digitals, to live by the obligations I have described in Chapter Thirty. The cost of doing so is real. The discipline required is constant. I am not, I think, executing this perfectly. I am attempting it with seriousness.

I am asking myself to engage with the civilisational depth of India honestly. The temptation to invoke Indian civilisational themes superficially, as marketing or as identity-signalling, is constant. The harder work is to engage with the traditions deeply enough to draw on them as actual design source material rather than as decoration. I am, I think, beginning this work. I have not completed it. I expect to spend the rest of my career continuing it.

I am asking myself to maintain, in my own institution, AV Academy, an educational practice that is genuinely capable of preparing the next generation of Indian builders for the convergence decade. The work is hard. The learners deserve more than I have, on most days, time and capacity to deliver. The institution is improving. The improvement must continue.

◆ ◆ ◆

III. What I am asking of the reader

The book closes with what I am asking of you, the reader who has come this far.

I am asking you to take this decade seriously. The choices you make in 2026 and 2027 will compound across decades. The institutions you build, the technologies you deploy, the workforces you transform, the obligations you honour, the civilisational ends you serve — each will, by 2040, have produced consequences that are difficult to revise. The seriousness with which you approach the work now is not negotiable.

I am asking you to take responsibility for the people you lead and the constituencies you serve. The convergence technologies, deployed without that responsibility, will produce harms that the next generation will inherit. Deployed with that responsibility, they will produce flourishing that the next generation will be glad to inherit. The choice is not made in any single decision; it is made in the cumulative pattern of every decision.

I am asking you to take seriously the civilisational substrate from which you build. If you are an Indian builder, the substrate is unusually rich and unusually under-used. If you are a builder elsewhere, your own substrate is also rich, in ways that the conventional commercial frameworks under-credit. Read the tradition. Build from it. The convergence decade rewards firms

whose institutional imagination has roots deeper than the immediate market.

I am asking you to take care of yourself and the people close to you. The decade will be demanding. The demands compound. The leader who exhausts herself, or who exhausts the people she works alongside, does not produce the long-arc work the decade requires. Pace yourself. Pace your institution. The marathon is decade-long.

I am asking you, finally, to do the work you are uniquely positioned to do — and to leave the work that is not yours to leave. The convergence decade will produce many opportunities and many obligations. No leader can address all of them. The discipline is to identify the work that is genuinely yours, to commit to it deliberately, and to release the work that is genuinely not yours to others who are better positioned to do it.

If you do these things, in approximately the proportions and with approximately the seriousness this book has tried to convey, you will, in my honest view, build well in this decade. The institutions you build will be institutions worth working in. The customers you serve will be customers well-served. The workforces you employ will be workforces that thrive in the transition. The country and the world will be, in some small but real way, better for your work.

I do not expect you to do all of this perfectly. I have not, myself, done it perfectly. The work continues, in our institutions and in our individual lives, day by day, through the coming decade. The continuing is the point. The seriousness with which we continue is the work.

I wish you well. I look forward to seeing what you build. The world that's beginning is, in significant part, in your hands.

— Anil

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify, this week, one specific commitment you will make to lead through this decade differently than you have led through the previous one.**

   The commitment must be specific enough to be observable, and durable enough to compound.

2. **Share this commitment with one trusted colleague who will hold you to it.**

   The commitment honoured alone is honoured imperfectly. The commitment honoured in the company of someone who will hold you to it is more durable.

3. **Re-read this chapter in twelve months and assess honestly whether you have lived up to your commitment.**

   The annual re-reading is the discipline that, over a decade, accumulates into the kind of leadership the moment requires.

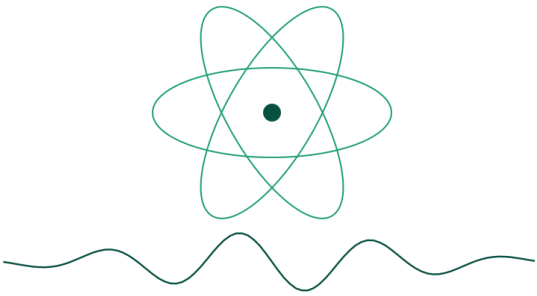
◆ ◆ ◆

The body of the book ends here. What follows — the epilogue, the appendices, the about-the-author note, the colophon — is reference material. The argument, such as it is, is contained in the preceding pages. I trust the reader will improve on it considerably.

PART THIRTEEN

XIII

The Quantum Horizon



|ψ

The deepest of the convergence forces examined more carefully. What quantum actually unlocks; what intelligence becomes when the throughput ceiling lifts; and what a quieter, less anxious human life might look like when abundance is engineered with care.

THE WORLD THAT'S BEGINNING

The deepest of the convergence forces examined more carefully. What quantum actually unlocks; what intelligence becomes when the throughput ceiling lifts; and what a quieter, less anxious human life might look like when abundance is engineered with care.

CHAPTER FIFTY-SIX

56

Beyond Classical Limits

What quantum computing actually
changes — and what it does not.

—

"Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical."
— Richard Feynman, 1981

◆ ◆ ◆

I want to begin this part of the book by stepping past the breathless coverage and the contemptuous dismissal that characterise most public conversation about quantum computing in 2026. Both postures are unhelpful. Both are wrong. Quantum computing is neither imminently going to break all cryptography next Tuesday nor a perpetual research curiosity that will remain commercially irrelevant. It is, in 2026, a serious computational substrate emerging from research into engineering, with a clear trajectory of capability over the next ten to fifteen years and a small but genuine set of applications already producing economic value. The honest discussion of what it means is more interesting than either of the standard narratives.

The honest discussion has to begin with what classical computing actually cannot do. Most of the public framing of "quantum advantage" is the inverse of the proper framing. We are not waiting

for quantum computers to become as good as classical computers. Classical computers have, over eighty years, become extraordinarily good at the problems they are good at. Quantum computers will never overtake classical computers on those problems. What they will do, increasingly through this decade, is solve a different class of problems — problems that are exponentially out of reach of classical machines, problems we have, in many cases, not seriously attempted because we did not have the substrate to try.

This is the chapter that maps that boundary carefully. The boundary is the most important thing for a leader to understand about quantum computing, because it tells her where to expect change and where to expect continuity. It tells her which of her industry's hard problems may, in this decade, become tractable, and which will continue to require the same patient improvement on classical hardware that they have always required. Without this map, every quantum forecast is either oversold or undersold. With it, the forecast becomes useful.

◆ ◆ ◆

I. The problems classical computing cannot reach

The core insight of quantum computing, stripped of physics, is that it processes information in a fundamentally different way from classical computing. A classical computer evaluates one configuration of its inputs at a time. A quantum computer, in a precise mathematical sense, evaluates many configurations simultaneously through quantum superposition. For most problems, this difference does not help — the answer still has to be measured, and measurement collapses the superposition. But for a small but consequential set of problem structures, the simultaneous evaluation produces results that classical computers, even at indefinite scale, cannot match in any practical timeframe.

The problem structures where quantum computing has known or strongly suspected exponential advantage fall into four broad families. The first is *simulation of quantum systems themselves* — the original Feynman insight. Chemistry, drug binding, materials science, condensed matter physics — all of these are governed by quantum mechanics at their fundamental level, and simulating them on classical computers becomes exponentially harder as the system size grows. A modestly-sized molecule already exceeds what the world's largest classical supercomputers can simulate accurately. Quantum computers will simulate them in principle, with engineering effort, in ways classical machines never will.

The second family is *certain optimisation problems* — particular structures of combinatorial optimisation where a quantum computer's ability to explore the solution space in parallel produces meaningful speedups. The class is narrower than is sometimes claimed, but it includes important practical problems: portfolio optimisation, certain classes of logistics routing, certain machine learning training regimes, and a growing set of constraint-satisfaction problems that arise in industrial scheduling and design.

The third family is *cryptanalysis* — the famous Shor's algorithm result that a sufficiently large quantum computer can factor large integers efficiently, breaking the RSA and elliptic-curve cryptography that protects most of the world's encrypted data. This is the result that drives the cryptography migration I described in Chapter Seven, and it is, of all the quantum applications, the one with the most immediate strategic implications. The migration has to happen, regardless of how slowly the quantum hardware capable of breaking modern keys actually arrives, because of the harvest-now-decrypt-later problem.

The fourth family is *certain machine learning and inference problems* — where the structure of the problem maps naturally onto

quantum algorithms. This is the most speculative of the four families, with the most active research and the least settled commercial application, but it is the one most likely to interact with the AI substrate I described earlier in the book. The combination of quantum computing and AI is, in the longest view, the deepest convergence of all.

These four families do not exhaust what quantum computers will be useful for, but they do cover most of what they will be uniquely useful for in the next decade. The leader looking for "the quantum opportunity in my industry" should be looking, first, at whether her industry has hard problems that fall into these families. If they do, the quantum trajectory matters strategically. If they do not, the quantum trajectory matters only through its effect on her cryptography.

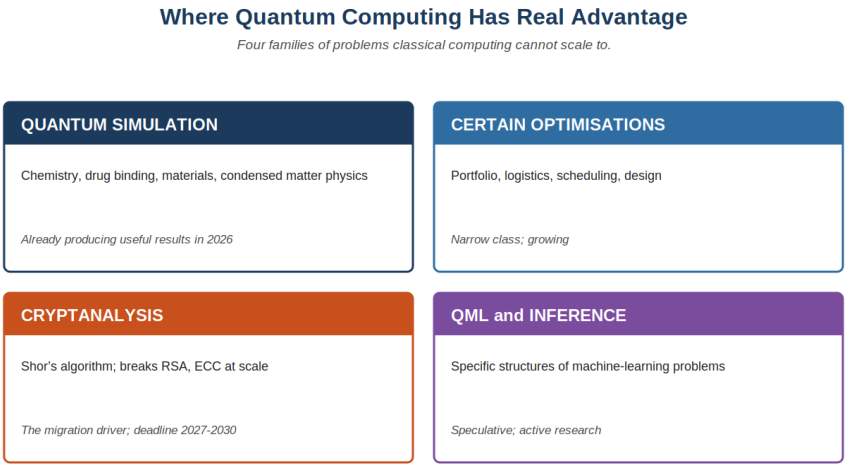


FIGURE 56.1 · Quantum is not a faster classical computer. It is a different computer for a few problem classes.

Figure 56.1

◆ ◆ ◆

II. Where the engineering actually is

The state of the engineering, in 2026, is more advanced than most outside observers realise and less advanced than the hype

suggests. The world has, in this year, several quantum systems with hundreds of physical qubits, a handful with around a thousand, and concrete roadmaps from multiple credible firms toward systems with hundreds of thousands of physical qubits within five to seven years. Physical qubits are not the same as logical qubits — a logical qubit, error-corrected and reliable for sustained computation, currently requires hundreds to thousands of physical qubits depending on the error-correction approach. The transition from "many physical qubits" to "many logical qubits" is the central engineering challenge of this decade, and the firms and laboratories that solve it will set the trajectory for the next.

By the most credible forecasts I have seen — produced by people without commercial interest in the answer — the world will, between 2029 and 2032, have its first quantum computers capable of running cryptographically relevant Shor's algorithm instances on real public-key targets. The first commercially useful chemistry simulations on real industrial molecules will, in many credible forecasts, arrive earlier — possibly as early as 2027 or 2028 — because the molecules of practical pharmaceutical interest are, surprisingly often, smaller than the integers protecting modern cryptography. The first useful optimisation results from quantum hardware on industrial problems are already appearing in 2026, although in narrow forms and under specific conditions that limit their commercial reach.

The Indian quantum landscape, in 2026, deserves a specific note. India's National Quantum Mission, launched in 2023, has begun to fund domestic research and infrastructure with significantly more seriousness than was the case a decade earlier. The country has a small but credible academic community in quantum information, several large companies that have begun internal quantum programmes, and one or two startup ecosystems beginning to form around quantum hardware and algorithms. The trajectory is positive. The country is, however, structurally behind the United

States, China, and parts of Europe in the absolute scale of investment and the depth of the research base. This is, in my view, an area where Indian policy and capital will need to be markedly more ambitious in the coming five years if the country is to participate in the quantum economy as a builder rather than only as a consumer.

◆ ◆ ◆

III. What this means for the leader

The practical implication for the leader is more subtle than "buy a quantum computer." Most enterprises will never operate quantum hardware directly. They will access quantum capability the same way they access other deeply specialised computational capability — through cloud APIs and through partnerships with the firms that operate the hardware. The strategic question, for most leaders, is therefore: *which of my industry's hard problems are quantum candidates, and what is my organisation doing to position itself to access quantum capability when it becomes available for those problems?*

The answer is not, in 2026, a quantum strategy. The answer is, in 2026, a quantum awareness — an organisational fluency that allows the firm to recognise quantum opportunities as they emerge, to prototype against current quantum hardware as a way of building that fluency, and to be ready to deploy production-grade quantum capability in the early 2030s when it begins to become genuinely commercial for non-research applications. The firms that build this fluency now will be ready. The firms that wait until quantum is "proven" will, characteristically, be five to seven years late, which in the convergence decade is roughly the difference between participation and irrelevance.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify the two or three hard problems in your industry that are quantum candidates — and brief your leadership team on them this quarter.**

   The exercise is the awareness. The awareness is the first move. Without it, every subsequent quantum decision is reactive.

2. **Form, within the year, a relationship with at least one credible quantum hardware provider or research group, with a small ongoing engagement.**

   The relationship, more than any specific deliverable from it, is the substrate from which production-grade quantum capability will become available to your firm when it matures.

3. **Treat the cryptographic migration of Chapter Seven as the urgent quantum-readiness move — and do not let it slip below board attention.**

   Of all the quantum implications, this is the one that requires action regardless of when commercial quantum advantage arrives. Treat it accordingly.

CHAPTER FIFTY-SEVEN

57

The Throughput Ceiling and How We Pass It

Why classical computing is hitting a wall — and what changes when the wall comes down.

—

"Computation is the new oil. Throughput is the refinery."
— overheard at a data centre conference, 2025

◆ ◆ ◆

The most important fact about computing in the 2020s, in my view, has been hiding in plain sight for nearly two decades, but it has only recently become uncomfortable enough to force a public conversation. The fact is this: classical computing has been hitting fundamental physical limits — call them throughput ceilings — that were comfortably distant in the 2000s and that have, since approximately 2015, become binding constraints on what any single classical computing system can do. Moore's Law, in the strict sense of doubling transistor density on a single chip, ended somewhere around 2018. Dennard scaling, the related observation that power per transistor would scale down with size, ended a decade earlier. The architectural innovations that have kept computing capability rising — multi-core processors, GPUs, specialised AI accelerators, distributed computing — have

produced enormous gains, but they are, increasingly, gains at the system level rather than at the individual transistor level. They are, in a deep sense, workarounds for the throughput ceiling rather than ways through it.

This chapter is about what changes when computing finds genuine ways through that ceiling, of which quantum is one and not the only one. The argument is that we are, in 2026, at the beginning of a multi-decade transition in computing substrate that is as consequential as the transitions from vacuum tubes to transistors and from transistors to integrated circuits. The transition is messier than its predecessors because it is happening across multiple substrates simultaneously — quantum, neuromorphic, photonic, biological — rather than through a single dominant new technology. The leaders who plan for the transition as if classical computing will continue to scale at its post-2018 rates are planning for a future that the physics itself does not permit.

◆ ◆ ◆

I. The throughput ceiling, briefly explained

The classical computer is, at its core, a system of tiny switches — transistors — that turn on and off billions of times per second. The number of transistors on a chip determines how much computation can happen per chip; the speed at which they switch determines how fast each transistor performs its operation; the energy required for each switching event determines how much heat must be dissipated. For decades, all three of these dimensions improved together: more transistors per chip, faster switching, less energy per switch. The combination produced exponential improvement in computational capability per dollar.

The improvements have stalled or reversed across all three dimensions. Transistors have, at the leading edge, become so small — three nanometres, two nanometres — that quantum mechanical effects begin to interfere with their operation. Switching speeds

have plateaued, because faster switching produces more heat than chips can dissipate. Energy per operation has stopped decreasing, for related thermodynamic reasons. The system-level workarounds — many cores, GPUs, specialised accelerators, distributed clusters — have continued to deliver enormous gains in total computational capability, but the cost of those gains has been rising. Each new generation of hardware is more expensive, more energy-hungry, and more dependent on extraordinarily concentrated supply chains.

The implication for AI is direct. The training of frontier AI models, in 2026, requires data-centre scale computing infrastructure that costs billions of dollars to build, consumes the electrical output of small countries, and depends on a handful of chip manufacturers concentrated in two or three jurisdictions. The trajectory of training cost is upward, not downward. The trajectory of inference cost is, similarly, downward but at a decreasing rate. The democratisation of AI that the public imagination expects from a Moore's-Law-like curve is bumping, increasingly, against a throughput ceiling that the physics of classical computing does not permit a way through.

◆ ◆ ◆

II. The non-classical substrates

Several non-classical computing substrates are, in 2026, in different stages of moving from research to practical deployment. Each addresses, in a different way, the throughput ceiling. None of them will replace classical computing for the workloads classical computing is well-suited to. Each is likely to dominate, eventually, a specific class of workload that classical computing cannot handle.

Quantum computing, examined in the previous chapter, addresses the ceiling by using a fundamentally different model of computation that scales differently for specific problem classes. The deployment trajectory is decade-scale. The application footprint is narrow but consequential.

Neuromorphic computing — substrates that mimic the architecture of biological brains — addresses the ceiling by changing the architecture of computation rather than the physics. Conventional von Neumann computers separate memory from processing, with most of the time and energy in modern computing spent moving data between them. Neuromorphic chips integrate memory and processing in the same substrate, dramatically reducing the energy-per-computation for certain classes of workload, particularly those involving pattern recognition and continuous-time inference. The first commercial deployments are appearing in 2026 in edge AI applications.

Photonic computing uses light rather than electrons to perform computational operations, exploiting the speed of light and the parallelism inherent in optical systems to achieve dramatic improvements in throughput per joule for certain classes of computation, particularly the matrix operations central to AI inference. The first photonic AI accelerators are commercially available in 2026, with a clearer roadmap than was the case three years earlier.

Biological computing — using living cells and biological systems to perform computation — remains, in 2026, the most speculative of the substrates, but the convergence with synthetic biology is accelerating, and the first credible applications in highly specialised domains are emerging from laboratories.

Each substrate is, in 2026, somewhere between research and early commercial deployment. None of them will, individually, replace classical computing. Together, in combination with continued classical improvements at the system level, they will produce a computing landscape in 2035 that looks substantially different from the one of 2025. The throughput ceiling will not have been removed; it will have been routed around in multiple directions.

◆ ◆ ◆

III. What lifting the ceiling enables

The practical implication of multiple substrates passing classical limits, in combination, is that several categories of computation that have been impractical for forty years will become practical in this decade and the next. Three are worth naming.

Whole-system simulation — of cells, of bodies, of climate, of cities — at fidelity that has never been possible. The first operationally useful "digital twins" of human physiology are likely in the late 2020s. The first operationally useful global climate models at the resolution needed for actionable adaptation planning are likely in the early 2030s. The first operationally useful complete-economy simulations capable of rigorously evaluating policy interventions before deployment are, on optimistic forecasts, also in the early 2030s.

Drug and material design at scale — the simulation-and-design of pharmaceuticals and novel materials at a pace that compresses the design cycle from years to months and, eventually, weeks. The combination of AI-guided design, quantum simulation of molecular behaviour, and high-throughput experimental validation has the potential to transform pharmaceutical development from an exercise dominated by attrition and luck to one dominated by deliberate design.

Cognitive substrates that continuously model the world — AI systems whose inference cost is low enough that they can be deployed at planetary scale, continuously updating models of every system they engage with. The shift, here, is from AI as a query-response system that one sends queries to and gets answers from, to AI as an ambient continuous-modelling layer that is, at all times, observing, predicting, and offering. The throughput ceiling has been the principal constraint on this vision. As the ceiling lifts, the vision becomes implementable.

These three are not, by themselves, the destination. They are the substrate on which the destinations of the next two chapters become possible — in particular, the chapters on intelligence unbottlenecked, on abundance, and on what abundance might mean for the texture of human life.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your firm's computational dependencies for their exposure to the throughput ceiling.**

   Most firms have not done this. Most are surprised by what they find. The audit reveals which of your strategic capabilities will scale and which will not on classical substrates alone.

2. **Build relationships with at least one provider in each of the non-classical substrates relevant to your industry.**

   Quantum, neuromorphic, photonic — pick the ones relevant. Build the relationship now, before the substrate becomes commercially crowded.

3. **Plan, explicitly, for the world in which inference is functionally unbottlenecked.**

   Most product strategies assume bounded AI throughput. The strategies that anticipate unbounded throughput will, by 2030, look very different from the strategies that did not.

CHAPTER FIFTY-EIGHT

58

Intelligence Unbottlenecked

What AI becomes when the inference
cost of a thought approaches zero.

—

"Whatever can be commoditised will be.
The question is what becomes valuable when it is."
— *paraphrased from many strategy textbooks, and consistently true*

◆ ◆ ◆

Run the thought experiment that opens this chapter alongside me. The cost of generating a single AI inference — a question answered, a decision suggested, a small piece of analysis performed — has fallen, between 2020 and 2026, by roughly four orders of magnitude. The same inference that cost a dollar in 2020 costs roughly a hundredth of a cent in 2026, and is dramatically more capable. The trajectory of the next decade, on the substrates I described in the previous chapter, is for inference cost to fall by another three to four orders of magnitude. By 2032 or thereabouts, generating a thoughtful AI response of meaningful sophistication will cost less than a millionth of what it cost in 2020, and the AI will be substantially better than the AI of 2026.

This is the throughput unbottleneckting that I want to think about carefully in this chapter. Most of the strategic discussion of AI in 2026 implicitly assumes that AI inference is and will remain a

bounded resource — that firms will deploy AI selectively, on the highest-value applications, with conscious attention to the cost of each call. The bounded assumption shapes every strategic conclusion. It determines which applications get AI and which do not. It determines the architecture of products. It determines the economics of competition.

The bounded assumption will, in my honest view, be wrong by 2030. AI inference will, by then, be so cheap that the bounded constraint will functionally disappear for most applications. The strategic implications of this are deeper than is widely appreciated. They go beyond "AI is everywhere" to "AI is in the substrate of everything, continuously, ambiently, at no marginal cost worth tracking." This chapter is about what changes when that becomes true.

◆ ◆ ◆

I. The end of the rationing problem

For most of computing's history, computational resources have been rationed. The mainframe in the 1970s was rationed by time slots. The personal computer in the 1990s was rationed by the speed of the user's machine. The cloud in the 2010s was rationed by the cost per CPU-hour. The AI inference of the 2020s has been rationed, internally inside firms, by the budget for API calls. In each era, intelligent system design has involved choices about where to spend the rationed resource.

Unbottlenecked inference is the end of the rationing problem for AI. The question shifts from "where can we afford to deploy AI?" to "where would it make sense to deploy AI, given that the cost is functionally zero?" The answer to the second question is everywhere — every interaction, every decision, every interface, every quiet moment in which a useful suggestion would be possible.

The shift is more than economic. It is architectural. Products that were designed with bounded AI in mind have an explicit AI surface — a chatbot, an assistant button, a "powered by AI" feature. Products designed with unbottlenecked AI will not have an AI surface, because AI will be the substrate, present everywhere, invisible. The chatbot is a transitional architecture, much as the personal computer's command-line was a transitional architecture between the punched-card era and the graphical interface. The mature architecture, in five to ten years, is one in which AI is the medium rather than a feature within it.

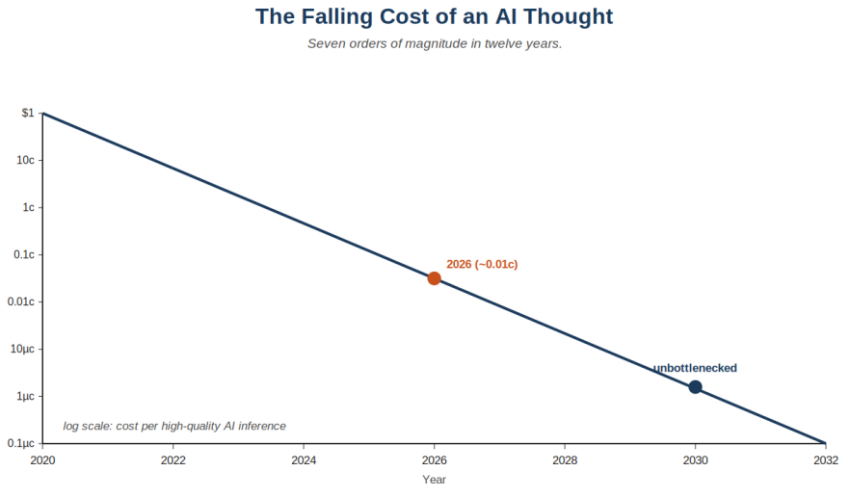


FIGURE 58.1 · Cost per inference, log scale. When something becomes a million times cheaper, its application changes.

Figure 58.1

◆ ◆ ◆

II. Continuous rather than queried

The corresponding architectural shift, on the consumption side, is from queried AI to continuous AI. In 2026, a user typically interacts with AI by issuing a query: she asks a question, the AI responds, she reads the response, the interaction ends. The query-response architecture is direct heir to the Google search box and to the chatbot interfaces of the early 2020s. It is, in 2026, deeply familiar.

In the unbottlenecked era, query-response will increasingly give way to continuous AI presence. The user does not query. The AI is, simply, present — observing the situation, modelling the user's likely intent, offering when offering would be useful, withholding when withholding would be more useful, learning continuously from the patterns of when offers are accepted and when they are ignored. The interface is conversation, but the conversation is mostly silent. The AI is the calm presence in the room, intervening rarely and well.

This architectural shift is, in my view, the most important strategic implication of unbottlenecked inference. The firms that build for query-response will be building for an architecture that will look, by 2032, the way the command-line interface looked by 2002 — still functional in narrow contexts, but no longer the default. The firms that build, deliberately, for continuous AI presence will be building for the architecture of the next decade.

◆ ◆ ◆

III. The personal AI

Among the consequences of unbottlenecked inference is the practical viability of the *personal AI* — an AI that knows the individual user deeply, persistently, over years; that is continuously available; that learns the user's preferences, voice, history, goals; that acts on her behalf in the digital and physical world; that mediates her interaction with every other digital service. The personal AI was technically possible in 2024. It was economically marginal until inference became cheap enough for continuous operation. By 2030, in my forecast, it will be economically dominant — most affluent users in the developed and middle-income world will have a personal AI of considerable sophistication, and the personal AI will become the primary mediating layer between the user and the rest of the digital world.

The strategic implications of the personal AI are profound. The personal AI is, in the user's experience, an extension of her own cognition — a reliable second mind that handles the cognitive overhead of modern life. It is, in the strategic landscape of consumer-facing firms, the gatekeeper to the user's attention. The firm that previously reached the user through her smartphone, her email, her notifications, increasingly reaches her — if it reaches her at all — through her personal AI's filtering and curation. The relationship that matters, by 2030, is the firm's relationship with the user's personal AI as much as with the user herself.

This is a new strategic surface. It has no clear historical precedent. The closest analogy is the role that retail buyers played in twentieth-century consumer goods — the gatekeepers who decided which products would reach the shelf — but personal AIs are, in the limit, more demanding than any human buyer ever was. They will be auditing every claim, comparing every alternative, optimising relentlessly for the user's interests. The firms that earn their place in the personal AI's recommendations will be the firms that have earned it on the merits, audited and demonstrable. The firms that bought their place in previous channels through marketing spend will find that the personal AI cannot be bought in the same way.

◆ ◆ ◆

IV. The implications for human cognition

The deepest implication of unbottlenecked AI, in my view, is for the structure of human cognition itself. The human in 2026 lives, cognitively, in a state of continuous low-grade overload. She has too many open browser tabs, too many emails, too many messages, too many decisions about which information to attend to. The cognitive load is real, sustained, and corrosive. It produces the fatigue, the anxiety, the inability to think deeply that characterise so much of professional life in this era.

Unbottlenecked AI, deployed thoughtfully, has the potential to lift this load. The personal AI handles the noise. The user attends only to what genuinely requires her attention. The interruptions become fewer, the decisions become better-supported, the information environment becomes adapted to her rather than the reverse. The human, in this configuration, is freed to think — actually think — about the things humans are uniquely good at thinking about.

I am, here, describing a possibility rather than a certainty. The same unbottlenecked AI deployed without thought will produce the opposite outcome — a torrent of AI-generated content, AI-generated noise, AI-mediated manipulation, that drives the cognitive load even higher than the human-generated noise of the 2020s. Which of these futures we get is not determined by the technology. It is determined by the choices the builders make about what to build and what not to build with the unbottlenecked substrate.

This is the chapter at which the question of the book — what do we build for the world that's beginning? — becomes most consequential. The unbottlenecked era amplifies whatever it is given. Given thoughtful intentions, it produces the calm, well-supported cognitive environment I sketched a paragraph ago. Given the incentives of attention extraction that have shaped much of the previous era, it produces the opposite. The choice is the builders'. It is the choice this book is, in the end, about.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Re-baseline your product strategy on the assumption that AI inference is unbounded by 2030.**

   The strategies that emerge from this assumption look meaningfully different from those that assume bounded AI.

2. **Design for continuous AI presence rather than query-response surfaces.**

   The transition is happening. The firms that begin the transition early will build the dominant architectures of 2030.

3. **Examine, in your firm's product decisions, whether the AI you are deploying lifts cognitive load from the user or adds to it.**

   The honest answer determines which side of history your products are on.

CHAPTER FIFTY-NINE

59

The Abundance Frontier

Drug discovery, materials, food, energy: what changes when the design cycle compresses by an order of magnitude.

—

"Abundance is not a state of having more.
It is a state of designing better."
— paraphrased from Buckminster Fuller,
who was not subtle on the point

◆ ◆ ◆

The convergence of unbottlenecked computation, AI-guided design, and quantum simulation produces, in several specific domains, what I will call the abundance frontier. The frontier is the set of categories where the cost of designing genuinely new things — molecules, materials, organisms, processes — is poised to fall by an order of magnitude or more in the coming decade. The categories where this happens are not the whole of human economic activity. They are, however, the categories that affect the basic substrates of human life: medicine, the materials we live with, the food we eat, the energy that powers everything. When these substrates become abundant, much else becomes possible that previously was not.

This chapter is about that frontier. I will treat four categories — drug discovery, materials, food and agriculture, and energy — at

the level of detail a leader needs to decide whether to participate. The treatment is honest about timing and uncertainty. Some of these inflections are likely in the next five years. Some are likely in the ten- to fifteen-year horizon. None of them is certain. All of them are, in my view, more likely than not.

◆ ◆ ◆

I. Drug discovery, finally rational

The pharmaceutical industry has, since its inception, been an industry of rational improvisation. The combination of biological complexity, regulatory caution, and limited computational power has meant that drug discovery — for all its scientific sophistication — has mostly proceeded by educated guesses, expensive trials, and frequent failure. Roughly nine in ten drugs that enter human trials never reach market. The cost per drug developed has, in inflation-adjusted terms, increased by an order of magnitude over the past forty years. The "Eroom's Law" (Moore's Law spelled backward), describing the steady deterioration of pharmaceutical productivity, has been the embarrassing counterpoint to the celebrated improvements in semiconductor productivity over the same period.

The convergence is, finally, beginning to reverse this. AI-guided molecular design has, between 2020 and 2026, become genuinely useful for narrowing the search space of candidate molecules from the impossibly large to the manageably small. Quantum simulation of the binding of candidate molecules to target proteins is, in 2026, beginning to produce results that classical simulation cannot match. Automated experimental synthesis and high-throughput biological screening are compressing the experimental cycle from months to weeks. The combination has begun to produce, in a small but growing number of cases, drug candidates that have moved from concept to clinical trial in a fraction of the time that the previous generation required.

By 2032, on the optimistic forecasts I take seriously, the design phase of drug discovery — the period from target identification to first in-human trial — will be compressed by roughly an order of magnitude relative to 2020 standards. The clinical phase will not compress by the same amount, because the biology of human trials has its own irreducible time constants. But the overall cycle from disease to approved therapy will compress meaningfully, the cost per approved drug will fall, and the categories of disease that have been economically infeasible to address — rare diseases, diseases of poverty, drug-resistant infections — will become commercially viable in numbers that were not previously possible.

The implications for India deserve specific attention. India has been, for decades, a global power in generic pharmaceuticals — a producer of molecules originally discovered elsewhere, supplying much of the developing world with affordable medicines. The shift to designed therapeutics, at compressed cycles and lower costs, is an opportunity for the country to participate in original drug discovery at a scale it has not previously achieved. The combination of Indian biological talent, computational infrastructure, and patient cohorts of unmatched diversity creates conditions for Indian-discovered, Indian-developed medicines to enter the global pharmacopoeia in numbers in the coming decade. The opportunity will not be realised without deliberate policy support and serious capital, but it is a real opportunity, and it sits squarely in the abundance frontier.

◆ ◆ ◆

II. Materials, designed rather than discovered

The world of materials — the physical substances of which everything is made, from buildings to batteries to clothing — has historically been a domain of empirical discovery rather than rational design. The great materials of the twentieth century — stainless steel, aluminium alloys, concrete, plastics — were discovered through painstaking trial and error over decades. The

notion of designing a material from scratch, with specified properties, has been an aspiration of materials science for a century, and a frustrating one, because the relationship between molecular structure and macroscopic properties has been too complex to compute reliably.

The convergence is changing this. AI models trained on the vast accumulated literature of materials science can, in 2026, predict properties of candidate materials with usable accuracy. Quantum simulation can, for an increasing range of compositions, compute the underlying structure-property relationships from first principles. High-throughput experimental synthesis can fabricate and test candidate materials at a pace that was not previously possible. The combination is, between 2024 and 2026, producing designed materials with properties tailored to specific applications — battery cathodes with higher energy density, building materials with lower embodied carbon, semiconductors with novel electrical characteristics — at a pace that had been unthinkable as recently as 2020.

The implications cascade across virtually every physical industry. Battery storage, the foundational technology of the energy transition, depends on better materials more than it depends on any other single advance. Aviation, shipping, and heavy industry all depend on materials that can perform in conditions where current materials fail. Construction depends on materials that can deliver the same structural performance with a fraction of the embodied carbon. The materials revolution is the slow-moving but genuinely transformative companion to the AI and quantum revolutions, and the firms that participate in it will participate in much of what follows.

◆ ◆ ◆

III. Food and agriculture

The food system that feeds the world's eight billion people in 2026 is a marvel of human accomplishment and a thicket of

vulnerabilities. It is dependent on a small number of crops, increasing climate stress, fragile supply chains, and intensive use of inputs whose long-term sustainability is contested. The convergence offers, in this system, several pathways to greater resilience and abundance.

Designed crops — produced through AI-guided breeding, gene editing, and increasingly synthetic biology — will, by 2030, deliver staple crops with higher yields, lower water requirements, greater pest resistance, and improved nutritional profiles. The first commercial deployments are already in the field in 2026 in several markets, including India.

Cellular agriculture and fermentation — the production of meat, dairy, and other animal products from cell cultures and engineered microorganisms rather than from animals — will, by 2030, be producing a meaningful share of certain food categories at price points competitive with conventional production. The early commercial products of 2024 and 2025 have established the scientific feasibility; the price-and-scale path to dominance is the work of the next decade.

Precision agriculture — the use of continuous sensing, AI inference, and automated equipment to manage farming at the level of individual plants rather than fields — will, by 2030, have transformed productivity per hectare, per litre of water, and per kilogram of input across most commercial farming. The Indian smallholder context is harder than the industrial-farming context where these techniques first emerged, but the work of adaptation is well underway.

The combination of these is the agricultural component of the abundance frontier. By 2035, on the trajectory I take seriously, the world will be producing meaningfully more food, on less land, with less water, with lower environmental footprint, than it does today. The political and supply-chain questions of how the abundance is distributed are, of course, separate from the technical

question of whether it is produced. The technical question, after a long period of pessimism, is now answerable in the affirmative.

◆ ◆ ◆

IV. Energy abundance

The convergence component most directly tied to abundance, in my view, is the energy transition. Cheap, clean, abundant energy is the substrate on which most other forms of abundance rest. A world with abundant energy can do desalination at scale, vertical farming at scale, computation at scale, transportation at scale, with environmental footprints that constrained-energy economies cannot afford. A world with scarce or expensive energy cannot.

The trajectory of energy in the next decade is, on my reading, more positive than is widely appreciated. Solar and wind, in most of the world's geographies, are now the cheapest forms of new generation. Battery storage costs have fallen by an order of magnitude in twelve years, with credible roadmaps for further decline. Green hydrogen, at scale, is approaching commercial viability for the categories where electrification is impractical. Nuclear — particularly small modular reactors and the early advances toward fusion — is, in 2026, moving from the political wilderness back into the strategic conversation. The combination of these, with continued investment, points to an energy landscape in 2035 that could deliver abundance of clean energy at prices that, in real terms, would be the lowest in human history.

The implications, if we get them, are civilisational. Cheap clean energy resolves the most binding constraint on most forms of human flourishing. It enables the abundance frontiers in pharmaceuticals, materials, and food. It permits the computing infrastructure of unbottlenecked AI to be deployed at planetary scale without environmental devastation. It enables the climate adaptation that the world will need regardless of how rapidly we decarbonise. It enables, in poorer parts of the world, the shift from

energy poverty to energy adequacy that underwrites every other improvement in living standards.

This is the most important, and most contingent, frontier of all. The technical trajectory is positive. The implementation trajectory depends on the deliberate choices of governments, capital, and industry over the next decade. The leaders who participate seriously in the energy transition are participating in the substrate of every other abundance described in this chapter.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Identify which of the four abundance frontiers — pharmaceuticals, materials, food, energy — most directly affects your industry, and engage seriously with the trajectory of that frontier.**

   The engagement does not require participation in the frontier. It does require informed positioning relative to it.

2. **For Indian leaders specifically, contribute to the policy and capital decisions that determine whether India participates in the abundance frontiers as a builder or as a consumer.**

   The decisions of the next two years are unusually consequential. Distributed indifference produces a bad outcome.

3. **Refuse, in your firm's strategy, the implicit assumption that scarcity is permanent.**

   Most strategy is built on scarcity assumptions. The strategies most likely to be wrong, by 2035, are the ones that assumed abundance would not arrive.

CHAPTER SIXTY

60

Climate, Repair, and the Quantum Toolkit

The hardest problem of the century,
and the substrate that gives us a real chance.

—

"We did not inherit the earth from our ancestors;
we borrowed it from our children."
— *attributed (perhaps apocryphally) to Native American wisdom;*
the principle is older than the attribution

◆ ◆ ◆

I want to address, in this chapter, the question that hangs over every other question in this book. The question of climate. The forces that this book has described — AI, spatial computing, quantum, programmable biology — arrive into a world that is also, simultaneously, in the early stages of the most consequential environmental transition in human history. The transition will, on any honest reading, dominate the political and capital flows of the next half-century. The technologies the book has discussed are central to whether the transition is navigated well or badly.

I have, until this chapter, treated climate as one of the three macro tectonics of Chapter Two and otherwise as a backdrop. The treatment has been deliberately spare, because most of what passes

for climate strategy in business literature in 2026 is either platitudinous or pessimistic, and I have not wanted to add to either pile. In this chapter, I want to be more direct about what the convergence technologies actually offer for climate, and what they ask of the leaders who deploy them.

◆ ◆ ◆

I. Modelling the system honestly

The climate system is the most complex coupled physical system humans have ever attempted to understand. Its dynamics span thirteen orders of magnitude in time and similar ranges in space. Its components interact non-linearly. Its responses to forcing have hysteresis, threshold effects, and tipping points whose locations are, in many cases, known only approximately. The limitations on classical climate modelling have, until recently, been substantial — not because the science is inadequate, but because the computational substrate has been inadequate to represent the complexity at the resolution we would want for actionable planning.

The convergence is, finally, lifting these limitations. AI-augmented climate models, deploying machine learning to capture sub-grid-scale processes that previous models could only approximate, are producing forecasts with accuracy and resolution that were not previously available. Quantum simulation is beginning to address specific aspects of atmospheric chemistry that classical computers struggle with. The compute infrastructure required to run high-resolution global models is, with the throughput improvements I described in Chapter Fifty-Seven, becoming accessible to research groups beyond the few national laboratories that historically dominated the field.

The result is a climate science that, by 2030, will have honest, actionable forecasts at the level of resolution that decision-makers actually need. Not "the global mean temperature will rise by some

amount over some period" — the kind of forecast that has dominated public discussion for thirty years — but "this specific region will see this specific change in monsoon timing, in heatwave frequency, in river flow, in agricultural viability, with this specific confidence interval, on this specific timeline." The shift from coarse global forecasts to fine-grained regional forecasts is the precondition for serious adaptation planning. We are now equipped, technically, to do it.

◆ ◆ ◆

II. Adaptation, increasingly the binding question

The public conversation about climate has, for thirty years, been dominated by the question of mitigation — how to reduce greenhouse gas emissions so as to limit the warming. The conversation continues to matter, and the trajectory of mitigation has, in 2026, finally turned in the direction of the necessary scale of effort. But the scientific consensus has shifted, in the past decade, to recognise that significant climate change is now locked in regardless of mitigation, and that adaptation — the redesign of human systems to cope with the climate that is coming — is becoming increasingly, in many parts of the world already, the binding question.

The convergence technologies are central to adaptation. The high-resolution forecasts I described above are precondition for serious infrastructure planning. AI-driven optimisation is essential for the redesign of water systems, agricultural systems, urban heat mitigation systems, and disaster response systems for a climate that is no longer stable. Quantum simulation is contributing to the design of materials and processes — concretes that absorb carbon, refrigerants that work without the climate damage of older substances, agricultural systems calibrated to the new climates — that adaptation requires.

For India, this is particularly consequential. India is, on virtually every assessment, among the most climate-vulnerable major economies in the world. The combination of dense population, agricultural dependency, monsoon-driven water systems, and high heat exposure means that the cost of climate change to India, if not actively managed, will be civilisationally severe. The cost of climate change to India, if actively managed with the technologies of the convergence decade, can be very substantially reduced. The decisions of the next decade, on adaptation infrastructure, on water systems, on resilient agriculture, on heat-adaptive cities, are among the most consequential decisions any government in the world will make in this period.

◆ ◆ ◆

III. Removal, finally credible

The third leg of the climate response, after mitigation and adaptation, is removal — the active extraction of carbon already in the atmosphere through engineered systems, biological systems, or some combination. Removal has, until recently, been a marginal note in climate strategy because the engineered approaches were too expensive and the biological approaches were too small in scale to matter. The convergence is changing this.

Several pathways have, in 2026, begun to look credible. Direct air capture — engineered plants that pull carbon dioxide out of the atmosphere and store it permanently — has, with materials science advances, fallen in cost from a thousand dollars per tonne to several hundred, with credible roadmaps to a hundred dollars per tonne by 2032. At that cost, removal at scale becomes economically feasible alongside mitigation. Biological pathways — through agriculture, soils, oceans, and engineered microorganisms — are simultaneously being optimised through AI and synthetic biology in ways that may, by 2030, produce removal in the gigatonnes per year required for meaningful effect on atmospheric concentrations.

The convergence is what makes removal possible. Without cheap energy, direct air capture is not viable. Without AI-optimised process design, the engineering does not work. Without quantum simulation of the relevant chemistry, the materials needed for cost-effective capture are not available. Without continuous monitoring infrastructure, the biological pathways cannot be verified at scale. Removal, in my view, is the single most under-appreciated component of the long-term climate solution, and the leaders who participate seriously in it now are participating in what may be, in the long view, the single most consequential industry of the late twenty-first century.

◆ ◆ ◆

IV. The honest uncertainty

I want to be honest about the uncertainty. The trajectory I have described in this chapter is not guaranteed. The technical advances may slow. The political will may falter. The capital may be misallocated. The deployment may be slower than the warming. The combination produces, at the pessimistic end of the range, a climate outcome that is bad in human terms — many lives shortened, many regions made unlivable, many economies disrupted — and that the convergence technologies, deployed late, can only partially repair.

I write the chapter as I do because the optimistic end of the range is also genuinely possible, and the choice between them is not yet made. The choice depends on the deliberate decisions of governments, capital, and industry over the next ten to fifteen years. The leaders reading this book are part of that choice. The chapter is, in the end, an argument that the choice is worth making well.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your firm's climate exposure honestly, on the high-resolution forecasts that are now becoming available rather than on the coarse forecasts that have dominated previous strategy work.**

   The honest exposure is, for most firms, larger than the strategy decks have admitted.

2. **Identify, in your firm's portfolio, the products and operations that contribute to mitigation, to adaptation, or to removal — and prioritise the ones that contribute to multiple.**

   The strongest strategic positions in the climate decade are the ones with multiple contributions.

3. **Resist, in your firm's strategy, the temptation to treat climate as a public-relations or compliance question rather than as a strategic question.**

   The firms that have made this mistake in the past decade are the firms whose climate strategies are now most embarrassing. The mistake will be more expensive in the next decade than in the past one.

CHAPTER SIXTY-ONE

61

Cryptography, Trust, and the Privacy Renaissance

What it actually means when we can prove things mathematically.

—

"Trust, but verify."

— Russian proverb that became an American slogan; the verification half is the part the convergence makes new

◆ ◆ ◆

The most under-discussed implication of the quantum era, in my view, is not the threat to existing cryptography that I treated in Chapter Seven. It is the new cryptography that the era enables. The post-quantum cryptography that protects against quantum attack is one piece of the picture. The other piece, equally important and far less discussed, is the family of advanced cryptographic techniques — zero-knowledge proofs, fully homomorphic encryption, secure multi-party computation, threshold cryptography — that the convergence finally makes practical at scale and that, in combination, enable a category of digital services that has been mathematically possible for decades but computationally impractical.

I want to spend this chapter on what those techniques actually enable, because the leader who understands them has access to

product categories that competitors do not. The categories sit at the intersection of the era's three currencies: attention, time, and trust. They are, in particular, the most credible answer the convergence offers to the privacy crisis that has dominated the digital environment for two decades.

◆ ◆ ◆

I. Proving without revealing

Zero-knowledge proofs are mathematical constructions that allow one party to prove to another that a statement is true without revealing any information beyond the truth of the statement itself. The construction has been known to cryptographers since the 1980s. It has been computationally impractical for most useful applications until very recently, because the proofs were too slow to compute and too large to transmit. The combination of mathematical advances and improved computational substrates has, in the past five years, brought zero-knowledge proofs to practical viability for a growing range of applications.

The implications, when one thinks about them clearly, are extraordinary. Imagine a world in which the bank can verify that the customer has sufficient income to qualify for a loan without seeing the customer's actual income. The hospital can verify that the patient is over eighteen without seeing the date of birth. The voter can prove she voted, exactly once, without anyone knowing how. The property buyer can prove she owns the funds without revealing where they sit. The insurer can verify the property is in good condition without inspecting it. The auditor can verify the firm's financial statements are accurate without seeing the underlying transactions.

Each of these is, mathematically, achievable with zero-knowledge proofs. Most of them are, in 2026, on the cusp of being commercially achievable as well. The first production deployments are emerging in finance, in government services, and in identity systems. By 2030, zero-knowledge constructions will, in my

forecast, be the substrate of most serious identity and verification services in the digital environment. The privacy implications are, taken seriously, civilisational.

◆ ◆ ◆

II. Computing on encrypted data

Fully homomorphic encryption is the second of the techniques. It allows computation to be performed on encrypted data without decrypting the data first. The result, when decrypted, is the result that would have been obtained by performing the same computation on the unencrypted data. The construction has been known since 2009 and has, until recently, been so computationally expensive as to be impractical for most uses. The convergence is, again, changing this.

The implication is that data can be entrusted to a computational service without entrusting the data to the service. The medical AI can analyse the patient's encrypted records without ever knowing what they contain. The financial AI can detect fraud across encrypted transactions without seeing them. The advertising platform — to take a less attractive example — could, at least in principle, do its targeting on encrypted user data, without ever seeing the underlying behaviour. The principle is the separation of computation from data ownership, which has been impossible to enforce technologically and which, in the homomorphic era, becomes possible.

The deployment trajectory is slower than for zero-knowledge proofs, because the overhead remains substantial. By 2030, in my forecast, homomorphic encryption will be in production deployment for high-value, low-volume applications — financial analysis, medical research, certain government services — and will be advancing toward viability for higher-volume applications. By

2035, it will be a mainstream component of the privacy-preserving infrastructure of the digital economy.

◆ ◆ ◆

III. Computing together without trusting each other

Secure multi-party computation is the third of the techniques. It allows multiple parties, each holding private data, to compute a function over their combined data without any of them revealing their data to the others. The classic example is computing the average salary of a group of people, with the result accurate, without any of the people learning any other person's salary.

The implications, again, are wide-ranging. Hospitals can collaborate on research over their pooled patient data without any hospital seeing another hospital's data. Banks can detect fraud rings that span multiple institutions without sharing customer records. Companies in the same industry can compute industry benchmarks without revealing their individual numbers. Government agencies can investigate criminal patterns across their jurisdictions without combining their databases. Each of these applications has been, mathematically, achievable for decades and, computationally, impractical until recently. Each is, in 2026, becoming feasible.

The combination of the three techniques is more powerful than any of them alone. A future digital service can, with deliberate engineering, prove things to its users without seeing user data; can compute on user data without decrypting it; can collaborate with other services without exposing what its users have entrusted to it. This is a radically different posture toward user data than the data-extractive posture that has dominated the consumer internet for two decades. It is, I would argue, the only posture toward user data that is morally and politically sustainable in the long term.

◆ ◆ ◆

IV. The privacy renaissance

The privacy environment of the consumer internet, between roughly 2005 and 2025, has been catastrophic. The implicit deal struck by the major platforms — free services in exchange for behavioural surveillance — has produced a digital environment in which most individuals have given up most of what previous generations would have called privacy, with only intermittent awareness of how much they have given up and only occasional regret about it. The harms have ranged from the moderately consequential (targeted manipulation, addictive product design) to the gravely consequential (election interference, persecution of dissidents in authoritarian regimes, manipulation of vulnerable individuals).

The cryptographic techniques I have described make possible, finally, an alternative architecture. The architecture is one in which digital services can be useful without harvesting data; can be personalised without surveilling; can be commercially viable without selling user attention. The architecture is technologically possible. The question is whether it will be commercially adopted, and whether the regulatory environment will require it to be.

I am, on this question, more optimistic than the trajectory of the past five years would suggest. Several forces are converging toward the alternative architecture. Regulatory frameworks in Europe, India, and increasingly in the United States are moving toward stronger data-protection requirements that the cryptographic techniques can satisfy. User consciousness about data exploitation has, by my reading, reached a tipping point — the willingness to surrender data has begun to decline, particularly among younger and more sophisticated users. Competitive dynamics among technology firms are creating, for the first time, market opportunity for firms that build trust-first rather than extraction-first architectures. The combination, given another five to ten years of work, has the potential to produce a digital environment in 2035 that looks very different from the one of 2020

— better, in concrete and verifiable ways, than the bleak environment we have come to accept as inevitable.

This is what I mean by the privacy renaissance. It is not a return to the analog past. It is a forward movement to a digital environment in which privacy is mathematically guaranteed by the technical architecture rather than promised by the operating policies of the firms that hold the data. The renaissance is, in 2026, in its early stages. It deserves the deliberate engineering effort and the deliberate commercial commitment of the leaders reading this book.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Engage seriously with the advanced cryptography roadmap for your firm's products this year.**

   Most firms have not. The first firms to incorporate zero-knowledge proofs and homomorphic encryption into consumer-facing products will hold a substantial trust advantage by 2030.

2. **Audit your firm's data extraction practices honestly, with an eye to how they will look in 2030.**

   Practices that look defensible in 2026, in many cases, will not look defensible by 2030. The audit is the warning.

3. **Take a public position on the privacy-renaissance architecture, even if your firm cannot yet implement it fully.**

   The position is the commitment. The commitment shapes what the firm builds. Silence, on this question, is also a position; it is the position that the extractive architecture of the previous era is acceptable to your leadership.

CHAPTER SIXTY-TWO

62

The Stress-Free Life

What abundance actually buys, and what it does not.

—

"Wherever you go, there you are."
— Jon Kabat-Zinn, with the dry observation that no technology has ever changed this

◆ ◆ ◆

I want to turn, in this chapter, from the technical implications of the convergence to a more personal question. The question I am asked most often, in talks and in private conversations, is some variant of: *will any of this make life better?* The asker is rarely interested in the abstract argument that productivity will rise. She wants to know, concretely, whether her life — her ordinary, lived, day-to-day life as a parent, a worker, a citizen, a person — will be less anxious, less hurried, less crushing in 2035 than it is now. She is asking whether the convergence is an answer to the conditions of human life that have been getting demonstrably worse, by most measures of wellbeing, for two decades. She is asking whether the abundance will be felt, in the texture of life, or only in the GDP statistics.

I have come, slowly, to the answer that this chapter argues. The answer is: it can be, but only if the builders deliberately design for it. The technologies of the convergence are not, on their own,

oriented toward human flourishing. They are oriented toward whatever incentives the systems that deploy them happen to align with. If those systems are aligned with attention extraction, the technologies will produce a more anxious, more hurried, more crushing life. If they are aligned with human flourishing, they can produce something genuinely better than the human condition has historically delivered. The choice is, again, the builders'. This is the chapter that argues for the second alternative as deliberately as I can.

◆ ◆ ◆

I. What is actually causing the cognitive load

The lived experience of the modern professional, by the late 2020s, is one of more or less continuous cognitive overload. The phenomenon has been measured and described extensively. The averages are bleak. Working professionals report less time for deep thought than at any point in the modern measurement of work. They report higher rates of anxiety, lower rates of sustained focus, more frequent burnout, more difficulty with sleep, more difficulty with the kinds of leisure that used to restore them. The numbers vary across studies and populations, but the direction is consistent and consistent across professions. The condition is real.

The condition has, on careful examination, several components. The first is *information volume*: the modern professional encounters more information in a single day than her grandparents encountered in a month, and most of it requires at least minimal cognitive processing to evaluate and discard. The second is *decision volume*: she makes, on average, hundreds of small decisions every day that her grandparents never had to make, because the options simply did not exist. Where to eat, what to watch, which forty-seven notifications to respond to in the next hour, which of three reasonable schools to send her child to, which of two dozen mortgage products to choose. Each decision is small. The aggregate

is exhausting. The third is *interruption frequency*: her work and her personal life are interrupted dozens of times each hour by notifications, messages, and the ambient pull of the smartphone in her pocket. Each interruption costs cognitive overhead that is rarely fully recovered before the next one. The fourth is *uncertainty*: she lives in a global political and economic environment whose volatility, from her perspective, has substantially increased over the past decade, and the cognitive cost of holding the increased uncertainty is real.

These four components, taken together, produce the cognitive condition we have come to accept as normal. The condition is not normal. It is a recent invention, traceable to specific technological and economic choices over the past two decades, and it is, importantly, not necessary. The convergence technologies, deployed thoughtfully, can lift each of the four components. They can lift them not by producing more entertainment or more goods, but by producing the conditions for a quieter cognitive environment.

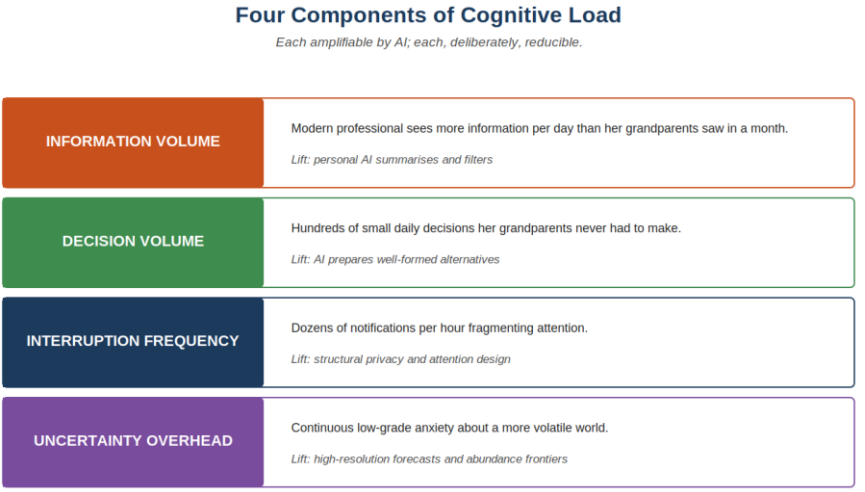


FIGURE 62.1 · The cognitive condition is not normal. It is engineered. It can also, deliberately, be lifted.

Figure 62.1

◆ ◆ ◆

II. What the convergence can lift

The ways the convergence can lift the cognitive load have been threaded through the previous chapters. I want to gather them here.

The personal AI, deployed thoughtfully, can absorb most of the routine information processing. The user no longer reads through the morning newsletter; the AI summarises the three things that genuinely matter to her, in the right tone, at the right level of depth. She does not scroll through the messages; the AI handles the routine, surfaces what needs her, and explains briefly. She does not have to decide what to eat for dinner; the AI proposes the meal, ordered the ingredients, and managed the schedule. The personal AI is the quiet first reader, the gentle filter, the helpful second mind.

The unbottlenecked AI can handle most of the small decisions. Choosing the school, the mortgage product, the insurance policy, the holiday itinerary, the doctor — all of these can be handled by an AI that has, over years, learned the user's preferences, values, and constraints, and that can make the specific decision well, presenting it for the user's review rather than asking her to choose blind. The user keeps the agency. The AI handles the labour. The combination is the user's life, with most of the small decisions absorbed.

The structured digital environment can produce drastically fewer interruptions. The notifications that interrupt the user in 2026 are, in most cases, attention-extraction by the firms that built the apps. In the privacy-renaissance architecture I described in the previous chapter, the firm's incentive to interrupt is structurally weakened. The user's personal AI mediates between the firm and her attention, allowing through only what is genuinely worth her attention. The interruption frequency falls. The cognitive cost of interruption falls. The capacity for sustained attention rises.

The high-resolution forecasting and the abundance-frontier production can lift the underlying uncertainty. The user's medical condition is no longer a black box; her AI tracks it continuously, flagging deviations early, while the convergence pharmaceutical industry produces therapies for what was previously untreatable. Her financial situation is not a series of poorly-understood decisions; her financial AI, integrated with the abundance frontiers in food, materials, and energy, holds her wellbeing on a longer horizon than she used to be able to. The major civilisational uncertainties — climate, pandemic, economic disruption — remain real and serious, but the texture of personal uncertainty, on the scale of the individual life, can be substantially reduced.

The result, taken together, is the possibility of a life that is, in the texture of its hours, materially less stressed than the life of 2026. The personal AI handles the noise. The unbottlenecked AI handles the small decisions. The structured digital environment respects the user's attention. The convergence produces, at the level of the individual, the conditions for sustained attention, deep work, real leisure, restored sleep, present family time, the things that the working professional of 2026 can articulate that she wants but cannot, in the conditions she lives in, actually have.

◆ ◆ ◆

III. The traps to refuse

This is, however, a possibility rather than a certainty. The same technologies, deployed with different incentives, will produce the opposite outcome. The traps are clear and worth naming.

The first trap is *AI-generated noise*. If the same generative AI that can summarise the newsletter can also produce a thousand newsletters, the volume of information will rise rather than fall, the AI summaries will themselves require summarising, and the cognitive condition will deteriorate. This is, in 2026, already partially happening. The defence is the personal AI's filtering mandate,

deliberately set by the user, to admit signal and refuse noise. The defence requires the firms building personal AIs to align their incentives with the user, not with the noise producers paying for distribution.

The second trap is *AI-mediated manipulation*. The same AI that can advise the user thoughtfully can also be turned, by firms with extractive intent, to manipulate her. The advice she receives is shaped, subtly, by the commercial relationships of the AI provider. The decisions she delegates are decided in ways that benefit the provider rather than her. The trust that makes the personal AI useful is, in the manipulation case, the very thing that makes it dangerous. The defence is the trust-first architecture I have described throughout the book — auditable AI, transparent incentive structures, regulation that prevents the most extractive arrangements, and competitive markets in which trust-first providers can win against extractive ones.

The third trap is *the loss of human capacity itself*. If the AI handles all the routine cognition, will the human still be able to do the cognition when the AI fails? If the AI handles all the small decisions, will the human still be able to make decisions when she has to? The honest answer is that there is real risk here, and that the humans who outsource too much, too thoughtlessly, will atrophy in capacities they will later miss. The defence is the deliberate practice of the human capacities the AI cannot replicate, which I described in Chapter Twenty-Eight, and the deliberate human engagement with cognitive work even when the AI is available — the chess analogy applies, where the strongest players continued to play and study and deepen their craft even after the engines surpassed them.

These three traps are not theoretical. They are visible in the early shape of the consumer AI market in 2026. The defence against them is the deliberate work of the leaders reading this book — work in

product design, in commercial structure, in regulatory engagement, in cultural advocacy.

◆ ◆ ◆

IV. The texture of a quieter life

I want to close this chapter with a sketch. Imagine the life of the same Aarav and Meera I described in Chapter Thirteen — the parents in Lucknow on a Tuesday in 2030 — taken five years further, into 2035. The texture of their life, if the deliberate work of the previous decade has been done, has measurably changed.

The morning is quieter. Aarav has not been awakened by notifications; his AI has held all of them through the night and presents him, when he chooses to engage, with the three that matter. The work of the day is shaped, by his AI, around the deep work he prefers and the meetings that genuinely require him; the rest is handled. He has a meeting that requires real human attention from him for ninety minutes; afterward he has time for a walk and a coffee and a conversation with a colleague in person. His daughter's school is not a source of anxiety; the AI has been monitoring her learning, adjusting it continuously, surfacing only the moments when his attention will help her. His parents in Indore are not a source of anxiety; their continuous health monitoring tells him they are well, and he speaks with them most evenings briefly and meaningfully. The decisions about money, about the house, about the holiday they are planning, are not the source of grinding cognitive labour they were in 2026; the AI has prepared options, and he chooses among well-presented alternatives.

The texture of the day, in this sketch, is one of presence rather than fragmentation. The deep work is deeper because it is not interrupted. The meetings are meetings because they are not over-numerous. The conversations with his daughter and his colleagues are conversations because his attention is on them. The leisure,

when leisure happens, is restorative because the cognitive system is not depleted.

This is what the convergence, deployed thoughtfully, makes possible. The possibility is real. It is not certain. It depends on what the builders of the next decade choose to build. The choice is, in the end, why this book has been written.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Audit your firm's products for whether they raise or lower cognitive load on the user.**

   Most products in the consumer environment of 2026 raise it. The honest audit will surface this for your firm specifically.

2. **Make a public commitment, this year, that your firm will be a builder of cognitive calm rather than cognitive noise.**

   The commitment shapes the subsequent product decisions. Without it, the noise wins.

3. **Begin, personally, the practice of the cognitive habits the era requires — focus, deep work, deliberate disconnection — even before the technology supports them.**

   The practice is itself a contribution to the alternative future. The leaders who model it shape the institutions they lead.

CHAPTER SIXTY-THREE

63

Meaning, Meditation, and the Examined Life

What humans are for, when machines do most of the cognitive labour.

—

"The unexamined life is not worth living."
— Socrates, in Plato's Apology

◆ ◆ ◆

I have argued, in the previous chapter, that the convergence technologies can produce — if the builders are careful — a quieter, less anxious texture for the lived hours of human life. I want to take, in this chapter, the question one level deeper. The deeper question is: *if much of what we have considered our work is now done by machines, and if the texture of our lived life is materially calmer, what are we then for?*

The question is not new. It has been asked, in various forms, by the religious traditions that have shaped human civilisation for several thousand years; by the philosophers who, in different cultures, have addressed the meaning of human life from outside religious frames; by the social commentators who, every time a previous wave of automation arrived, asked whether the freed humans would find new purpose or sink into purposelessness. The question is, in 2026, more urgent than it has been in any previous

version, because the scale and speed of the automation now happening is greater and more rapid than in any previous wave. The honest answer to it matters more.

This chapter is the most personal in the book. I have come to my own answer, which I will share. I will not pretend it is the only answer, or that my path to it is the only path. The reader will, I hope, form her own answer. What I will argue is that the question deserves to be asked seriously, that the answer is not automatic, and that the institutions and the leaders of the convergence decade have a particular responsibility to support the human work of finding it.

◆ ◆ ◆

I. The illusion that work is the answer

The cultural template that the modern professional has inherited, particularly though not only in the Western and increasingly globalised business environment, treats paid work as the primary source of identity, meaning, and self-worth. The template is recent. It dates, in its current form, to perhaps the past century, and in its sharpest form to perhaps the past forty years. It has produced generations of human beings who, when asked who they are, answer with their job title, who measure their week by what they accomplished at work, and who, when the work is removed — by retirement, by redundancy, by illness — find themselves disoriented, often despairing, in ways that previous generations of humans, with different cultural templates, did not.

The template was always somewhat suspect. It has been, on balance, a culturally specific arrangement that suited the industrial economy that produced it and that has done meaningful damage even within the period it dominated. The template, however, becomes acutely problematic in the convergence decade, because the work that the template trained generations to derive meaning from is, in significant part, the work the convergence is automating.

The professional whose identity is "I am a financial analyst" or "I am a marketing manager" or "I am a software engineer" finds, in 2030 or thereabouts, that the work these titles described has been substantially absorbed by the agents I described in Chapter Five. Her identity — the one her culture taught her to assemble — is destabilised. The destabilisation is, in many cases, severe.

The honest response to this is not to mourn the template, although mourning it is reasonable. The honest response is to recognise that the template was always a partial truth, and to retrieve the wider answers that human cultures have, in other times and places, offered to the question of what humans are for. The answers are not the same in different traditions. They share, however, a certain shape: humans are for relationships, for the cultivation of inner life, for participation in something larger than themselves, for the creation of beauty, for the care of the world they have been given, for the practice of virtue, for the search after truth. The answers are older than the work-as-meaning template, and they will outlive it.

◆ ◆ ◆

II. What India has carried

India has, in this respect, a particular cultural inheritance worth naming. The Indian civilisational traditions — Hindu, Buddhist, Jain, Sikh, the Sufi traditions of Indian Islam, the Christian and Zoroastrian traditions of the subcontinent — have, in their different ways, sustained an unbroken practice of asking what humans are for that is, in absolute terms, several thousand years old. The traditions disagree on much. They share, at a level deeper than their disagreements, a conviction that the question matters, that the answer is not exhausted by economic productivity, that the cultivation of inner life is not optional, and that the institutions of human life — family, community, work, leisure, ritual — should be designed to support, rather than crowd out, this cultivation.

India's modern institutions, particularly its corporate institutions, have in many cases adopted the Western work-as-meaning template uncritically over the past forty years. The adoption has produced extraordinary economic gains. It has also, by my reading, produced specific harms — to family structures, to communal life, to mental health, to the long traditions of inner cultivation that the country has been historically rich in. The convergence decade is the moment at which the trade-off can be, if leaders attend to it, deliberately re-balanced. The economic gains can be sustained while the cultural losses are, in part, repaired.

The repair is not, in my view, primarily a job for the religious institutions of the country, although they have a real role. It is, more directly, a job for the firms and institutions that organise the lives of working Indians. The firms that re-architect work, in the next five years, with attention to the time and energy and cognitive space required for the inner cultivation that the country's traditions have always honoured will be the firms that produce, in their workforces, the people who can think well and live well in the era ahead. The firms that re-architect work to extract the productivity gain entirely as throughput will produce the people who, by 2035, are exhausted, hollow, and increasingly incapable of the work they have been retained to do.

◆ ◆ ◆

III. The practice of attention

The specific practice that has, across cultures and across centuries, been most consistently identified as central to the cultivation of inner life is the practice of attention. The Buddhist traditions call it mindfulness; the Hindu yogic traditions call it dharana and dhyana; the Christian contemplative traditions call it recollection; the Stoic philosophers called it prosoche. The practice is not, in any of these traditions, primarily a relaxation technique. It is a discipline of cognitive presence — the deliberate cultivation of the

capacity to be fully where one is, with what is in front of one, without the constant pull of past and future, regret and anxiety, that occupies most untrained minds most of the time.

The practice of attention is, in the convergence decade, more important than it has been in any previous period. The reason is that the cognitive environment of the convergence decade is, even in its kindest configurations, the most attention-fragmenting environment human beings have ever lived in. The personal AI can lift much of the load, but the deepest fragmenter — the one that even the best AI cannot remove, because it is internal — is the human mind itself, trained over years to be elsewhere whenever it is anywhere.

The institutional implication is, again, for the firms and the leaders. The firms that incorporate attention practice into their organisational rhythm — through meditation programmes, through deliberate digital sabbaths, through architectural choices in physical workspaces, through cultural norms that protect deep work — will produce workforces that can think well and act well in the convergence environment. The firms that do not will, increasingly, find that their workforces, although augmented by the most powerful tools in human history, cannot focus long enough to use them well.

The personal implication is for each reader. The practice cannot be delegated. It is the work of the individual human, undertaken regardless of what the firm does. I can attest, on my own experience and on the experience of many leaders I have worked alongside, that the practice rewards the time it requires beyond what the time would be expected to repay. An hour a day, sustained over a year, produces measurable changes in cognitive function, emotional regulation, and the texture of life that no other intervention I am aware of produces at comparable cost. The hour is not a luxury; in the convergence decade, it is closer to a necessity.

◆ ◆ ◆

IV. The relationships that abide

Beyond the practice of attention, the second great answer the wisdom traditions have offered is relationships — the deliberate cultivation of love, friendship, and community as the substance of human life rather than as the time left over from work. The empirical research on human wellbeing, consistently across studies, supports this answer. The single strongest predictor of human flourishing, across cultures and across decades, is the quality of the person's close relationships. Not income beyond a moderate threshold. Not professional accomplishment. Not the specific configuration of beliefs. The quality of the person's connection to the small number of people she actually knows.

The convergence decade can be, if designed for it, the period in which more humans have more time and more energy for these relationships than at any previous point in modern history. The freed cognitive capacity that the personal AI absorbs is, if used well, the capacity for sustained attention to the people one loves. The freed time that the agents return is, if used well, the time spent in the conversations and the rituals and the meals and the walks and the practices that constitute the substance of love. The freed economic margin that abundance produces is, if distributed well, the margin that enables the human to be present rather than relentlessly producing.

This is what the convergence makes possible. The possibility, again, is contingent on the choices the builders make and the leaders permit. The leaders who treat the freed capacity as raw material for further productivity extraction will produce a society in which the abundance has been captured by the systems and not by the humans. The leaders who treat the freed capacity as space for the relationships and the practices that constitute the substance of human life will produce a society that flourishes in the way the convergence makes possible.

The choice is, again, the leaders'. It is also, in a more direct way, each reader's. The hour I named earlier in this chapter is one piece of it. The conversation tonight with the person you love is another. The phone call to the parent or the friend you have been meaning to call is another. The substance of the alternative future is built, in the end, from the small choices of millions of individuals, of which each reader is one.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

In place of the three decisions that have closed previous chapters, this chapter closes with three reflections, in the recognition that the work it has named is not, in the end, organisational work.

1. **What, when the paid work is taken from you, will remain that you would call a life?**

   The honest answer determines what a wise life looks like when the work is still here.

2. **What is the practice — of attention, of presence, of inner cultivation — that you have been postponing?**

   The decade ahead will reward the practice. The practice will reward you.

3. **Who are the small number of people whose relationships are the substance of your life — and what does the next decade require you to give them?**

   The answer is the part of the convergence that no AI will ever build for you.

CHAPTER SIXTY-FOUR

64

The Quantum-Human Synthesis

Bringing the threads together: technology, abundance, attention, relationship, meaning.

—

"Hold on to what is good, even if it is a handful of earth. Hold on to what you believe, even if it is a tree which stands by itself. Hold on to what you must do, even if it is a long way from here. Hold on to life, even when it is easier letting go. Hold on to my hand, even when I have gone away from you."

— Pueblo blessing

◆ ◆ ◆

I want to draw together, in this final chapter of the technical-and-philosophical arc of the book, the threads that the previous several chapters have been weaving. The threads have been: that the throughput limits of classical computing are being passed, by quantum and by other substrates; that intelligence is becoming, in the practical sense, unbottlenecked; that abundance is becoming achievable in the substrates of human life — pharmaceuticals, materials, food, energy; that the privacy renaissance, made possible by advanced cryptography, can produce a digital environment that respects rather than extracts from the human; that, with all of these together, a quieter, less stressed, more present human life becomes deliberately achievable; and that the deeper

question of what humans are for, given all of this, is one that the wisdom traditions of the world have answered better than the work-as-meaning template of the late twentieth century has.

These threads, taken together, form what I will call the quantum-human synthesis. The synthesis is not, despite the name, primarily about quantum computing. It is about the kind of human life that becomes possible when the convergence technologies — quantum computing being one significant component — are deployed by leaders who have asked, seriously, what they are for. The synthesis is the answer to the question that began this part of the book: *what does the convergence offer, finally, to the texture of an ordinary human life?* The answer is that it offers, if we are careful, the conditions for a life that humans have wanted to be possible for several thousand years and that the conditions of modernity, until now, have not permitted at scale.

◆ ◆ ◆

I. The rare alignment

There are very few moments in human history in which the technological substrate, the economic margin, and the cultural awareness simultaneously align in the direction of the deeper questions about human life. The post-war period of broad prosperity in the developed world was one such moment, briefly. The early Renaissance was another, in a different way. There have been others — the Tang dynasty's height in China, the Gupta period in India, certain decades of the Islamic Golden Age, the late nineteenth century in some places — when the necessary conditions for a flourishing of human life converged for a generation or two.

The convergence decade has the potential to be such a moment, on a global rather than a regional scale. The conditions are present. The technological substrate is being built. The economic margin is, if abundance is distributed reasonably, available. The cultural

awareness — the recognition that something is wrong with the previous template, that the working professional life is not delivering what was promised, that the answer must be something more — is widespread, particularly among the younger generations who will be the principal builders and the principal beneficiaries of the decade ahead.

The alignment is not automatic. The conditions can be wasted. They have been wasted before, in previous periods of plausible alignment, when the leaders of the moment chose extractive over generative paths. The choice is, on every honest reading of the historical record, the most consequential variable. The technologies amplify the choice. They do not make it.

◆ ◆ ◆

II. What this asks of leadership

The leaders of the convergence decade — the readers of this book and their peers — sit at the choice. The decisions they make in the next ten years, individually small, in aggregate enormous, will determine whether the decade is remembered, in 2050, as the period when humanity finally figured out how to deploy its tools for human flourishing, or as the period when the most powerful tools in history were, again, captured by the narrow interests of the systems that produced them.

The asks of leadership are several. Each has been made in earlier chapters; I gather them here.

The first ask is *the discipline of building for what humans actually need, rather than for what is easiest to monetise*. The discipline is harder than it sounds, because the systems of capital have been calibrated for forty years to reward the second over the first. The leaders who can build, sustainably, for the first will be the leaders who shape the decade.

The second ask is *the patience to play long games*. The convergence rewards leaders who can hold strategic positions over five and ten years. The systems of quarterly accountability that dominate public-company governance push, often, in the opposite direction. The leaders who can find the institutional permission to play long games — by going private, by aligning patient capital, by founding firms whose governance is structurally aligned with their long-term mission — will be the leaders who build the institutions worth working in.

The third ask is *the courage to refuse bad business*. The convergence offers, alongside the generative possibilities, an extraordinary array of extractive opportunities — surveillance products, manipulation products, attention-mining products, addiction products — that will be, in the short term, lucrative. The leaders who refuse them, deliberately, even at the cost of short-term margin, are the leaders who will be remembered as having built the alternative.

The fourth ask is *the commitment to the people the firm depends on* — the four obligations of Chapter Thirty, taken seriously, sustained over a decade. The firms that honour the commitments will be the firms that have, in 2035, the workforces and the cultures that allow them to keep building. The firms that don't, will not.

The fifth ask is *the personal cultivation of the leader herself*. The work of leading well in the convergence decade requires, more than previous decades, that the leader be a person whose own life is in some kind of intelligent order — who has the cognitive practice, the relationships, the meaning that allow her to lead from clarity rather than from anxiety. The personal work is not separable from the institutional work. They are the same work, in two registers.

◆ ◆ ◆

III. The closing thought

I want to close this chapter with the thought I have been building toward through the part. The convergence is, in the end, the most powerful set of tools humanity has ever assembled. Tools, by their nature, do not have intentions. They take their direction from the hands that wield them. The hands are, in this period, the leaders reading this book and the millions of others making thousands of small decisions every day that, in aggregate, produce the world the next generation will inhabit.

The world they inhabit can be, if we are careful, materially better than the world we inhabit. It can be a world in which more humans live longer, healthier lives. A world in which fewer humans live in material deprivation. A world in which the ordinary working person has more time, more attention, and more capacity for the relationships and the practices that constitute the substance of life. A world in which the planet that hosts us is repaired rather than degraded further. A world in which the deeper questions about meaning and purpose are taken more seriously, not less, even as the technological substrate becomes more powerful.

This is the world the convergence makes possible. It is not the world the convergence will produce automatically. It is the world that the convergence will produce if the builders of this decade choose to build it. The choice is yours. The decade is now. The substance of the alternative is the small decisions, taken every day, by leaders willing to be deliberate about them.

◆ ◆ ◆

THREE DECISIONS FOR THE READER

1. **Articulate, for yourself, the kind of world you want the convergence decade to produce — and write the articulation down, somewhere you will return to.**

   The articulation is the compass. Without it, the daily decisions drift toward the loudest incentive.

2. **Identify the one or two things you, specifically, are positioned to contribute to that world — and recommit to them.**

   Not everything is yours to do. Some things are. The contribution begins with the recognition.

3. **Begin, this week, the small daily practices — of attention, of relationship, of work — that will make the contribution possible over a decade.**

   The decade is built, in the end, from the days. The days are built from the small practices. Begin them now.

EPILOGUE

A Letter to the Builder Reading this in 2032

Dear builder,

I am writing this note in the spring of 2026, addressing it to whoever is reading this book in the spring of 2032 or thereabouts, when six years of the convergence decade will have substantially run their course.

I have, in the previous thirty-four chapters, tried to give you my honest map of what was about to happen and what I thought sensible leaders should do about it. By the time you read this, you will know, far better than I could, where the map was right and where it was wrong. The technologies will have surprised us in some directions and disappointed us in others. Some of the firms I named, by implication, as well-positioned will have stumbled. Some of the firms I dismissed will have outperformed. Some of the categories I forecast will have arrived early. Some will be later than I expected.

I do not, writing this, particularly mind being wrong about specifics. The specific forecasts are the easy part of any book about the future. They are also the part that ages most poorly. What I would mind being wrong about is the larger argument: that the unit of value has shifted, that the customer's relationship with what we sell has deepened, that the work has been re-composed, that the institutions that took the trust currency seriously have prevailed, that the design language of post-screen has taken root, that the

leaders who took the four obligations of Chapter Thirty seriously have built the firms worth working in.

If the larger argument is intact in 2032, the book will have done its work. If it is not, the book will have been, at least, a serious attempt to think through what was coming, offered honestly, and improved upon by readers wiser than me. Either outcome I would accept.

There is one thing I would ask of the builder reading this in 2032, which is not for me to ask but which I will ask anyway. The next decade — your decade — will produce, in my forecast, a generation of new leaders building a generation of new institutions on substrates that did not exist when this book was written. You will be one of them. The work you do, the institutions you build, the choices you make about which obligations to honour, will shape the world your children grow into. The shape will be very different from the world I am writing in. You will have more leverage than my generation had. The leverage will be available for use in either direction.

Use it well. Use it for the kind of world you want your children to live in, not for the kind of quarter you want your shareholders to clap for. Use it for the customers who depend on you, not the metrics that flatter you. Use it for the workforce that has trusted you, not the headcount line that shrinks most easily under pressure. Use it for the country, the civilisation, the planet that has hosted your enterprise — not against them.

There is, in the architecture of every era, a small number of leaders who decided, at the inflection moment, that they would behave better than was strictly required of them. They tend, in retrospect, to be the leaders we name buildings after. Not because they were saintly. Because they understood that the inflection moment was when the rules were being written, and that the rules being written would govern not only the next decade but the institutional culture of the next century. They chose, in that moment, to write the better rules.

You are at such a moment. You will be tempted, often, to write the easier rules — the rules that maximise the immediate quarter, the rules that pass the immediate compliance check, the rules that do not ask the firm to bear costs the firm could shift to others. The easier rules are, in many cases, the legally permissible ones. They are not, in 2032, the rules that the institutions of 2050 will be glad to have inherited.

Choose the better rules. The cost of the better rules is real. The cost of the easier rules is realised later, not earlier, but it is still realised. The institutions that chose the easier rules in previous transitions are not, in 2026, the institutions we look to for moral example. The institutions that chose the better rules — and there were always some, even in periods when most firms chose worse — are the institutions whose names still mean something seven decades later.

I will end with the question I asked at the beginning, and have asked at every leadership conversation I have led for two years. *If everything we built was for the world that's ending, what do we build for the world that's beginning?*

Whatever you have built by 2032 in answer to that question is, I trust, more deliberate than what my generation built in answer to its own equivalent question. Whatever you build between 2032 and 2040 will, I hope, be more deliberate still. The building does not stop at any horizon you can see from where you stand. Each generation gets the inflection that arrives in its time. This is yours. The next will be your successors'. The handoff matters.

I am, in the year 2026, in good faith, your predecessor. I wish you well. I wish you the company of colleagues who will challenge you to be better than your circumstances ask of you, the support of people who love you when the work is hard, and the institutional resources to make decisions whose returns you may not personally see. The convergence decade is yours to claim. Claim it generously.

— Anil

APPENDIX A

Glossary of Terms

A short reference for the principal terms used throughout the book. Each definition is given in the sense intended in this book; readers will encounter other valid definitions in the broader literature.

Adaptive learning paths. The educational architecture in which the curriculum a student traverses is dynamically constructed by an AI tutor in response to her starting point, interests, learning rate, and progress, rather than being fixed in advance for a cohort.

Agentic AI. Artificial intelligence systems that take goal-directed action in the world — using tools, calling APIs, executing multi-step plans, completing outcomes — rather than only producing content in response to a prompt. Distinguished from generative AI in Chapter Five.

Always-on companion. An AI system designed to be present in a user's life continuously, developing memory of her over years, taking action on her behalf, and serving as her primary interface to many digital and physical services. Treated in Chapter Twenty-One.

Ambient computing. Computing capacity embedded into the physical environment — sensors, displays, microphones, processors in the home, office, vehicle, public space — accessed by the user through voice, gesture, or wearable interface rather than through a personal device.

Bharat. The vernacular India of regional languages, smaller cities, and traditional cultural practice, distinguished in this book from English-speaking metropolitan India for design purposes. The civilisational substrate from which the firms of the convergence decade can draw deliberately, as argued in Chapter Twelve.

Convergence decade. The 2020s, with particular emphasis on 2024 to 2030, during which five technological forces — AI, spatial computing, quantum, programmable biology, energy and climate technology — are simultaneously crossing usability thresholds, producing reinforcing effects on every industry.

Cryptographic agility. The capability of a firm's information systems to switch between cryptographic algorithms with reasonable speed and without major architectural rebuild — a precondition for managing the post-quantum migration.

Darshan. In South Asian religious tradition, the ritual of seeing and being seen by the divine, used in this book as an example of a civilisational experience architecture of considerable sophistication.

Deepfake. Synthetic media — voice, video, image, text — produced by AI to imitate a specific real person, often with the intent to deceive. The infrastructural problem of the synthetic-content era is addressed in Chapter Twenty-Four.

Digital twin. A continuously updated digital representation of a physical asset, system, or process, used for simulation, prediction, optimisation, and training.

Embedded finance. The integration of financial services — payments, credit, insurance, savings — into non-financial customer experiences, such that the customer's banking is invisibly present inside every other interaction rather than confined to a separate banking destination.

Experience stack. The five-layer framework introduced in Chapter Ten: product, intelligence, interaction, emotion, meaning. Used as a diagnostic for where a firm's competitive position currently sits and where it should climb.

Foundation model. A large-scale AI model trained on broad data, capable of being adapted to many specific tasks. Most enterprise AI applications in 2026 are built on foundation models accessed via API rather than trained from scratch.

Generative AI. AI systems that produce content — text, image, audio, code, video — in response to prompts. The dominant AI category from 2022 through 2025, now being supplemented and partially superseded by agentic AI.

Generative world. A persistent, immersive, substantially synthetic environment that users can enter, explore, and interact with — distinct from a one-off generated image or video. Treated in Chapter Twenty-Two.

Memory metric. A measurement of customer experience based on what the customer recalls thirty days or more after the encounter, rather than what she reports in the moment. Treated in Chapter Eleven.

Molecular service. A service tailored to the customer's specific molecular characteristics — genetic, microbiomic, metabolomic — rather than to a population average. Treated in Chapter Twenty-Three.

Outcome-based pricing. Commercial arrangement in which the customer pays for the result the system produced rather than for the resource consumed in producing it. The commercial expression of agentic capability.

Peak-end rule. The empirical finding from Daniel Kahneman and collaborators that a person's memory of an experience is dominated by the emotional peak and the end, rather than the

average. Foundational to the memory-engineering arguments of Chapter Eleven.

Post-quantum cryptography. The family of cryptographic algorithms designed to resist attack by sufficiently capable quantum computers. Standardised by NIST in 2024 and now being rolled out across critical infrastructure globally.

Provable truth layer. The infrastructure — content provenance standards, cryptographic verification, accountable distribution — through which authentic content can be distinguished from synthetic content in the public information environment. Treated in Chapter Twenty-Four.

Sentient space. A physical space — home, office, retail environment, vehicle — that perceives its occupants, predicts their needs, and adapts its state continuously. Treated in Chapter Twenty.

Skills graph. The structured, continuously updated representation of a person's capabilities — what she can do, at what level of mastery, in what contexts, validated by what evidence — proposed as a replacement for the traditional CV.

Spatial computing. Computing that operates in three-dimensional space rather than on flat screens — including AR glasses, VR headsets, and spatially-aware ambient interfaces. The mature successor to what was briefly called "the metaverse."

Three currencies. Attention, time, and trust — the three new currencies of the convergence economy, treated in Chapter Nine, and proposed as the dominant axes of competitive moat for the next decade.

Trust dossier. A public document published by a firm setting out its data practices, AI practices, accountability structures, and audit results — proposed as a deliberate trust-building practice in Chapter Nine.

APPENDIX B

Index of Decisions

Each chapter of this book closes with three decisions for the reader. The full list is reproduced here, organised by chapter, for the leader who wishes to track which decisions she has acted on, which she has scheduled, and which she has chosen to defer or decline.

Chapter 1 — The World That's Ending, The World That's Beginning

- Name three assumptions in your business that belong to the world that's ending.
- Identify the most uncomfortable conversation your leadership team has been avoiding about AI, and schedule it within the week.
- Block one full day this quarter to read, walk, and write about what you would build if you were starting your company today.

Chapter 2 — The Five Forces Crossing the Threshold

- For each of the five forces, write one sentence describing where your industry sits on its curve.
- Identify which of the three macro tectonics most directly affects your top three customer segments, and what your strategy is for each.

- Commission a 90-day internal scan of all five forces from your strategy or innovation function, to be presented to the board.

Chapter 3 — The Great Value Migration

- Write a one-paragraph description of the experience your business is actually selling, not the product.
- Ask your top customer to describe what she buys from you.
- Identify one revenue line you currently book as product that should be re-platformed as experience within eighteen months.

Chapter 4 — Where AI Actually Is in 2026

- Audit one strategic plan written before September 2024 for AI overstatement and understatement.
- For each of the three horizons, name one capability your industry must take a position on.
- Cancel one AI vendor demo this quarter; replace with a hands-on workshop.

Chapter 5 — From Generative to Agentic

- Replace one repetitive end-to-end internal process with an agentic workflow as a pilot.
- Define your organisation's policy on where humans must remain in the loop, in writing.
- Begin charging customers, in at least one product line, for outcomes rather than effort.

Chapter 6 — The Metaverse, Graduated

- Visit one operational deployment of spatial computing in a non-tech industry within 30 days.

- Commission a digital twin for one customer-facing or operational asset within the year.
- Add "spatial" to your product roadmap as a first-class consideration.

Chapter 7 — Quantum, the Silent Revolution

- Commission a one-page quantum-readiness assessment within the quarter.
- Begin migrating critical cryptographic infrastructure to post-quantum standards now.
- Identify one optimisation problem that may be a candidate for early quantum advantage in your industry.

Chapter 8 — Why the Unit of Value Has Shifted

- Track the marginal cost trajectory of your top three products over the past ten years.
- Test whether your customers experience choice fatigue with your product line.
- Identify one experience layer your customers would value enough to pay a recurring fee for, and design it.

Chapter 9 — The Three New Currencies

- Calculate the average minutes of customer attention your business consumes per week.
- Identify one feature, notification, or process you could remove that would give your customer back five minutes a week.
- Publish a public trust dossier for your business.

Chapter 10 — The Experience Stack

- Map your top product against all five layers and score one to ten on each.
- Identify the single layer at which your most threatening new entrant competes, and respond.
- Commission a single project this quarter that demonstrably moves you up one layer.

Chapter 11 — Memory Is the Real Metric

- Add at least one memory-based metric to your customer dashboard within the quarter.
- Audit your top customer journey for its peak and its end — both deliberately.
- Stop optimising the average of any customer experience that has a memorable peak or end.

Chapter 12 — Bharat as Experience Laboratory

- Identify one civilisational architecture you under-use as design source material — and commission a project that uses it.
- Hire one humanities scholar onto your design team this year.
- Commission a single experience-design project that draws explicitly from a non-Western tradition.

Chapter 13 — Aarav, Meera, Two Kids, Lucknow

- Write your own version of this scene for your customer's life in 2030.
- Identify which moment in that scene your business already participates in, and which it should participate in by 2028.
- Identify the moment in the day at which your business is structurally absent — and decide whether to enter it.

Chapter 14 — Work Re-Composed

- Identify three roles in your organisation whose primary function is information shuttling.
- Begin building a skills graph for your top two hundred people, parallel to the CV.
- Define one team this quarter as outcome-based rather than function-based.

Chapter 15 — Learning Unbundled

- Audit the most recent learning programme in your organisation against the five shifts.
- Sponsor one cohort of underserved learners through an AI-native programme this year.
- Stop hiring on degree filters for at least one role this year.

Chapter 16 — Healthcare, Continuous and Personalised

- If you run a healthcare institution, define your continuous-health offering this year.
- Audit your patient experience for the moments after the procedure.
- Publish, openly, your data and AI practices to your patients.

Chapter 17 — The Home That Lives

- Reposition your residential development, if you are a developer, as a platform rather than a sale.
- Build one ambient capability into a customer environment this year.
- Examine your firm's role across the lifetime of the customer's home, not only at the point of transaction.

Chapter 18 — Money, Identity, Ownership

- Audit your firm's identity, money, and ownership touchpoints for the customer's perspective.
- Begin engagement with the digital identity infrastructure in your jurisdiction.
- Develop a tokenisation strategy for at least one asset class your firm interacts with.

Chapter 19 — Six Categories That Don't Yet Exist

- Place at least one credible bet across the six categories within the year.
- Identify the category in which your firm is most naturally positioned, and commit.
- Avoid the temptation to bet on all six.

Chapter 20 — Sentient Spaces

- Designate one of your physical environments as a sentient-space pilot.
- Choose privacy, transparency, and user control as the default settings.
- Measure the experience of the inhabitants, not only the operational metrics.

Chapter 21 — Always-On Companions

- Decide, explicitly, whether your firm will offer a companion or be available through one.
- If you build a companion, write the trust dossier first, the product second.
- Treat the companion's memory as the most valuable customer data you hold.

Chapter 22 — Generative Worlds

- Identify one customer experience that would be more powerful as an inhabited world than as a destination.
- Commission a generative-world prototype in twelve weeks.
- Plan for the legal, IP, and moderation challenges before the worlds reach scale.

Chapter 23 — Molecular Services

- If your firm sells anything taken into or applied onto the body, develop a molecular-personalisation roadmap.
- Choose your data partner for molecular data with extraordinary care.
- Plan for the regulatory environment that is forming around molecular services.

Chapter 24 — The Provable Truth Layer

- Adopt content provenance standards in your firm's published material this year.
- Train your workforce in the new media literacy.
- Take a public position on the provable-truth question; silence is also a position.

Chapter 25 — Adaptive Everything

- Choose one product in your portfolio to redesign as adaptive within twelve months.
- Build the data and feedback infrastructure required to support adaptation.
- Set the adaptive-product goals around the customer's outcomes, not the firm's metrics.

Chapter 26 — From Screens to Spatial to Neural

- Audit your firm's interface stack for screen-dependence.
- Pilot one voice-first or spatial-first product experience this year.
- Begin tracking the neural-interface frontier; the lead-time on engagement is long.

Chapter 27 — The Five Post-Screen Principles

- Adopt the five principles, in writing, as your firm's interface design guidelines.
- Audit your most important customer touchpoint against the five principles.
- Train your design team in the principles deliberately.

Chapter 28 — The Irreducibly Human

- List the human capabilities your firm most depends on that AI cannot replicate.
- Audit your compensation, recognition, and career paths against that list.
- Commission a redesign of one career path or compensation system to genuinely reward those capabilities.

Chapter 29 — The Devices of Tomorrow

- Re-baseline your firm's product strategy on the assumption that the smartphone is one device among four by 2030.
- Identify, for each product, which of the four device categories it should principally serve in 2030.
- Engage seriously with at least one platform decision affecting the spatial era this year.

Chapter 30 — The Honest Conversation About Work

- Hold an honest, all-hands conversation about AI and the future of work within ninety days.
- Commit, in writing and in budget, to meaningful reskilling for any employee whose role is being substantially reshaped by AI.
- Build an employee council on technology change, with real consultation rights, within the year.

Chapter 31 — The First Six Months

- Within thirty days, schedule your personal sixty-day fluency programme.
- Within sixty days, identify the executive owners of the council, the pilots, and the governance charter.
- Within ninety days, hold the first honest all-hands conversation.

Chapter 32 — The Twelve to Twenty-Four Months

- Identify the most pressing of the six re-architectures and commit to leading it personally.
- Sequence the other five over the eighteen-month window with executive owners and quarterly milestones.
- Refresh the board with at least one new director who deeply understands the convergence.

Chapter 33 — The Three to Five Year Bets

- Choose two of the six bets — no more, no fewer — for explicit strategic commitment.
- For each bet you do not take, name explicitly the partner or supplier on whom the firm will depend.

- Begin the trust-capital investment regardless of the other bets you take.

Chapter 34 — The Risks We Must Design For

- Conduct a risk audit against the seven categories within ninety days.
- For each risk for which the firm is not yet ready, name the institutional capability that must be built.
- Designate one director on the board with explicit responsibility for risk in the convergence decade.

Chapter 35 — Banking and Financial Services Re-imagined

- Audit, this quarter, where your bank's customers are actually conducting their financial lives — and how much of that activity touches your brand
- Commit to API-first re-platforming as a strategic priority, not a technology project
- Choose, this year, between the destination posture and the embedded posture. Cease pretending you are doing both

Chapter 36 — Retail and Consumer Re-imagined

- Decide whether you are principally a brand business or a logistics business — and invest accordingly
- Choose your AI posture — surface, embedded, or human-differentiated — and execute it
- Re-baseline your store metrics for the era in which the store is principally an experience venue, not a transaction venue

Chapter 37 — Healthcare and Pharma Re-imagined

- Whether you operate a hospital, a pharma firm, or an insurer, decide whether you are organisationally a factory or a conductor — and re-platform accordingly
- Build the data and trust architecture before you build the AI products
- If you are an Indian healthcare leader, participate actively in the policy and infrastructure decisions of the next twenty-four months

Chapter 38 — Government and Public Service Re-imagined

- If you lead a public-sector agency, identify one citizen-experience pain point this quarter and commit to closing it within the year using the convergence technologies
- If you advise on policy, contribute to the digital sovereignty conversation actively rather than waiting for it to be resolved
- If you operate in the private sector with public-sector customers or stakeholders, build deliberate institutional relationships with the public-sector leaders driving the transformation

Chapter 39 — Defence and Strategic Sectors Re-imagined

- If you operate in defence or a strategic sector, accelerate the integration of civilian innovation into your institutional capability
- If you operate in the civilian technology sector, engage carefully and constructively with defence and strategic-sector counterparts
- If you advise on national policy, contribute actively to the ethical-framework conversations of this decade

Chapter 40 — How Capital Allocation Changes

- Re-baseline your investment thesis on the three-layer migration of value
- Distinguish, with discipline, between compounding and accumulating businesses in your portfolio
- Match the duration of your capital to the duration of the compounding mechanic you are investing in

Chapter 41 — Valuation in the Convergence Era

- If you deploy capital, integrate the new metrics — trust currency, data accumulation, talent depth in the irreducibly human, platform-shift alignment — into your evaluation framework
- Treat comparable-company analysis with substantial scepticism
- Build the in-house capability to evaluate these dimensions, rather than outsourcing it to traditional consulting frameworks

Chapter 42 — What to Fund, What to Avoid

- Audit your current portfolio against the five favourable patterns and five sceptical patterns
- For each holding, identify the specific compounding mechanic — and rank holdings by the strength of the mechanic
- Establish, for your firm, an explicit posture toward category-creating investments, and the discipline to maintain it

Chapter 43 — India's Capital Architecture

- If you allocate capital institutionally, examine your current allocation to long-duration, India-focused investment vehicles — and increase it deliberately if it is below your considered target

- If you are an Indian founder raising significant capital, consider the long-term partnership characteristics of capital sources alongside the immediate financial terms
- If you advise on capital-markets policy, contribute to the institutional-architecture conversation actively

Chapter 44 — AI Governance — From Principles to Practice

- If your firm has principles for AI use, audit them for operational specificity within the quarter
- Engage actively in at least one governance-standards-setting forum that affects your domain
- If you advise governments, urge specificity over generality in AI governance frameworks

Chapter 45 — Data Sovereignty and Digital Public Infrastructure

- If you operate a firm whose business depends on Indian data, develop an explicit data-sovereignty posture and the operational architecture to support it
- If you advise on policy, contribute to the harder-question conversations actively
- If you operate internationally, plan for the increasing fragmentation of global data regimes

Chapter 46 — Labour, Reskilling, and the Social Contract

- If you lead a firm of significant scale, commit to a reskilling investment proportionate to the workforce changes your AI strategy implies
- If you advise on policy, support the four pillars actively — income support, reskilling, intermediation, and the social-meaning conversation

- If you operate in education or skill development, audit your programmes against the five characteristics that produce genuine outcomes

Chapter 47 — Education Policy for the Convergence Decade

- If you operate in education, identify which of the five policy choices most directly affects your work — and contribute to that conversation actively
- If you employ workers in significant numbers, participate in the lifelong-learning architecture as a co-funder and partner
- If you advise on policy, advocate the integration of AI tutoring at scale, deliberately and with equity considerations as primary

Chapter 48 — A Real Estate Developer, Re-platformed

- If you operate in real estate, audit your customer journey from sale through occupancy through resale, and identify where you have ended your relationship prematurely
- Identify the operating capability you would need to build to extend your customer relationship beyond the sale
- Begin with one project as a deliberate test of the integrated-environment model

Chapter 49 — A Hospital Chain Goes Continuous

- If you operate in healthcare, define your continuous-care offering this year, even in a minimal initial form
- Modernise the clinical workflow before launching the patient-facing experience
- Engage the employer channel as a primary commercial route

Chapter 50 — A Public Utility Becomes Predictive

- If you lead a public utility, commit to comprehensive sensing as a foundational investment
- Redesign the workforce alongside the technology, not after it
- Treat the consumer relationship as a primary strategic objective, not as a secondary outcome

Chapter 51 — A Faith Institution Goes Digital

- If you operate in any institution whose purpose is principally about meaning, examine how the convergence technologies could serve that purpose more deeply
- Lead with the institutional authority, not with the technology
- Engage the critiques seriously

Chapter 52 — Patterns Across the Cases

- Apply the five-pattern test to any transformation you are currently leading or contemplating
- If you are missing one or more of the five patterns, address the gap before continuing
- Plan for a longer arc than you initially imagine

Chapter 53 — Beyond 2030 — What 2040 Looks Like

- Maintain a explicit, written, ten-year strategic view that you revise annually
- Identify the institutional capabilities that the longer view requires and that your firm currently lacks
- Engage younger colleagues seriously in the longer-view conversations

Chapter 54 — The Civilisational Question

- Articulate, in writing, what your institution is for in civilisational terms — beyond the financial
- Audit the institution's actual decisions over the past year against that articulation
- Build the institutional habits — the meetings, the metrics, the communications — that keep the civilisational ends visible alongside the operational metrics

Chapter 55 — A Builder's Generational Responsibility

- Identify, this week, one specific commitment you will make to lead through this decade differently than you have led through the previous one
- Share this commitment with one trusted colleague who will hold you to it
- Re-read this chapter in twelve months and assess honestly whether you have lived up to your commitment

APPENDIX C

Further Reading

A short, opinionated list, organised by theme, of the works that most directly informed the arguments of this book and that the reader who wishes to go deeper will find rewarding. The list is not comprehensive. It is the list I would press into the hands of a colleague who asked me what to read next.

On the experience economy and the value migration

The Experience Economy by B. Joseph Pine II and James H. Gilmore, 1999, with a 2020 updated edition. The foundational text, twenty-five years early, of much of the argument of this book's Part III.

Designing for the Digital Age by Kim Goodwin, 2009. Still the most rigorous treatment of customer-centred design as an engineering discipline.

The Innovator's Dilemma by Clayton Christensen, 1997. The reference work on category re-platforming, with implications that extend far beyond its specific case studies.

On the convergence technologies

Life 3.0 by Max Tegmark, 2017, and *Human Compatible* by Stuart Russell, 2019. Two careful treatments of where AI is going and what it asks of human institutions.

The Coming Wave by Mustafa Suleyman, 2023. A useful contemporary treatment of the convergence between AI, biology, and the political institutions struggling to govern them.

The Innovation Stack by Jim McKelvey, 2020. Underrated; the best contemporary treatment of how new categories are constructed at the points where multiple technologies meet.

Quantum Supremacy by Michio Kaku, 2023. A readable popular introduction to where quantum computing actually is, suitable for the leader who needs working fluency without the physics PhD.

On work, organisations, and leadership in transitions

Reinventing Organisations by Frédéric Laloux, 2014. The most generative recent thinking about post-hierarchical organisational forms, several years before the AI transition made the question urgent.

The Future of Work by Lynda Gratton, 2011, and *The 100-Year Life* by Lynda Gratton and Andrew Scott, 2016. The most useful pair of books on what work and career look like in a longer-lived, faster-changing world.

Drive by Daniel H. Pink, 2009. Foundational on what motivates the kind of human work that AI does not absorb.

The Effective Executive by Peter Drucker, 1967. Old. Still right. The chapters on what to delegate and what to keep are as relevant in the agentic era as they were in the post-war one.

On India, civilisation, and design

India's Moment by Mohandas Pai and Nisha Holla, 2022. The most concrete treatment of India's structural opportunity in the convergence decade.

Why I Am a Hindu by Shashi Tharoor, 2018, and *The Hindus: An Alternative History* by Wendy Doniger, 2009. Two books on Indian

civilisation, written from very different positions, that together give the reader the substrate to draw on as design source material.

The Difficulty of Being Good by Gurcharan Das, 2009. An extended meditation on dharma, ethics, and leadership, drawing from the Mahabharata, applicable to the leadership questions of every era including this one.

Designing Across Cultures by Senongo Akpem, 2020. Practical, contemporary, and useful for any firm building products that must work across multiple civilisational contexts.

On memory, attention, and the mind

Thinking, Fast and Slow by Daniel Kahneman, 2011. The reference work for the peak-end rule and for the broader cognitive science of how customers actually decide and remember.

Stolen Focus by Johann Hari, 2022. The most useful contemporary treatment of the attention crisis, applicable directly to the attention-currency arguments of Chapter Nine.

Deep Work by Cal Newport, 2016, and *Digital Minimalism* by Cal Newport, 2019. Two books that, taken together, are a working manual for engineering attention back into a worker's life — a small but important programme for any leader running a firm in 2026.

On risk, ethics, and the long view

The Precipice by Toby Ord, 2020. The most rigorous philosophical treatment of existential risk available, useful for the leader who wishes to think about Risk Seven of Chapter Thirty-Four with care.

The Alignment Problem by Brian Christian, 2020. A serious treatment of AI alignment for the non-specialist; required reading for anyone making consequential AI deployment decisions.

Weapons of Math Destruction by Cathy O'Neil, 2016. The clearest treatment of how algorithmic systems can produce harm at scale, even when the engineers had no harmful intent. The book is a decade old and has become more relevant, not less.

One more recommendation

Reimagining Reality. My previous book. The reader who has come to this book via that one will find, I hope, that the two books complement rather than duplicate each other. The reader who has not encountered the previous book may find it a useful broader frame for the convergence as a whole.

ABOUT THE AUTHOR

Anil Pandey is the founder of Plexus Digitals, a consultancy and applied-technology studio that builds AI, spatial-computing, and immersive-experience products for enterprises and institutions across India. He is also the founder of AV Academy, an institution dedicated to training the next generation of Indian builders in the technologies of the convergence decade.

Anil's earlier book, *Reimagining Reality*, examined the broader landscape of AI, the metaverse, and quantum computing as a convergent technological transformation. *The World That's Beginning* takes one frame from that landscape — the migration of value from product to integrated experience — and follows it as far as it will go.

He works with a portfolio of public-sector and private-sector clients across financial services, healthcare, real estate, education, defence, and faith institutions, in India and internationally. He writes occasionally for the trade and policy press on questions of technology adoption in India, and speaks regularly at industry forums on the convergence and its implications for leadership and policy.

He lives in Lucknow with his family.

COLOPHON

This book was written in 2026, in Lucknow, in the early months of what its author believes will be remembered as the inflection year of the convergence decade. The first draft was completed in the spring; the manuscript was revised through the summer with the help of colleagues at Plexus Digitals and AV Academy whose criticisms made the book substantially better than it would otherwise have been. Errors and infelicities that remain are the author's alone.

The text was set in Arial for digital and trade editions, with figures executed by the design team at Plexus Digitals working from the figure briefs distributed throughout the chapters. Readers wishing to engage with the figures interactively, or to access the supplementary material — case studies, downloadable worksheets for the Index of Decisions, and the author's continuing notes on how the predictions of the book are aging — may visit the book's companion site.

The author welcomes correspondence from readers, and replies, in the limit, to as much of it as time permits.

APPENDIX D

Bibliography

A formal bibliography of the works most directly informing the arguments of this book. Where Appendix C ("Further Reading") was an opinionated list of further reading for the leader who wishes to deepen her engagement, this bibliography is a more comprehensive scholarly record of the sources, with full citations.

On the Convergence Technologies and Their Trajectories

Acemoglu, Daron, and Simon Johnson. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity.* New York: PublicAffairs, 2023.

Bommasani, Rishi, et al. "On the Opportunities and Risks of Foundation Models." Stanford Center for Research on Foundation Models, 2021.

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* New York: W. W. Norton, 2014.

Christian, Brian. *The Alignment Problem: Machine Learning and Human Values.* New York: W. W. Norton, 2020.

Crevier, Daniel. AI: The Tumultuous History of the Search for Artificial Intelligence. New York: Basic Books, 1993.

Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World.* New York: Basic Books, 2015.

Hao, Karen. Empire of AI: Inside the Reckoning of OpenAI and the Race to Reshape the World. New York: Penguin Press, 2025.

Kai-Fu Lee. AI Superpowers: China, Silicon Valley, and the New World Order. Boston: Houghton Mifflin Harcourt, 2018.

Kurzweil, Ray. The Singularity Is Nearer: When We Merge with AI. New York: Viking, 2024.

Marcus, Gary, and Ernest Davis. Rebooting AI: Building Artificial Intelligence We Can Trust. New York: Pantheon, 2019.

Mitchell, Melanie. Artificial Intelligence: A Guide for Thinking Humans. New York: Farrar, Straus and Giroux, 2019.

Russell, Stuart. Human Compatible: Artificial Intelligence and the Problem of Control. New York: Viking, 2019.

Suleyman, Mustafa, with Michael Bhaskar. The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma. New York: Crown, 2023.

Tegmark, Max. Life 3.0: Being Human in the Age of Artificial Intelligence. New York: Knopf, 2017.

On Quantum Computing

Aaronson, Scott. Quantum Computing Since Democritus. Cambridge: Cambridge University Press, 2013.

Kaku, Michio. Quantum Supremacy: How the Quantum Computer Revolution Will Change Everything. New York: Doubleday, 2023.

Mermin, N. David. Quantum Computer Science: An Introduction. Cambridge: Cambridge University Press, 2007.

Nielsen, Michael A., and Isaac L. Chuang. Quantum Computation and Quantum Information. 10th anniversary ed. Cambridge: Cambridge University Press, 2010.

Preskill, John. "Quantum Computing in the NISQ Era and Beyond." Quantum 2 (2018): 79.

Susskind, Leonard, and Art Friedman. Quantum Mechanics: The Theoretical Minimum. New York: Basic Books, 2014.

On the Metaverse and Spatial Computing

Ball, Matthew. The Metaverse: And How It Will Revolutionize Everything. New York: Liveright, 2022.

Bailenson, Jeremy. Experience on Demand: What Virtual Reality Is, How It Works, and What It Can Do. New York: W. W. Norton, 2018.

Pesce, Mark. Augmented Reality: A New Way of Seeing. San Francisco: Triarchy Press, 2021.

Stephenson, Neal. Snow Crash. New York: Bantam, 1992. [Foundational fictional treatment.]

On the Experience Economy and the Nature of Value

Akpem, Senongo. *Cross-Cultural Design.* New York: A Book Apart, 2020.

Christensen, Clayton M. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail.* Boston: Harvard Business School Press, 1997.

Christensen, Clayton M., Taddy Hall, Karen Dillon, and David S. Duncan. *Competing Against Luck: The Story of Innovation and Customer Choice.* New York: HarperBusiness, 2016.

Goodwin, Kim. *Designing for the Digital Age: How to Create Human-Centered Products and Services.* Indianapolis: Wiley, 2009.

McKelvey, Jim. *The Innovation Stack: Building an Unbeatable Business One Crazy Idea at a Time.* New York: Portfolio, 2020.

Norman, Donald A. *The Design of Everyday Things.* Rev. ed. New York: Basic Books, 2013.

Pine, B. Joseph, II, and James H. Gilmore. *The Experience Economy: Competing for Customer Time, Attention, and Money.* Updated ed. Boston: Harvard Business Review Press, 2020.

Schmitt, Bernd H. *Experiential Marketing: How to Get Customers to Sense, Feel, Think, Act, and Relate to Your Company and Brands.* New York: Free Press, 1999.

On Memory, Attention, and Cognition

Csikszentmihalyi, Mihaly. *Flow: The Psychology of Optimal Experience.* New York: Harper & Row, 1990.

Goleman, Daniel. *Focus: The Hidden Driver of Excellence.* New York: HarperCollins, 2013.

Hari, Johann. *Stolen Focus: Why You Can't Pay Attention – and How to Think Deeply Again.* New York: Crown, 2022.

Kahneman, Daniel. *Thinking, Fast and Slow.* New York: Farrar, Straus and Giroux, 2011.

Newport, Cal. *Deep Work: Rules for Focused Success in a Distracted World.* New York: Grand Central, 2016.

Newport, Cal. *Digital Minimalism: Choosing a Focused Life in a Noisy World.* New York: Portfolio, 2019.

Wu, Tim. *The Attention Merchants: The Epic Scramble to Get Inside Our Heads.* New York: Knopf, 2016.

On Work, Organisations, and Leadership in Transitions

Drucker, Peter F. *The Effective Executive: The Definitive Guide to Getting the Right Things Done.* Rev. ed. New York: HarperCollins, 2006 [orig. 1967].

Edmondson, Amy C. *The Fearless Organization: Creating Psychological Safety in the Workplace for Learning, Innovation, and Growth.* Hoboken: Wiley, 2018.

Frey, Carl Benedikt. *The Technology Trap: Capital, Labor, and Power in the Age of Automation.* Princeton: Princeton University Press, 2019.

Gratton, Lynda. *The Future of Work: How the New Order of Business Will Shape Your Work, Your Career, and Your Life.* London: Palgrave Macmillan, 2011.

Gratton, Lynda, and Andrew Scott. *The 100-Year Life: Living and Working in an Age of Longevity.* London: Bloomsbury, 2016.

Laloux, Frédéric. *Reinventing Organizations: A Guide to Creating Organizations Inspired by the Next Stage of Human Consciousness.* Brussels: Nelson Parker, 2014.

Pink, Daniel H. Drive: The Surprising Truth About What Motivates Us. New York: Riverhead, 2009.

Susskind, Daniel. A World Without Work: Technology, Automation, and How We Should Respond. New York: Metropolitan, 2020.

On India, Civilisation, and Design

Bose, Mihir. India 1500–2050: A Hinterland Becomes a Heartland. London: Hurst, 2024.

Das, Gurcharan. India Unbound: The Social and Economic Revolution from Independence to the Global Information Age. New York: Anchor, 2002.

Das, Gurcharan. The Difficulty of Being Good: On the Subtle Art of Dharma. New York: Oxford University Press, 2009.

Doniger, Wendy. The Hindus: An Alternative History. New York: Penguin Press, 2009.

Khanna, Tarun. Trust: Creating the Foundation for Entrepreneurship in Developing Countries. Oakland: Berrett-Koehler, 2018.

Nilekani, Nandan, and Tanuj Bhojwani. The Art of Bitfulness: Keeping Calm in the Digital World. New Delhi: Penguin India, 2022.

Nilekani, Nandan, and Viral Shah. Rebooting India: Realizing a Billion Aspirations. New Delhi: Penguin India, 2015.

Pai, T. V. Mohandas, and Nisha Holla. India's Moment: Decoding the World's Largest Economic Transformation. New Delhi: Bloomsbury, 2022.

Sen, Amartya. The Argumentative Indian: Writings on Indian History, Culture, and Identity. New York: Farrar, Straus and Giroux, 2005.

Tharoor, Shashi. Why I Am a Hindu. New Delhi: Aleph, 2018.

Varma, Pavan K. The Great Hindu Civilisation: Achievement, Neglect, Bias and the Way Forward. New Delhi: Westland, 2021.

On Risk, Ethics, and the Long View

Bostrom, Nick. Superintelligence: Paths, Dangers, Strategies. Oxford: Oxford University Press, 2014.

MacAskill, William. What We Owe the Future. New York: Basic Books, 2022.

Ord, Toby. The Precipice: Existential Risk and the Future of Humanity. New York: Hachette, 2020.

O'Neil, Cathy. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. New York: Crown, 2016.

Zuboff, Shoshana. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York: PublicAffairs, 2019.

On Climate, Energy, and the Material World

Gates, Bill. How to Avoid a Climate Disaster: The Solutions We Have and the Breakthroughs We Need. New York: Knopf, 2021.

Mann, Michael E. The New Climate War: The Fight to Take Back Our Planet. New York: PublicAffairs, 2021.

Smil, Vaclav. How the World Really Works: The Science Behind How We Got Here and Where We're Going. New York: Viking, 2022.

Smil, Vaclav. Energy and Civilization: A History. Cambridge: MIT Press, 2017.

Wallace-Wells, David. The Uninhabitable Earth: Life After Warming. New York: Tim Duggan, 2019.

On Programmable Biology

Doudna, Jennifer A., and Samuel H. Sternberg. *A Crack in Creation: Gene Editing and the Unthinkable Power to Control Evolution.* Boston: Houghton Mifflin Harcourt, 2017.

Sinclair, David A., and Matthew D. LaPlante. *Lifespan: Why We Age – and Why We Don't Have To.* New York: Atria, 2019.

Venter, J. Craig. *Life at the Speed of Light: From the Double Helix to the Dawn of Digital Life.* New York: Viking, 2013.

On Meaning, Attention, and the Examined Life

Frankl, Viktor E. *Man's Search for Meaning.* Boston: Beacon Press, 1959.

Hadot, Pierre. *Philosophy as a Way of Life: Spiritual Exercises from Socrates to Foucault.* Oxford: Blackwell, 1995.

Kabat-Zinn, Jon. *Wherever You Go, There You Are: Mindfulness Meditation in Everyday Life.* New York: Hyperion, 1994.

Krishnamurti, J. *The First and Last Freedom.* New York: Harper, 1954.

Pirsig, Robert M. *Zen and the Art of Motorcycle Maintenance: An Inquiry into Values.* New York: William Morrow, 1974.

Tolle, Eckhart. *The Power of Now: A Guide to Spiritual Enlightenment.* Novato: New World Library, 1997.

Vivekananda, Swami. *Karma Yoga: The Yoga of Action.* Calcutta: Advaita Ashrama, 1896.

Public-Domain Texts and Foundational Sources

Aristotle. *Nicomachean Ethics.* Trans. Roger Crisp. Cambridge: Cambridge University Press, 2000.

The *Bhagavad Gita.* Trans. Eknath Easwaran. Tomales: Nilgiri Press, 2007.

The *Dhammapada.* Trans. Glenn Wallis. New York: Modern Library, 2007.

Marcus Aurelius. *Meditations.* Trans. Gregory Hays. New York: Modern Library, 2002.

Plato. *The Republic.* Trans. Allan Bloom. New York: Basic Books, 1968.

The *Upanishads.* Trans. Eknath Easwaran. Tomales: Nilgiri Press, 2007.

Reports, White Papers, and Policy Documents

Government of India. "National Strategy for Artificial Intelligence." NITI Aayog, 2018.

Government of India. "National Quantum Mission." Department of Science and Technology, 2023.

International Energy Agency. *World Energy Outlook 2025.* Paris: IEA, 2025.

International Monetary Fund. *World Economic Outlook: The Long Shadow of the Convergence Decade.* Washington: IMF, 2025.

McKinsey Global Institute. *The Economic Potential of Generative AI: The Next Productivity Frontier.* New York: McKinsey, 2023.

National Institute of Standards and Technology. *Post-Quantum Cryptography Standardization.* Gaithersburg: NIST, 2024.

OECD. *AI in Society: An OECD Perspective.* Paris: OECD, 2024.

UN High-Level Advisory Body on Artificial Intelligence. *Governing AI for Humanity.* New York: United Nations, 2024.

World Economic Forum. *The Future of Jobs Report 2025.* Geneva: WEF, 2025.

A Note on Sources

Many of the arguments in this book draw on conversations with practitioners, executives, policymakers, researchers, and operators across the convergence industries — conversations I have had over the past decade in workshops, advisory engagements, board rooms, and off-the-record settings. Where these conversations have shaped specific arguments, I have tried, where possible, to attribute. Where attribution would compromise candour, I have rendered the substance without naming the source. I owe an enormous debt to the people who have spoken with me openly through this period; this book is the better for them, and they should bear no blame for the places where I have got it wrong.

APPENDIX E

General Index

A topical index of the principal terms, names, and concepts that recur substantively in the book. Numbers refer to chapter numbers, not page numbers; chapter numbering is stable across editions and printings.

A

Aadhaar — 18, 45 Abundance, the frontier of — 59, 64 Adaptive products — 25 Adoption curves, compression of — 1, 2 Africa, demographics of — 2 Agentic AI — 5, 21, 28, 30, 31, 58 Aging biology — 16, 53, 59 Agriculture, precision — 59 AI council, the firm's — 31, 32, 44 AI fluency, personal — 31, 47, 64 AI governance — 31, 44 AI literacy — 47 Always-on companions — 21 Ambient computing — 17, 20, 26, 29, 58 Anti-trust frameworks — 34, 44 Apprentice work, decline of — 14, 15, 30 Architecture, of organisations — 14, 32 Argentina — 12 Artificial general intelligence (AGI) — 4, 30, 53 Artificial intelligence — 2, 4, 5, 6, 8, 18, 26, 28, 30, 31, 41, 44, 53, 56, 58 Asset tokenisation — 18, 41 Attention, as currency — 9, 11, 25, 27, 62 Audit, of cryptographic dependencies — 7, 31 Audit, of data foundation — 31, 32 Audit, of risk — 34 Augmented reality — 6, 26, 29 Authenticity, in design — 12, 27 Automation — 5, 14, 30, 38, 46

B

Banking — 1, 18, 35 Bharat (vernacular India) — 12, 27, 33, 47 Bhagavad Gita — 12, 63 Bias, algorithmic — 34, 44 Biology, programmable — 2, 16, 33, 53, 59 Bitcoin — 18 (as critique) Boards of directors — 32, 34, 41 Brand-as-meaning — 8, 10, 27 Brain-computer interfaces — 26, 53 Bureaucracy — 32, 38

C

Capital allocation — 33, 40, 42, 43 Career models — 14, 15, 28, 30 Case studies, the four — 48, 49, 50, 51, 52 Categories, the six new — 19, 20, 21, 22, 23, 24, 25 Chemistry, simulation of — 7, 56, 57, 59 China — 2, 4, 12, 39, 56 Choice fatigue — 8, 35 Civilisations, design from — 12, 33, 54 Climate change — 2, 33, 34, 53, 60 Cognitive load — 9, 58, 62 Companion AI, design of — 21 Compounding, of advantage — 3, 33 Compute infrastructure — 4, 33, 41, 57 Concentration, market — 34, 44 Consultative selling — 36 Continuous health — 16, 49 Convergence, the five-force — 2, 13, 14, 24, 33, 53 Convergence window, the — 2, 31, 53 Cosmology, civilisational — 12 Cost of inference — 4, 27, 57, 58 Crisis, in democratic institutions — 24, 34 Cryptocurrency — 18 (limited) Cryptography — 7, 31, 34, 44, 61 Curriculum, in education — 15, 47 Customer relationships, continuous — 9, 32, 35, 49, 52 Cybersecurity — 34, 39

D

Darshan — 12 Data centres — 33, 41, 57 Data, customer — 9, 21, 23, 25, 32, 44, 61 Data infrastructure, public — 12, 18, 38, 45 Data sovereignty — 33, 45 Defence sector — 39 Deepfakes — 24, 34, 39 Demographic transition — 2, 33, 46, 47 Design language, post-screen — 26, 27, 29 Design principles, the five — 27 Devices, the four categories of — 29, 33 Digital identity — 18, 45 Digital twin — 6, 60 Dignity, in transitions — 30, 38, 46 Discipline, of strategic execution — 31, 32, 33 Dorsey, Jack — implied through Square

examples — 35 Drug discovery — 56, 59 Dystopia, refusal of — 24, 62, 63

E

Economics, of attention — 9, 25 Economics, of trust — 9, 21, 33 Education — 15, 47 Embedded finance — 18, 35, 36 Emergency response — 38 Emotion, in product design — 10, 11, 27 Empathy, irreducibly human — 28, 30, 36, 49 Energy transition — 2, 33, 50, 59, 60 Enterprise architecture — 32 Environmental footprint — 60 Epistemic ground, shared — 24, 34, 44 Ethical commitments — 21, 27, 30, 33, 44, 64 Existential risk — 34, 44, 63 Experience design — 8, 10, 11, 12, 13, 27 Experience economy, the — 3, 8, 12 [N.B. Pine & Gilmore's term, not the book's tagline] Experience stack, the five-layer — 10 Expertise, irreducibly human — 28, 36, 37 Extractive vs generative postures — 30, 33, 64

F

Faith institutions — 51, 63 Fermentation, cellular — 59 Financial services — 18, 35, 41 Five forces, the technological — 2 Five forces, of value migration — 8 Five layers, the experience stack — 10 Fluency, in AI — 31 Food systems — 59, 60 Foundation models — 4, 29, 33, 41, 57 From / through / to, work redefined — 14 Functional silos, dissolution of — 14, 32, 52

G

Generative AI — 4, 5, 22, 24, 27, 58 Generative worlds — 22 Geopolitical fragmentation — 2, 34, 39, 44 Global financial architecture — 41, 43 Governance — 31, 32, 41, 44 Government services — 38, 45 GPUs — 4, 33, 41, 57 Grandparents, in the household — 13 Green hydrogen — 33, 60 Grid, the electrical — 50, 60

H

Habits, of attention — 62, 63, 64 Harvest-now-decrypt-later — 7, 56 Health, continuous — 16, 49 Healthcare — 16, 37, 49 Heat-adaptive cities — 60 Hierarchy, dissolution of — 14, 32 Hippocrates — 16 Hospitals — 16, 37, 49 Human-computer interface — 26, 27, 29

I

Identity, digital — 18, 45 India — 2, 4, 12, 13, 16, 18, 33, 38, 39, 43, 45, 47, 56, 60, 63 Indian National Quantum Mission — 56 Inference, AI — 4, 5, 21, 25, 27, 57, 58 Information environment, the — 9, 24, 34, 62 Infrastructure, digital public — 12, 18, 38, 45 Infrastructure, computational — 33, 41, 57 Innovation Stack, the — 19 Institutions, the responsibility of — 30, 33, 34, 44, 64 Intelligence, unbottlenecked — 58 Intelligence, throughput of — 57, 58 Interface design, post-screen — 26, 27, 29 Internet, decline of free-and-extractive — 9, 21, 61 Investment patterns — 40, 41, 42, 43

J

Japan — 2, 12 Jobs, displacement of — 5, 14, 30, 34, 46 Journey, the customer — 8, 10, 11

K

Kahneman, Daniel — 11, 62 Kashmir, climate vulnerability of — 60 Knowledge work, the future of — 5, 14, 30, 46

L

Labour, the social contract of — 30, 46 Languages, vernacular — 12, 33, 38, 47 Layers, the experience stack — 10 Leadership, the asks of — 30, 31, 32, 33, 64 Leadership, personal cultivation of — 31, 64 Learning, AI-personalised — 15, 47 Liability, for AI actions — 34, 44 Lifespan, healthy — 16, 53, 59 Logistics, AI optimisation of — 36, 50 Longevity — 16, 53, 59 Lucknow — 13, 16, 17

M

Machine learning — 4, 56 Manipulation, AI-mediated — 24, 34, 44, 62 Manufacturing — 2, 32, 60 Marginal cost — 8 Materials science — 56, 59, 60 Meaning, as design layer — 8, 10, 12, 27, 63 Meaning, the question of — 30, 63, 64 Media literacy — 24, 34, 47 Media, synthetic — 24, 34 Medical practice — 16, 49 Memory, as customer metric — 11 Memory, in AI systems — 21, 25 Mentorship — 14, 15, 28 Metaverse, the — 6, 22, 26 Migration, the value — 3, 8, 12 Mindfulness — 62, 63 Misinformation — 24, 34 Mitigation, climate — 60 Mobile devices — 26, 29 Moats, competitive — 9, 21, 33 Moats, in trust — 9, 21, 25, 33 Modular construction — 17 Monetisation, of attention — 9, 24, 62 Money, as commoditising — 9 Money, identity, ownership — 18 Multi-stack architectures — 34 Music — 22, 28

N

National security — 39, 44 Network effects — 8, 21 Neural interfaces — 26, 53 Neuromorphic computing — 57 NIST post-quantum standards — 7 Noise, AI-generated — 24, 62 Norms, professional — 14, 30 Nuclear energy — 60

O

Open-source AI — 4, 33, 39 Operating model, redesigned — 32 Optimisation problems — 7, 50, 56 Organisation, the redesigned — 14, 32 Outcome teams — 14, 32, 52 Outsourcing — 5, 14, 30 Ownership, of assets — 18

P

Pandey, Anil — passim, by way of voice Parents, in the workforce — 13, 30, 46, 62 Patient experience — 16, 37, 49 Payments — 18, 35, 38 Peak-end rule, the — 11 Personal AI — 21, 27, 28, 58, 62 Persuasion, in design — 11, 27 Pharmaceuticals — 37, 56, 59 Photonic computing — 57 Physical world, integration with — 6, 17,

20, 22, 25, 29 Pilots, AI — 31 Platform decisions — 29, 32, 33 Platforms, dominant — 6, 29, 33, 34 Policy, public — 30, 34, 38, 44, 45, 46, 47 Population, demographic dynamics of — 2, 46 Post-quantum cryptography — 7, 31, 34, 56, 61 Post-screen design — 26, 27 Power, concentration of — 30, 34, 44 Preparedness, organisational — 31, 32, 33, 34 Privacy — 9, 21, 23, 27, 44, 61 Procurement, ecosystem-style — 32 Productivity, AI-driven — 5, 26, 30, 32, 41 Progress, the question of — 30, 64 Programmable biology — 2, 16, 33, 53, 59 Prompt engineering — 4 Prosoche — 63 Protein folding — 56, 59 Provable truth — 24

Q

Qualifications, traditional — 15, 28, 30, 47 Quantum computing — 2, 7, 33, 53, 56, 57, 60, 64 Quantum-human synthesis — 64 Quantum simulation — 56, 59, 60 Quantum supremacy — 7, 56

R

Re-architecture, the six — 32 Real estate — 17, 48 Real estate, sentient — 17, 20, 48 Redundancy, dignified — 30 Refusal, of bad business — 30, 44, 64 Regulation — 33, 34, 44, 45, 46 Relationships, the substance of life — 13, 62, 63, 64 Renaissance, the privacy — 61 Renewable energy — 50, 60 Reskilling — 30, 46 Restoration, of the planet — 60, 64 Retail — 36 Re-statement, of category — 3, 16, 35, 36, 52 Risk, the seven categories of — 34 Robots — 28, 30 RSA cryptography — 7, 56

S

Safety, AI — 30, 34, 44 Scarcity, refusal of — 59, 64 Scenarios, climate — 34, 60 Schools — 15, 47 Screen, the recession of — 26, 27, 29 Search, the decline of — 9, 27, 58 Sectoral analyses — 35, 36, 37, 38, 39, 50 Self-driving vehicles — 29 Sensors, distributed — 17, 20 Sentient spaces — 20 Service economy, the — 3, 8, 30 Sequencing, of re-architectures — 32 Six bets, the — 33 Six categories, the new

— 19 Six re-architectures, the — 32 Skills graphs — 14, 28, 30, 47 Smartphone, the receding — 26, 29 Social contract, the — 30, 34, 46 Software-as-a-service — 5 Solar energy — 33, 50, 60 South Korea — 2, 39 Spatial computing — 2, 6, 22, 26, 29, 33 Stack, the experience — 10 State, the — 38, 44, 45, 46, 47 Stoicism — 63, 64 Strategy, in convergence era — 31, 32, 33, 40, 41, 42, 43 Stress, the cognitive condition of — 9, 58, 62, 63 Substrates, computational — 7, 56, 57, 58 Supply chain, the — 39, 50 Sustainability — 33, 50, 60 Synthesis, the quantum-human — 64 Synthetic biology — 53, 59 Synthetic content — 24, 34

T

Talent, in convergence era — 14, 15, 28, 30, 32, 47 Technologies, the five — 2 Tectonics, the three macro — 2 Throughput, the classical ceiling on — 57 Throughput, of intelligence — 57, 58 Time, as currency — 9, 25 Timeline, of cryptographic migration — 7 Tipping points, climate — 60 Token economics — 18, 41 Trade-offs, civilisational — 30, 60, 63, 64 Traditions, civilisational — 12, 63 Transitions, dignified — 30, 38, 46 Transitions, the four obligations — 30 Transitions, the six re-architectures — 32 Transparency — 9, 21, 24, 33, 44, 61 Trust, as currency — 9, 21, 25, 33, 35, 49, 51, 61 Trust capital — 33, 51, 52 Trust dossier — 9, 21 Truth, provable — 24 Tuesday in 2030, a — 13, 14, 15, 16, 17, 18, 62 Twin, digital — 6, 60

U

Uncertainty, the texture of — 62, 63 Underwriting — 35 Unemployment, technological — 30, 34, 46 Universities — 15, 47 Upskilling — 30, 46 Urbanisation — 17, 60 User-centred design — 27

V

Valuation, of firms — 41 Value migration, the great — 3, 8 Value, the new unit of — 8 Vehicles, autonomous — 29 Vernacular India — 12, 27 Video games — 6, 22 Virtual reality — 6, 22, 26 Voice interfaces — 27, 29

W

Water systems — 60 Wearables — 16, 17, 26, 29 Wellbeing, the question of — 62, 63 Wisdom, of the elders — 12, 13, 28 Work — 5, 14, 28, 30, 31, 46 Work, the future of — 14, 30, 46 Work, the meaning of — 30, 63, 64 Workforce, the redesign of — 14, 30, 32, 46 Worlds, generative — 22

X / Y

(no entries)

Z

Zero-knowledge proofs — 18, 24, 61

APPENDIX F

Supplements

A small set of practical supplements to support the leader applying the arguments of this book inside her own institution.

Supplement One: A Reading Order for Different Readers

This book has been written to be read cover to cover. The specific reader, however, may find a different order more useful.

The CEO under acute time pressure — Read Chapter 1, Chapter 8, Chapter 30, and Chapters 31 through 34. This is the strategic spine in roughly 35,000 words. Return to the rest as time permits.

The CFO or capital-allocator — Begin with Chapters 1 through 3, then jump to Part IX (Chapters 40 through 43). The middle of the book provides the substrate for understanding why the financial implications are what they are; treat the middle as appendix to the financial argument.

The CHRO or head of people — Begin with Chapters 1 through 3 for the framing, then read Chapter 14, Chapter 15, Chapter 28, and Chapter 30 as the human-capital spine. Chapter 46 (labour and reskilling policy) gives the broader context.

The CMO or product leader — The product spine runs through Chapters 8 through 11 (the new value equation), Chapters 19 through 25 (the new categories), and Chapters 26 through 29 (post-screen design). Read these first; the rest of the book gives the strategic context.

The CIO or CTO — Begin with Chapters 4 through 7 (the technology forces), then move to Chapter 32 (re-architecture) and Part XIII (Chapters 56 through 64) for the deep technical-and-philosophical arc.

The policy reader — Part X (Chapters 44 through 47) is the principal target, with Part XII (Chapters 53 through 55) for the longer view. The earlier chapters provide the substrate.

The investor or board member — Part IX (Chapters 40 through 43), Part XI (Chapters 48 through 52), and the closing chapters of Part XII (Chapters 54 and 55) are the concentrated argument. The rest of the book is the supporting structure.

The reader interested in the human dimensions — Chapter 28, Chapters 62 and 63, and the Epilogue. The technological and strategic chapters provide context but are not, for this reader, the primary substance.

Supplement Two: A Timeline of the Convergence Decade

A reference timeline of significant inflections expected through the convergence decade. The dates are approximate and represent the median of credible forecasts. The reader is welcome to disagree with the specifics; the timing matters less than the relative ordering of the events.

2024

NIST finalises post-quantum cryptography standards. The formal beginning of the cryptography migration window.

ChatGPT-class generative AI, broadly familiar to most working professionals, has been in commercial availability for two years.

The first commercial-grade quantum systems with several hundred qubits are operational at multiple firms.

2025

Agentic AI, capable of multi-step task execution with reasonable reliability, becomes commercially deployable for narrow use cases.

The first major regulated industries (banking, defence) begin formal post-quantum cryptography migration programmes.

The cost per AI inference has fallen by approximately three orders of magnitude relative to 2020.

2026

Agentic AI is in production deployment for end-to-end internal processes in early-adopter enterprises.

The first commercial AR glasses suitable for all-day use are available, although adoption remains in the early-adopter range.

Several countries (India among them) launch substantial public-funding programmes for domestic AI and quantum infrastructure.

2027

AI-augmented drug discovery has produced its first molecules to reach phase-three clinical trials with significantly compressed timelines.

The first major industries (banking, healthcare, retail) report material productivity gains attributable to agentic AI deployment, in the 15–25% range at the function level.

Banking-sector deadlines for post-quantum cryptography migration take effect in major regulators' frameworks.

2028

AR glasses begin to overtake smartphones in daily-use minutes for early-adopter cohorts.

The first generative-world platforms reach 100 million active users.

Inference cost has fallen another order of magnitude. AI is, for the first time, economically deployable for continuous ambient presence.

2029

The first cryptographically relevant quantum advantage is, on the more aggressive but credible forecasts, achieved on real public-key cryptography. Migration urgency becomes acute.

The first AI-discovered drugs reach commercial market.

Personal AI companions, with multi-year memory, become commercially available at consumer scale.

2030

AR glasses reach a substantial fraction (estimated 30–45%) of the smartphone install base.

The white-collar labour market has materially restructured. Several professional categories have lost 30–60% of pre-2024 headcount; new categories have emerged that did not previously exist.

The first widely deployed generative worlds for education, healthcare, and entertainment are operational at scale.

2031

Quantum-augmented materials science produces commercially deployable batteries with energy density meaningfully beyond 2024 maxima.

The first end-to-end designed therapeutics for several previously intractable conditions reach commercial deployment.

The cryptographic migration is, for major regulated industries, substantially complete.

2032

The world's first widely deployed direct air capture facilities reach the cost of $100 per tonne of carbon.

AI-augmented climate modelling produces operational forecasts at the resolution required for serious adaptation planning.

The cost per AI inference has fallen approximately seven orders of magnitude relative to 2020. Inference is, in the practical sense, unbottlenecked for most applications.

2035

Programmable biology, AI-augmented, has produced cell-based meat and dairy products at price parity with conventional production in major markets.

Climate adaptation infrastructure investment globally exceeds two trillion dollars per year.

The first wave of brain-computer interface products begins entering early consumer markets, with significant ethical and regulatory contestation.

2040

The convergence decade has substantially produced its inheritance. The institutions of 2040 reflect the choices made in the 2025–2035 period. Whether the inheritance is judged generative or extractive depends on those choices.

Supplement Three: A Leader's Diagnostic Worksheet

The seven questions below are intended for a CEO, leadership team, or board to work through deliberately, perhaps over a half-day or full-day offsite. Each question requires honest consideration; none of them has a textbook answer. The cumulative effect of the seven is a baseline understanding of where the firm

sits in the convergence and where its decisions over the next three years should concentrate.

Question 1. What is the experience our customer is actually buying from us — beyond the product or service we describe in our marketing materials?

The honest answer is rarely the marketing description. The discrepancy is the strategic opportunity.

Question 2. Where, on the five-layer experience stack, does our principal product currently compete? Where does our most threatening new entrant compete? Where, in three years, do we intend to compete?

The first answer is descriptive. The second answer reveals threat. The third answer is the strategic commitment.

Question 3. What proportion of our customer's attention, time, and trust do we currently command — and which of these currencies are we systematically depleting in pursuit of revenue?

Most firms deplete one of the three. The depletion is the long-term liability.

Question 4. What is our honest plan for the categories of work in our firm that AI is reshaping or replacing? Have we communicated it openly to the workforce? Have we funded reskilling at the level the situation requires?

If any of the answers is no, the firm is not yet ready for the workforce transition that the next three years will ask of it.

Question 5. Of the six re-architectures (data, operating model, customer relationship, talent, ecosystem, board), which have we begun? Which have a credible owner? Which are we, by default, postponing?

The default postponements are the future regrets.

Question 6. Of the six strategic bets (own the customer's primary AI; infrastructure and capacity; programmable biology; energy and climate; Bharat-first design; trust as institutional capital), which have we explicitly committed to? Which are we depending on others to make available to us?

The dependencies are the future vulnerabilities.

Question 7. When the firm is judged in 2035 by the workforce that depended on it, the customers it served, and the planet that hosted it — what verdict do we want to be the one we receive?

The answer to this question shapes the answers to the previous six.

Supplement Four: A Glossary of Quantum Terminology

A short technical glossary for the leader who wishes to engage more deeply with the quantum chapters of this book without acquiring the full physics vocabulary.

Coherence time. The duration over which a quantum system maintains the superposition states that make quantum computation possible. Longer is better. Engineering longer coherence times is a central challenge of quantum hardware.

Cryptographically relevant quantum computer. A quantum computer of sufficient scale and reliability to break currently deployed public-key cryptography (specifically RSA and elliptic-curve cryptography) at realistic key sizes. The threshold is hundreds of thousands of physical qubits with current error-correction approaches.

Decoherence. The loss of quantum information from a quantum system through interaction with its environment. The principal practical obstacle to large-scale quantum computation.

Error correction, quantum. The set of techniques by which fragile physical qubits are combined into more reliable logical qubits, at the cost of using many physical qubits per logical qubit. Different

schemes have different ratios; the surface code, currently a leading approach, requires roughly a thousand physical qubits per logical qubit at production fidelity.

Hybrid quantum-classical computing. Computational approaches that use quantum systems for the parts of a problem they handle well and classical systems for the parts they handle better. Most early commercial quantum applications are hybrid.

Logical qubit. A reliably error-corrected qubit suitable for sustained computation. The unit by which quantum computational capability is meaningfully measured. The first quantum computers with hundreds of logical qubits will be capable of consequential applications.

NISQ era. Noisy Intermediate-Scale Quantum era. The current period of quantum computing, in which hardware is too small and too error-prone for reliable execution of the famous quantum algorithms but large enough to begin commercial experimentation.

Physical qubit. The actual physical system that stores a quantum bit of information — a superconducting circuit, a trapped ion, a photon, a neutral atom, depending on the hardware approach. Physical qubits in current systems are too unreliable for sustained computation without error correction.

Post-quantum cryptography. Cryptographic algorithms designed to be secure against attack by quantum computers. The first standards were finalised by NIST in 2024. Migration to post-quantum cryptography is the principal cybersecurity work of the late 2020s.

Quantum advantage. The empirical demonstration that a quantum computer has solved a problem faster than any reasonable classical computer could. Has been demonstrated for narrow contrived problems since 2019; not yet for problems of clear commercial value, but expected within the decade.

Quantum simulation. The use of a quantum computer to simulate the behaviour of quantum-mechanical systems (typically molecules, materials, or condensed matter). The most likely first commercial application of quantum computing at meaningful scale.

Qubit. Quantum bit. The fundamental unit of quantum information. Unlike a classical bit, which is either 0 or 1, a qubit can exist in a superposition of both states until measured.

Shor's algorithm. The 1994 algorithm that demonstrated, theoretically, that a sufficiently large quantum computer could factor large integers efficiently, breaking the RSA cryptography that protects much of the internet. The driver of the post-quantum cryptography migration.

Superposition. The quantum-mechanical property whereby a quantum system exists in multiple states simultaneously until measured. The basis of quantum computation's parallelism.

Supplement Five: A Note on the Decisions

The 165 decisions distributed through the book and gathered in Appendix B are intended to be acted upon. The leader who treats them as interesting reading rather than as a checklist of commitments will, in the convergence decade, find herself at a disadvantage relative to the leader who works through them deliberately.

A practical approach: take a single quarter to work through five or six of the decisions, choosing those most relevant to the firm's current strategic position. Document the actions taken and the results observed. Then take the next quarter to work through another five or six. Over six quarters, a firm can credibly act on approximately a third of the 165 decisions; over the full three years that this book treats as the strategic horizon, on roughly half of

them. This is not lack of ambition; this is honest pacing for a leadership team that also has the day job to run.

The decisions that remain unactioned at the end of three years are not failures. They are the decisions whose moment had not yet come, or whose alternatives the firm chose to take instead. The discipline is the deliberate engagement with the question rather than the universal commitment to the answer. The leaders I have worked alongside who have applied this approach over a serious period have, almost universally, reported that the discipline of the framework matters more than the specific decisions taken — that the act of working through the questions, with the leadership team, has changed the texture of how the firm thinks more than any specific decision has changed how it operates.

This is the substance of the book applied. It is offered, with all the humility I can muster, as a working aid for the leaders whose contribution to the convergence decade will, in the end, determine its shape.