

TEACHING THE NEXT INDIA

A Policy-and-Practice Handbook for
the AI Education Decade

A N I L P A N D E Y



BlueRoseONE^{.com}
S t o r i e s M a t t e r
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.



BlueRose ONE^{.com}
S t o r i e s M a t t e r
New Delhi • London

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

BLUEROSE PUBLISHERS

www.BlueRoseONE.com

info@bluerosepublishers.com

+91 8882 898 898

+4407342408967

ISBN: 978-81-69761-92-5

Cover design: Aafiya Saifi]

Typesetting: Namrata Saini

First Edition: July 2026

*For the teachers who will do this work,
and for the children who will benefit from it.*

A Note to the Reader



On names, numbers, and the nature of the argument.



This book is a work of argument, not of reportage. The reader should know, before proceeding, what that distinction means for the specific material she is about to encounter.

The named individuals who appear in the book's storyline — Sunita, the Class 3 child in Semra village; Kusum Devi, her teacher; Renu Balakrishnan, the Kerala mathematics teacher; Moirangthem Brojeswor, the Manipuri physics teacher; Lakshmi, the seven-year-old in Umaria district; and any other named teacher, student, parent, or official who appears in the narrative — are composites. They are drawn from my conversations and observations across many Indian classrooms over several years, but no individual named in this book corresponds to a single identifiable person in the real world. The names, the specific districts, and the biographical details have been chosen to make the argument concrete without exposing any real individual to the attention a book of this circulation may bring. If the reader recognises, in these composites, a teacher or a child she knows, the recognition is a sign that the composite has done its representational work, not that any specific person has been placed on these pages. Any resemblance to a specific named individual is, in the strict sense, coincidental.

The named institutions — NCERT, NCTE, NIEPA, the SCERTs, the DIETs, the Block Resource Centres, the Ministry of Education, MeitY, Bhashini, the CBSE, the state boards, NITI Aayog, UIDAI, the National Health Authority, and the various statutory bodies and ministries referenced throughout — are real. The descriptions of their mandates, their operational realities, and

their institutional cultures reflect my understanding as of the writing of this book in early 2026. I have made every reasonable effort to describe them accurately. Where I have characterised their performance or their tendencies in unflattering terms, I have done so on the basis of the public record and of conversations with officers who have served within them, and I accept responsibility for the characterisations. Where institutional behaviour has changed since my characterisation was written, the reader should update her understanding accordingly.

The numerical estimates offered throughout the book — fiscal envelopes, teacher workforce counts, device unit costs, content production volumes, per-school deployment costs, the Equity Fund's proposed quantum, the compensation increments, and the various other figures that populate the argument — are estimates. They are built up from assumptions I have tried to make as reasonable as the available evidence permits, but they are not the product of a detailed budgetary exercise of the kind a Cabinet note would require. They are intended to establish orders of magnitude and to give the policy reader a specific point of reference from which her own more rigorous analysis can begin. The reader should treat every specific number in this book as a starting point for refinement rather than as a conclusion. Where I have offered a point estimate, the honest confidence interval is usually wider than the presentation implies. I have tried to flag the most consequential uncertainties in the text. The reader who proposes to act on any of the numbers should commission her own analysis before doing so.

The institutional proposals — the National AI Education Mission, the National AI Education Content Repository, the National AI Teacher Institute, the Equity Fund, the Children's AI Safety Code, the Rapid Response Unit, the DPI platform for AI education, and the various other structures the book proposes — are my proposals. They do not exist at the time of writing. They reflect my best judgement about what the reform requires, offered for contestation, refinement, and, I hope, eventual adoption in

some form. The reader who has worked in Indian public administration will recognise that the proposals are offered with the knowledge that the final shape of any such institution is determined by political and administrative negotiation that this book cannot anticipate. I have tried to make the proposals concrete enough to be useful to that negotiation without pretending that my specifications will or should survive it unchanged.

The political characterisations in the book are unavoidable, because a reform of this scale has political dimensions that cannot be separated from its substantive design. I have tried, throughout, to present the political economy of the reform in terms that readers across the Indian political spectrum can engage with. Where I have expressed views on specific political choices, I have tried to do so on the basis of argument rather than of partisan preference. The reader will judge whether I have succeeded. I hold, in any case, that the reform's case does not depend on its being the initiative of any particular political formation; the reform is, in my view, a national undertaking whose success or failure will matter to Indians regardless of the political balance at the moment of its undertaking.

Finally, the book was written at a specific moment in the development of artificial intelligence as a technology, and its specific technical claims are time-bounded in ways that a non-technical book typically is not. The capabilities of AI tutors, the cost trajectories of the underlying models, the quality of Indian-language performance, the specific safety concerns that deserve regulatory attention — all of these will evolve substantially across the decade the reform is meant to occupy. I have tried to write in a way that the book's argument survives the specific technical evolution, by grounding the argument in the stable questions of pedagogy, institution, equity, and delivery rather than in the unstable answers of the current technical moment. The reader who encounters this book in 2030 or 2035 should treat its technical characterisations as a snapshot of early 2026 and should update them against the reality of her own moment.

With those disclaimers registered, the rest of the book proceeds on its argument.



Endorsements

Col. Lakshmikant Tiwari

“Teaching the Next India is not another book on Artificial Intelligence. It is a national blueprint that connects policy, technology, classrooms, teachers, institutions, and governance into one actionable framework. Anil Pandey combines the discipline of an engineer, the vision of a strategist, and the practicality of an execution leader to present a roadmap that every educator, policymaker, and institution head should study. This book deserves to become a reference for India’s AI education journey.”

Dr. Sarvesh Goel

Chairman, G.D. Goenka Public School, Lucknow & The Centrum Hotel

“Education is no longer about preparing students for examinations; it is about preparing them for a rapidly changing world. Teaching the Next India recognises this reality and provides a balanced framework that keeps teachers at the centre while embracing Artificial Intelligence responsibly. Anil Pandey has transformed a complex national challenge into an implementable action plan. Every educational institution serious about the future should read this book.”

Shri S. K. Jaggi

Chief Operating Officer, Eldeco Group

“India needs more than vision – it needs execution. Teaching the Next India stands out because it bridges that gap. It combines policy thinking with operational clarity, offering institutions a practical framework to adopt AI without losing sight of human values and educational outcomes. Anil Pandey’s experience in large-scale transformation is evident throughout the book, making it an important contribution to India’s education ecosystem.”

Dr. Vivek Tangri

“Artificial Intelligence will redefine education, but technology alone cannot transform learning. It requires visionary leadership, empowered teachers, institutional readiness, and a clear roadmap for execution. Teaching the Next India successfully brings these dimensions together into a practical national framework. Anil Pandey has written a timely and insightful guide that can help educational institutions confidently prepare learners for the opportunities and challenges of the AI era.”

Dr. Ashok Tiwari

“Teaching the Next India presents a rare blend of educational vision, policy insight, and technological pragmatism. Rather than treating AI as a future possibility, Anil Pandey provides educators and policymakers with an actionable roadmap for transforming learning at scale. This book is a significant contribution to India’s educational transformation and deserves a place on the desk of every institution leader committed to preparing the next generation.”

Contents

A Note to the Reader	v
Preface	xiii
Acknowledgements	xvii
List of Figures and Tables	xix
Introduction	xxi
PART I: The Promise & The Gap	1
Chapter One: The Class 3 in Barabanki	3
Chapter Two: What the NEP Promises, and What It Does Not	13
PART II: The Classroom & The Teacher	25
Chapter Three: The Grade-by-Grade Roadmap	27
Chapter Four: The Content Question – Who Writes What Our Children Will Read	46
Chapter Five: The Technology Question – Devices, Platforms, and Not Being Captured	68
Chapter Six: The Teacher Question	88
Chapter Seven: The Equity Question	119
Chapter Eight: A Letter to Parents	135
PART III: Making it Real at Scale	153
Chapter Nine: The Institution Head's Playbook	155
Chapter Ten: Building the Mission – A Statutory Proposal ..	174
Chapter Eleven: Paying for It – The Fiscal Architecture	197
Chapter Twelve: The Delivery Chain – From Shastri Bhavan to the Classroom	216
PART IV: The Execution Layer	235
Chapter Thirteen: Outcomes and Accountability	237
Chapter Fourteen: Risks and Failure Modes	248

Chapter Fifteen: The Phased Rollout	263
PART V: The Legal, Comparative, and Human Frame.....	275
Chapter Sixteen: The Legislative and Regulatory Layer	277
Chapter Seventeen: What India Should and Should Not Copy	286
Chapter Eighteen: The Manifesto — Returning to Sunita.....	299
APPENDICES	311
Appendix A: Glossary of Terms	312
Appendix B: Index of Decisions.....	316
Appendix C: A Model AI-Literacy Curriculum Specification ..	321
Appendix D: A Principal's Hundred-Day Checklist.....	324
Appendix E: Further Reading	326
Bibliography	328
Thematic Index.....	333
About the Author.....	340

Preface



Why this book, and why now.



I began writing this book in late 2024, in the specific form it now takes, after a sequence of classroom visits across five Indian states that made a particular idea unignorable to me. The idea was not that artificial intelligence would transform Indian education – that claim was already commonplace, and I would have been one of many voices restating it had I written a book on that theme. The idea was narrower, and I believe more useful. It was that the specific question of whether AI would reach the Indian child who most needs it – the child in the government primary school in the weaker-capacity state, the child whose home language is not Hindi or English, the child whose teacher has never held a tablet – was not going to be answered by the unfolding of the technology alone. It was going to be answered by specific institutional choices, made at specific moments, by specific people whose attention was, at the time of writing, still being claimed by other priorities. The book I wanted to write was a book for those specific people, made specific enough that it could actually influence what they decided.

The book has passed through four drafts over the fourteen months of its composition. The first draft was an essay, perhaps thirty thousand words, structured around the Class 3 child in Barabanki whose question still opens this book. The second draft was an attempt to turn that essay into a policy handbook, and it was, as policy handbooks often become, too long and too dry for anyone outside the specific policy audience to read. The third draft attempted to rescue the handbook from its dryness by restoring the

first draft's narrative register, and the result was a book that was neither a satisfying narrative nor a usable handbook. The fourth draft, which the reader now holds, tries to be both — a book with a narrative spine that a thoughtful general reader can follow from beginning to end, and a book with enough operational specificity that the Joint Secretary, the Chief Minister's Principal Secretary, the SCERT Director, and the institution head can find in it the specific artefacts their work requires. Whether the compromise succeeds is for the reader to judge. I have tried.

I should say, because the question has been asked of me by friends who read earlier drafts, why I have written this book at all rather than continued the direct practitioner work from which the book's material is drawn. The answer is that I came, gradually, to the view that the specific gap this book tries to fill was not being filled by the practitioner conversation on its own. The teachers, the institution heads, the state officers with whom I work regularly have views on these questions that are often more nuanced than the national policy conversation reflects, but their views are not being aggregated into a form that the national conversation can engage with. This book is my attempt to do some of that aggregation. The reader will form her own view of whether the attempt was worth making.

I should also say, in the interest of honesty, what the book is not. It is not a work of original research. It does not contain new empirical findings about AI-education outcomes in Indian classrooms, because those findings are properly the work of the evaluation researchers whose methods are more rigorous than mine. It is not a technical book about artificial intelligence — there are excellent books on that subject, and I have contributed in a small way to at least one of them. It is not a comprehensive survey of Indian education, a task that would have required several volumes of substantially greater scope than this one. It is, specifically, a book about one question, at one moment, in one country: how the AI-education reform India is now beginning

should be designed so that it reaches the children for whom it is ostensibly being undertaken. The narrowness of the question is deliberate. The alternative, of writing a book that attempted to address every adjacent question, was a book that would have addressed none of them well.

The book's first and most important debt is to the teachers who allowed me into their classrooms and who shared with me, across hundreds of hours of conversation, their specific experience of the work the reform will now ask them to extend. Kusum Devi, Renu Balakrishnan, Moirangthem Brojeswor — the composite names by which they appear in the book — reflect real conversations with real teachers whose names I protect for the reasons I have specified in the preceding disclaimer. What I have tried to honour, in the fictional naming, is the substance of what they taught me. The honour is imperfect, and the teachers themselves would recognise the imperfection. The attempt is the book.

The book's second debt is to the officers across the central and state administrations who spoke with me, often on the condition of confidentiality that a book of this kind must respect, about the specific institutional realities within which the reform will operate. Where the book's institutional analysis is accurate, the accuracy is theirs. Where it is inaccurate, the inaccuracy is mine. The officers I have relied on most extensively know who they are, and I hope that when they read these pages they find their own thinking reflected in the argument, even when my conclusions diverge from theirs. The divergence is the privilege of a book; the alignment, where it exists, is the debt.

The book's third debt is to the children themselves. The argument is, at its centre, an argument about what they deserve, and the specific child who stands behind the abstractions on every page is the child the argument is for. If this book contributes to her getting what she deserves, it will have earned the years it took to write. If it does not, I hope another book soon will.

I have tried to write clearly. Where I have failed, the failure is mine, and I apologise in advance for the specific effort the reader will have had to make to extract the argument from my sentences. Where I have succeeded, the success belongs, as all such successes do, to the editors and early readers whose patience with the earlier drafts shaped the version you now hold.

— Anil Pandey, Lucknow early 2026



Acknowledgements



A book of this kind accumulates debts on a scale that no acknowledgements page can fully discharge. What follows is an attempt at the principal ones.

To the teachers, institution heads, and officers across the states of Uttar Pradesh, Kerala, Manipur, Tamil Nadu, Karnataka, Madhya Pradesh, Odisha, Bihar, West Bengal, and several others whose hospitality across many classroom visits made this book possible, my thanks. You know, I hope, who you are. The specific people I would wish to name by name, I cannot, for the reasons registered in the disclaimer. The debt is not reduced by the anonymity of its acknowledgement.

To the officers of the Ministry of Education, of NCERT, of NIEPA, of the Department of School Education and Literacy, of MeitY, of NITI Aayog, and of the various state education departments who gave me their time in conversations that were often longer than their diaries could accommodate, my thanks. Several of you disagreed with parts of what I have written, and your disagreement improved the book in ways I could not have produced on my own.

To the evaluation researchers, academic specialists, civil-society practitioners, and technology-sector colleagues whose work has shaped my thinking across the years this book has been in composition, my thanks. The intellectual debts are too numerous to itemise, but a reader familiar with the field will recognise, in the book's arguments, the specific influences that shaped them. Where I have failed to credit a specific contribution adequately in the text, I apologise, and I commit to correcting the omission in any subsequent edition.

To the readers of the book's earlier drafts — who will not be named here, either, because the honesty of private early-reader feedback depends on the promise that it will not be publicly attributed — my particular thanks. The book is substantially better than the draft you first encountered. The improvements reflect your specific comments, even where I have not explicitly incorporated every suggestion. The residual faults are, in the conventional and accurate formulation, entirely my own.

To my editors and the production team who carried the book from manuscript to the form in which you hold it, my thanks. Your patience with a schedule that consistently underestimated the work remaining was more generous than the schedule deserved.

To my family, whose tolerance of the specific form of absence that writing requires has been, across fourteen months, both substantial and unreciprocated. The book is dedicated to teachers and children; the writing of it was made possible by the people at home, and the unwritten dedication is theirs.

A final word. The specific project this book describes — the reform it argues for — is larger than any single book, any single author, any single institution. If the project is undertaken, its success will be the work of millions of people whose names will not appear anywhere in its public record. If the book contributes, in however modest a way, to that undertaking, the contribution will have been worth the labour.



List of Figures and Tables



Chapter Two

Figure 2.1 | What the NEP says and does not — the three operational gaps the reform must close.

Chapter Three

Figure 3.1 | The grade-by-grade roadmap — three stages, with focus, strands, and expected outcomes at each.

Chapter Four

Figure 4.1 | The National AI Education Content Repository — commissioning, origination, adaptation, delivery.

Chapter Five

Figure 5.1 | The DPI architecture — a thin public layer, private innovation on top, no vendor lock-in.

Chapter Six

Figure 6.1 | Six components of teacher preparation — each substantial, none reducible to a workshop.

Figure 6.2 | The teacher-capacity delivery chain — from national institute to the classroom.

Chapter Nine

Figure 9.1 | The institution head's three horizons — foundations, building, scale.

Chapter Ten

Figure 10.1 | National AI Education Mission — statutory architecture.

Chapter Eleven

Figure 11.1 | The fiscal architecture — where the money goes, where it comes from, how it flows.

Chapter Twelve

Figure 12.1 | The seven-layer delivery chain — and the three disciplines that shorten it.

Chapters Thirteen — Fifteen (Part IV)

Figure 13.1 | The Seven National Outcomes — Metric, Cadence, Accountable Institution.

Figure 13.2 | The Six-Level Metric Framework — Classroom through National.

Figure 14.1 | Ten Failure Modes — Early-Warning Indicator and Pre-emptive Mitigation.

Figure 15.1 | The Phased Rollout — Years One through Five.

Figure 15.2 | Capex and Opex by Year — the fiscal envelope's composition.

INTRODUCTION

How to Read This Book



A short orientation for the policy reader, the practitioner, and the general reader.



This book tries to do three things at once, and the reader is entitled to a short orientation, before she begins the main argument, about which of the three she most wants and how to find it.

For the policy reader — the Joint Secretary drafting a Cabinet note, the Principal Secretary in a state department of education, the officer in the Prime Minister's Office or a Chief Minister's Office who will hold the reform politically — the book's most directly useful sections are Chapter Ten (on the statutory mission), Chapter Eleven (on the fiscal architecture), Chapter Twelve (on the delivery chain), and Part IV in its entirety — Chapters Thirteen (on outcomes and accountability), Fourteen (on risks and failure modes), and Fifteen (on the phased rollout). These six chapters, read in sequence, constitute a substantially complete picture of the institutional machinery the reform requires. The earlier chapters provide the pedagogical context that justifies the machinery. If time is short, begin with Part IV and work backward; if you have the time to read in order, the argument builds more naturally forward from Chapter One.

For the practitioner reader — the institution head, the SCERT director, the DIET principal, the teacher who has responsibility for her colleagues' professional development — the most directly useful sections are Chapter Three (the grade-by-grade roadmap),

Chapter Six (on teachers, the book's centre of gravity), Chapter Nine (the institution head's playbook), and the appendices, particularly Appendix D (the principal's hundred-day checklist). These sections offer specific operational guidance that is, I hope, immediately usable. The parent reader should go directly to Chapter Eight, which is a letter addressed to her, and which can be read on its own without the rest of the book.

For the general reader — the thoughtful Indian citizen who wants to understand what the AI-education reform is likely to involve and what she should make of it — the book is meant to be read in sequence from beginning to end. The narrative thread follows a specific child from Chapter One to Chapter Fifteen, and the argument develops cumulatively. Chapters that are operational in their emphasis are placed in the book's second half; the first half establishes the frame. A general reader who finds specific chapters too operational for her interest can skim them — the book is not written on the premise that every reader will engage with every chapter at the same depth.

The book has three parts. Part I (Chapters One and Two) establishes the problem: the specific gap between what Indian policy promises on AI education and what will actually reach the child who most needs the promise honoured. Part II (Chapters Three through Eight) addresses the classroom — what should be taught, how, with what content, using what technology, by which teachers, under what equity discipline, and with what engagement from the child's parents. Part III (Chapters Nine through Twelve) addresses the machinery through which the classroom changes are produced at national scale — the institution head's work, the national Mission, the fiscal architecture, and the delivery chain. Part IV (Chapters Thirteen through Fifteen) specifies the execution layer — the outcomes the reform will be held to, the risks it must guard against, and the phased rollout through which it is delivered. Part V (Chapters Sixteen through Eighteen) closes the argument — the legislative and regulatory layer, the international

comparisons, and the manifesto that returns to the child with whom the book began.

The book uses some recurring conventions the reader should recognise. At the end of most chapters, a section called "Three Decisions for the Reader" specifies the concrete actions the chapter's argument invites, addressed to different constituencies. The decisions are not a substitute for the argument; they are a handle by which the argument can be grasped quickly when time is short. Where specific numerical estimates or institutional proposals appear, they should be understood within the limits registered in the disclaimer at the front of the book. Where specific named individuals appear in the narrative, they should be understood as composites, also as registered there.

A final orientation. The book is long, and it is specific. Both of these are deliberate. A shorter book on this subject would have had to be less specific, and the specificity is, in my judgement, the book's most useful feature for the reader who actually needs to act. If the length seems excessive to the reader whose purpose the book does not directly serve, I apologise, and I encourage her to use the book selectively rather than comprehensively. The argument's integrity does not depend on any single reader reading every chapter.



PART I



The Promise & The Gap

CHAPTER ONE

The Class 3 in Barabanki

*A child, a question, and the distance between
Delhi and the classroom she sits in.*



Sunita is eleven years old. She is in Class 3 at the government primary school in Semra village, Barabanki district, Uttar Pradesh. She should, by the prescription of the National Education Policy and every framework document that descends from it, be able to read Hindi fluently, perform the four operations of arithmetic, and hold a conversation about the natural world around her. She can do none of these things. Her teacher, Kusum Devi, who qualified in 2011 and has taught at this school for nine years, knows this. Her mother, who never went to school past Class 5, suspects it. The District Education Officer, whose monthly returns show Semra's Class 3 at ninety-one percent attainment of the foundational literacy benchmark, does not.

I begin this book with Sunita because she is the reason the book exists. Every sentence in the pages that follow is, at some level, about her. Every policy argument, every curricular prescription, every piece of operational machinery I describe for the Centre or the states is, ultimately, measured against a single test: does it change what happens in Sunita's classroom on a Monday morning, or does it not?

The reform this book proposes — a national programme to teach artificial intelligence to every Indian schoolchild from Class 3 through Class 12, supported by a statutory mission, an adequate fiscal envelope, and a delivery chain that reaches every district — will succeed or fail on whether it reaches Sunita. Not on whether it

reaches the urban private schools whose children would have been fine anyway. Not on whether it reaches the Kendriya Vidyalayas or the Navodayas or the well-provisioned CBSE schools of the Tier 1 metros. On whether it reaches the Semras of Uttar Pradesh, the Chhapras of Bihar, the Khowais of Tripura, the Koraput of Odisha. On whether it reaches the child for whom this reform is ostensibly being undertaken, and not merely the child who was already going to benefit from whatever the country did next.

I should say clearly, now, what I believe. I believe the Indian state is serious about this reform. I have sat in enough meetings in Shastri Bhavan and in Yojana Bhavan to know that the people designing it are thoughtful, overworked, and in many cases unusually committed to its success. I do not write from cynicism. I write from the particular discomfort of a practitioner who has watched earlier education reforms — some excellent in their design — fail to reach the Sunitas of this country not because anyone wanted them to fail but because the machinery of translation, between Delhi's intention and Barabanki's classroom, was never adequate to the weight placed on it.



I. What this book is, and what it is not

This is a book about how to make the Indian state's AI-education reform actually work — for the schoolteacher who will deliver it, for the institution head who will host it, for the official who will fund and supervise it, for the parent who will live with its consequences in her own home, and, centrally, for the child it claims to serve. It is a handbook. It is an argument. It is, in places, a manifesto. It tries not to be any one of these at the expense of the others.

It is not a book that will tell you what artificial intelligence is, or what it can do, or why you should be excited about it. There are excellent books for that, and I have contributed in a small way to at

least one of them. This book assumes that the reader has formed, from reading and experience, some general picture of the technology — and that she has arrived at the question that matters next, which is: what should we, specifically, in India, be doing to prepare our children for the world this technology is creating?

It is also not a book that will tell you what to think. A policy handbook of this kind must respect the reader's own judgement. What I offer is a framework, a set of specific recommendations with their reasoning, and the evidence on which those recommendations rest. Where I believe something strongly, I say so, and I distinguish belief from fact. Where the evidence is contested, I name the contest. Where my recommendation would be costly or politically difficult, I acknowledge the cost rather than smuggling it in with softer language. A book that pretends its recommendations are free, or that the people who disagree with them are foolish, is a book that should not be trusted.



II. Who this book is for

I have written this book with one reader specifically in mind. She is the head of an educational institution in India — a school principal, a college dean, a training-institute director, a vice-chancellor — somewhere between forty-five and fifty-five years old, running an organisation of between eight hundred and four thousand students. She has been in her role for at least five years. She reads seriously but not constantly, and has perhaps three hours a week to spend on a book of this kind. She has heard about AI in education at every conference she has attended in the last two years. She has sat through at least three demonstrations of AI tutors by polished sales teams. She has, by the end of each demonstration, felt both intrigued and quietly overwhelmed, and she has, in each case, left the room without knowing what specifically she should do next.

She is not waiting for the government. She has been in Indian education long enough to know that waiting for the government on questions of this kind is a form of managed resignation. She wants to act, within her own institution, with whatever autonomy her role permits her. But she wants to act with judgement, not with panic, and she has, so far, been unable to find a source of guidance that is neither a sales pitch nor a policy paper written for somebody else.

This book is written for her. It is also, secondarily, written for three other readers. For the officer of the Indian state — a Joint Secretary, an additional secretary, a chairperson of a state education board, an adviser at NITI Aayog — who is responsible for some portion of the national reform and who needs a serious practitioner's view of what will and will not work in the classrooms her policies are meant to reach. For the founder or head of an organisation that builds, funds, or supports educational programmes — an EdTech company, a philanthropic foundation, a corporate CSR arm — who needs to understand where her work fits within the national architecture. And for the parent of a school-aged Indian child, who has been watching the headlines and has begun to worry that something important is being decided about her child's education without her being asked. I write a chapter directly to her later in the book. I hope she finds it useful.

If you are none of these four readers — if you are, say, a school teacher yourself, or a researcher in education policy, or an engaged citizen whose children are grown — you are still welcome here. Much of the book will be useful to you. But I have tried, in writing it, to remember that the institution head is my primary reader, and when I have had to choose between serving her and serving another audience, I have chosen her. This choice is, I believe, also the right choice for the country. Institution heads, acting with autonomy and good judgement in their own organisations, are the shortest path from policy to classroom that India's administrative

geography offers. A reform that mobilises them acts quickly. A reform that goes around them moves at the speed of Delhi.



III. What I believe, stated plainly

Before a reader invests her time in a book of this kind, she deserves to know where its author stands. Policy books that pretend to neutrality are, in my experience, the most likely to smuggle in their author's assumptions unexamined. I would rather state mine openly, so the reader can weigh them as she reads.

I believe artificial intelligence is the most consequential technology Indian children will encounter in their lifetimes, and that the current generation of Indian adults — parents, teachers, policymakers, institution heads — has roughly a decade in which to prepare those children for a world this technology will reshape. I believe this window is real and not a matter of fashionable urgency.

I believe the Indian state, despite its many limitations, is uniquely positioned to undertake a national AI-education reform of genuine ambition — more so than most of the states often cited as models. I believe this because of a specific combination of assets India possesses that are rare in the global picture: a federal constitutional framework that forces real engagement with scale; a Digital Public Infrastructure tradition that has already solved problems of identity, payment, and document-sharing at population scale and is therefore capable of solving comparable problems in education; a young population whose demographic weight gives a reform of this kind outsized consequence; and a tradition of public-spirited administrative practice that, whatever its well-documented failures, still produces, in each generation, a cohort of officers capable of extraordinary work when they are given the authority and the protection to do it.

I believe, also, that the reform will fail if it is treated as a technology programme rather than as an educational programme. The temptation to do the opposite — to lead with devices, with platforms, with vendor partnerships, with the visible furniture of modernity — is considerable. It has killed earlier reforms. It will kill this one too unless the people designing it hold the discipline that what we are attempting is, at its core, not a technology deployment but a pedagogical transformation, which happens to use technology. The distinction sounds philosophical. It is, in practice, the difference between a reform that reaches Sunita and one that does not.

I believe the teacher is the centre of the reform, not its instrument. This is the argument on which the whole book rests, and I will not belabour it here because the sixth chapter belabours it at sufficient length. I name it now only so that the reader knows, from the beginning, what this book is ultimately about.

I believe equity is not a separate chapter of the reform but its overall test. A reform that raises the educational ceiling for India's top quartile while the bottom quartile either stagnates or falls further behind is not a reform; it is an aggravation of an existing injustice under a new vocabulary. I will say more on this in Chapter Seven. I name it here because equity is the single lens through which I ask the reader to hold every other argument in the book. If an idea I propose would widen the gap rather than close it, I have failed my own test, and I would like the reader to hold me to it.

I believe, finally, that this reform is winnable. I hold this belief with some effort, against a fair number of contrary experiences, and I cannot offer the reader the grounds for optimism as proof. What I can offer is the specific observation that Indian education, despite the volume of its failures, has also produced remarkable reforms in my own lifetime — the Right to Education Act, the Mid-Day Meal Scheme, the Sarva Shiksha Abhiyan, the foundational literacy push under Nipun Bharat — each of which seemed unlikely before it succeeded. None of these reforms is complete. All

of them have produced outcomes that would have been dismissed, a decade before their success, as implausibly ambitious. The AI-education reform belongs in the same tradition of possibility. Whether it will belong in the same tradition of achievement is the open question that the rest of this book is an attempt to answer.



IV. How this book is organised

The book has three parts. The first part, of which this chapter is the opening, is short. It establishes the stakes and interrogates what the National Education Policy of 2020 says and does not say about artificial intelligence in schools. It runs to two chapters and is, in a sense, the ramp that earns the reader's attention for the harder work that follows.

The second part is the heart of the book. It contains six chapters that work through what an AI-education reform actually means inside a classroom — what children should learn, what content should be written for them, what technology should support them, what teachers need to make it real, what equity demands of the reform, and what parents should do at home while the institutions around them catch up. This second part is the longest section of the book. It is also, I hope, the most useful for the reader who has picked this volume up in order to decide what to do next in her own institution or her own home.

The third part is about scale. It takes the argument of the second part and asks what the national machinery must look like for the reform to reach every Indian child, not merely those whose institutions or parents are unusually capable. It contains chapters on the statutory design of the mission I propose, on fiscal architecture, on the delivery chain, on legal and regulatory questions, on what India should and should not copy from other countries' attempts at similar reforms, and, at the book's close, a chapter that is both a manifesto and a call to action.

Each chapter ends, as the final act of its work, with three decisions I invite the reader to take in her own context. The decisions are specific. They are framed so that they can be acted on within ninety days of the reader closing the book. They are not, in most cases, decisions that require the permission of a higher authority to execute. They are, for the most part, decisions that any serious reader can take on her own initiative, in her own institution, in her own home, starting on the Monday morning after she finishes the chapter. The three decisions at the close of the fifteenth chapter are the ones the book's argument ultimately rests on. I hope, by the time the reader reaches them, she will find them less a prescription than a clarification of what she already knew she wanted to do.



V. A word on how I came to write this

I write this book from Lucknow, where I have lived most of my adult life, and where I run a technology business whose work includes, among other things, the training of educators in artificial intelligence and allied subjects. I mention the business because the reader is entitled to know the author's vantage, and because some of the observations in this book are informed by what I have seen in classrooms and training rooms across North India over the past decade. I do not mention the business further in the pages that follow. The book is not about it. The book is about a national reform whose success depends on tens of thousands of institutions of which mine is an inconsequentially small one, and whose argument stands or falls on grounds unconnected to any specific organisation's work.

The book draws on field interviews conducted specifically for it — with teachers, principals, administrators, parents, and students across at least six Indian states — as well as on a wider body of observation and reading accumulated over years. Where I quote a teacher or a parent by name, I have secured her consent.

Where I describe a composite — the Sunita of this chapter is a composite, as I will disclose whenever it matters — I have drawn it from real people and real classrooms in ways that, I hope, are faithful to the country they inhabit, even where the individual is not herself real.

I have tried to write a book that would be useful to a reader who does not agree with all of it. I have my own views, strongly held in places, and I state them openly when they are mine. But a book that demands agreement is a book that is of limited use to a country whose education will be reformed by thousands of people holding a wide range of views, who must nonetheless find enough common ground to work together for the next two decades. I have tried, therefore, to write in a way that a reader who disagrees with my central recommendations could still find use in the observations that surround them. Whether I have succeeded is, like much else in this undertaking, a matter for the reader to judge.



Three decisions for the reader

1. Name the child this reform is for, in your own mind.

Before you close this chapter, give her a name. Picture her. Place her in a school you know. If you are a principal, picture a child in your own institution. If you are an officer, picture a child in the district you were last posted to. If you are a parent, picture your own child. Every time, in the chapters that follow, I use a phrase like the reform will succeed if, or the policy should require, measure what I have said against her. If a recommendation would not improve her education, it is a recommendation that has failed the book's own test, and you are entitled to reject it. The argument of this book is that specificity is the discipline that saves reforms from becoming rhetoric. Begin that discipline now, with one named child, before you read any further.

2. Identify the decision you are currently avoiding.

Somewhere in the institution you run, or the family you lead, or the office you administer, there is a decision about artificial intelligence in education that you have been avoiding for six months or longer. You know what it is. It is the decision you have told yourself you are not yet ready to take, because you are waiting for more information, or for a clearer policy signal, or for a better vendor to emerge, or for a moment when you will have more time. Write it down, in one sentence, before you continue reading. You do not have to take the decision yet. But name it. The rest of this book is written, in part, to give you the information and framework you need to take it. We will return to the sentence you wrote down, implicitly, many times in the chapters that follow.

3. Decide what you would need to read, in this book, to act.

Make a short list — three items, no more — of what a book on this subject would have to give you in order for you to change what you do in your institution, your office, or your home. Be specific. Not clarity on the curriculum, but a grade-by-grade roadmap I could hand to my curriculum coordinator on Monday. Not understanding of the policy landscape, but a brief on what the NEP actually commits the Centre to. Not evidence on what works, but two or three case studies of Indian schools where something comparable is already happening. Write your list down. Keep it to hand as you read. When you close this book, return to the list and judge the book against it. A policy-and-practice handbook is, at its core, a document whose worth is measured by whether the reader can act on it when she has finished reading. I would rather you held me to that test than to any other.



END OF CHAPTER ONE

CHAPTER TWO

What the NEP Promises, and What It Does Not



A close reading of the document on which everything else depends.



The National Education Policy of 2020 runs to sixty-six pages in its official English version. I have read it perhaps a dozen times, at varying distances — once as an interested citizen when it was published, several times since as a practitioner trying to work out what it actually requires of people like me, and most recently, in preparation for this book, as a reader trying to understand what it says about artificial intelligence specifically, and what it leaves for others to say after it. The document repays close reading. It is thoughtful, in places even visionary, and it was produced by a drafting process that included more serious educators than any comparable Indian policy document in my lifetime. I begin this chapter with that acknowledgement because what follows is, in part, a critique, and a critique has more weight when it comes from a reader who respects the document she is critiquing.

The NEP mentions artificial intelligence directly in seven places. I have counted. The mentions are scattered across sections on curriculum, on higher education, on vocational training, and on the use of technology in learning. Taken together, they signal an awareness that AI will matter to Indian education over the coming decades. Taken separately, each mention is brief, and none of them specifies what the state or any of its instruments must actually do as a consequence of the mention. The distance between those two

facts — between the signalled awareness and the absent specification — is the space this book is written into.

Let me take the mentions in turn, because the reader is entitled to see the text before she trusts my reading of it.



I. What the NEP says

In paragraph 4.27, the NEP states that coding will be introduced from Class 6 onwards. This is the single most widely quoted AI-adjacent commitment in the document, and it has driven a considerable amount of subsequent activity, including CBSE's revised curriculum and several state boards' adoption of coding modules for Classes 6 through 8. The paragraph is important and welcome. It is also silent on what comes before Class 6, what comes after Class 8, what teachers will teach the coding, what devices the children will code on, what the relationship of coding is to the larger question of AI literacy, and what the state is financially committing to make the provision real. These silences are not criticisms of the NEP. A policy framework is not the place where such specifications appear. But they are the specifications this book concerns itself with.

In paragraph 4.23, the NEP speaks of integrating cutting-edge areas such as artificial intelligence into the school curriculum at the secondary level, alongside holistic, integrated, enjoyable, and engaging learning. The paragraph treats AI as one of a list of contemporary subjects — alongside environmental awareness, financial literacy, and the three Rs — that the curriculum should address. It does not treat AI as a transversal competence that reshapes how every other subject is taught. It does not treat AI as a cognitive posture that children should develop the way they develop numeracy. It treats AI as a topic. I believe this is an insufficient framing, and I will make the argument for a different

framing in Chapter Three. But the reader should know what the NEP actually says before reading what I wish it had said.

In paragraph 11.8, the policy speaks of the use of technology — including AI — to support personalised learning, adaptive assessment, and intelligent tutoring. This is the paragraph closest to the reform this book is about. It acknowledges, correctly, that AI has potential to transform the learning experience. It is, inevitably, brief. It does not specify how such technology should be procured, what standards it must meet, how it should be regulated, what data it may collect, or how it should be paid for. Again, these are not criticisms. They are specifications the NEP leaves for other documents to supply. The trouble is that, five years after the NEP's publication, those other documents largely do not exist in the form the reform requires. They are the documents this book, in its third part, attempts to sketch.

In paragraphs 20 through 24, which address higher education, the NEP commits to the creation of new institutional capacity for AI research and education, including the establishment of a National Research Foundation and the promotion of multidisciplinary institutions that integrate AI with other fields of study. These are welcome commitments. They are also, as the reader will notice, commitments about the higher-education end of the system. They do not bear directly on the school-education reform this book is about, though they create the pipeline of AI-capable graduates who will, in the decade to come, become some of the teachers and officers whose absence the school system currently feels.

In paragraph 23.5, the NEP commits to the establishment of a National Educational Technology Forum as an autonomous body to provide independent evidence-based advice on technology in education. The Forum has been constituted. It has produced useful work. It does not, at the time of writing, possess the authority or the fiscal envelope to serve as the coordinating mission the reform this book proposes requires. It is the kind of body that might, with

enhanced mandate, become that mission; or it is the kind of body that a differently constituted mission would work alongside. I will return to this question in Chapter Ten.

And in several places across the document, the NEP speaks of the need for teacher capacity development in the use of technology, including AI, without specifying what that development entails, who is responsible for it, what it costs, or how its quality is assured. These mentions are honest in their recognition of the centrality of the teacher. They are, as we will see in Chapter Six, insufficient in their specification of what serious teacher preparation for an AI-enabled classroom requires.

Those are the direct AI mentions. I have listed them in full, or close to it, so that the reader can judge for herself whether my account is fair. A reader who wishes to check me against the original document is welcome to do so; I have not strained any interpretation.



II. What the NEP implies, without stating

A document's silences can be as instructive as its speech. The NEP's silences on AI fall into three categories, and each category shapes what the reform needs from the documents and decisions that must follow it.

The first category is specification. The NEP establishes direction; it does not specify magnitude. When the policy says coding will be introduced from Class 6, it does not say at what weekly pedagogical intensity, for how many years of a child's school career, delivered by teachers of what qualification, assessed against what standards. When the policy speaks of AI at the secondary level, it does not say whether AI is an elective, a compulsory subject, an integrated thread across other subjects, or a set of competencies to be demonstrated irrespective of how they are acquired. The specification work must be done somewhere —

by NCERT, by CBSE, by the state boards, by the SCERTs, or by a coordinating body that does not yet exist. Where it has been done, it has been done unevenly. The consequence is that two schools in two neighbouring states that each describe themselves as NEP-compliant in their AI provision may, on inspection, be teaching entirely different things to entirely different effect.

The second category is financing. The NEP is not a budget document. It does not state the quantum of resources required to implement its provisions, nor does it specify which provisions are first-order priorities in the event of a fiscal squeeze. The policy's headline commitment — that public education spending should rise to six percent of GDP — has not been met, and the gap between the aspiration and the actual allocation constrains everything downstream. Specifically for the AI-education reform, the absence of a dedicated financial line — whether in the Union Budget, in Samagra Shiksha, or in a ring-fenced mission allocation — means that even the provisions the NEP does specify are implemented unevenly across states, dependent on each state's own fiscal room and political will. I will return to the fiscal architecture in Chapter Eleven. For now it is enough to note that the NEP's silence on costing is structural rather than incidental.

The third category, and the most consequential, is coordination. The NEP describes a system in which the Centre, the states, the statutory bodies, and the private and philanthropic sectors all have roles to play in its implementation. It does not specify how those roles are to be coordinated, who holds the coordinating authority, and what happens when the coordination fails. In the absence of coordination, each actor does what its own mandate suggests it should do. NCERT develops curriculum. CBSE issues circulars. States implement — or do not — within their own political and financial constraints. Private schools operate with substantial autonomy. Philanthropies fund what they find interesting. The aggregate effect is activity, but it is activity that does not compound into reform. This is the coordination gap a

statutory mission is designed to close. I will propose the statutory mission in Chapter Ten.

These three gaps — specification, financing, coordination — are the operational deficit the reform must close. Each can, in principle, be closed without amending the NEP. The NEP is a framework; the reform is the detailed construction within the frame. I want to be clear that I am not, in this book, calling for the NEP to be rewritten. I am calling for the operational documents and institutional decisions that a working reform requires, and which the NEP's own architecture anticipates, to be produced with greater seriousness and specificity than has so far been achieved.




Figure 2.1 | What the NEP says and does not — the three operational gaps the reform must close.



III. Why these gaps matter for the reader of this book

I have dwelt on the NEP at this length because I want the reader, particularly the institution head or the officer, to understand what she is and is not operating within. She is operating within a framework that gives her broad direction and

substantial latitude. She is not operating within a detailed specification that tells her exactly what to do. In many contexts, Indian administrative tradition treats this state of affairs as a reason to wait for further guidance — to defer substantive action until the Centre, the state, or the statutory body produces the specification that the institution then implements.

I want to make the opposite argument. The fact that the NEP does not specify at a granular level what AI education looks like in Class 3 or Class 8 is not a reason to wait. It is, in the present moment, licence to act. An institution head who reads the NEP as permission to develop her own AI-education provision, consistent with the policy's direction but specified by her own judgement and her institution's context, is an institution head acting entirely within the policy's spirit. She is, moreover, the kind of actor whose work will eventually inform the national specification that is still to come. The Centres and the state boards learn from the institutions that have already begun. They do not learn from the institutions that are still waiting to be told what to do.

This is the pragmatic argument for the handbook this book attempts to be. If the specification is going to emerge from the field as much as from the Centre, then the book that equips a principal to produce a thoughtful, institution-specific specification is a book that contributes directly to the national reform. It is not a book that merely interprets the NEP for her. It is a book that invites her to become a participant in the NEP's continued elaboration.

I want to name one caution here, because the invitation is not without risk. An institution head who proceeds without the guard rails that a national specification provides is an institution head who may, with good intentions, deploy technology that harms the children in her care. She may buy AI tutors that collect data she should not permit them to collect. She may adopt content that reflects biases she has not examined. She may introduce screen time at a developmental stage where the evidence is still unsettled. The absence of specification is licence to act; it is not licence to act

carelessly. The chapters that follow are written, in part, to give the reader enough specificity that her action can be informed rather than merely enthusiastic. I would rather the reader close this book more cautious than she opened it, and still acting, than close it more confident and still waiting.



IV. What the Centre has done since 2020

A fair account of the reform must note what has actually been attempted in the five years since the NEP's publication, rather than treating the policy landscape as if it were static. The Centre has not been idle. I offer a compressed ledger, with neither the flattery of the departmental briefing nor the cynicism of the reflexive critic.

NCERT has developed a new National Curriculum Framework for School Education, completed in 2023, which takes the NEP's direction and elaborates it at greater specificity than the NEP itself supplies. The framework speaks of computational thinking, of data literacy, and of AI literacy in ways that bear directly on the reform this book is about. It is a significant document and is underutilised in the current national conversation. A serious institution head would be well served by reading it alongside this book.

CBSE introduced an Artificial Intelligence subject at the secondary level in 2019, ahead of the NEP, and has continued to refine it since. The subject is currently offered at over five thousand CBSE schools and has trained a first generation of secondary-level AI teachers. CBSE's work is the closest existing precedent for what a national reform at scale might look like, and it has accumulated institutional learning that the reform should draw on.

The Atal Tinkering Labs initiative, originally launched in 2016 and accelerated under NITI Aayog's Atal Innovation Mission, has established making-and-coding infrastructure in over ten thousand schools. The labs are of uneven quality — some are

exceptionally well-run, others are storerooms with untouched 3D printers — but the initiative has produced, in its best form, a proof of concept for what a national AI-education rollout could look like at the infrastructure layer. The Mission's own evaluation literature is honest about the gaps. It is worth the reader's time.

The India AI Mission, launched in 2024 under the Ministry of Electronics and Information Technology, has allocated approximately ten thousand three hundred crore rupees over five years for AI capability development across the country, including elements of skill development and education. Its education-specific allocation is smaller than the scale of the reform this book describes would require, and its institutional home in MeitY rather than the Ministry of Education reflects a federal structure in which AI education does not yet have a single bureaucratic owner. I will say more on this in Chapter Ten.

Bhashini — the India Language Technology Mission — has made substantial progress in building AI infrastructure for Indian languages, and this infrastructure is one of the indispensable components of any AI-education rollout that takes seriously the linguistic diversity of the country's schoolchildren. The Mission's work has advanced faster than any equivalent initiative in any country of comparable linguistic complexity. I will return to what it enables in Chapter Four.

PM-SHRI Schools, a Centrally Sponsored Scheme launched in 2022, has designated approximately fourteen thousand schools across the country as showcase institutions for NEP implementation, including technology integration. The scheme has had mixed results, varying substantially by state, and it has not yet become the vehicle for national AI-education rollout that its initial design anticipated it might. Whether it can be repurposed for that role, or whether the reform requires a different delivery vehicle, is a question I will address in Chapter Twelve.

This is a creditable ledger, taken as a whole. It is not, taken as a whole, a reform of the scale or coherence the moment demands. The individual efforts are real; they do not yet add up to a national programme. The argument of this book is that they can be made to add up — not by abandoning them, but by bringing them into a coordinated architecture with clear accountability, adequate financing, and a delivery chain that reaches every district. The ledger above is, in effect, the raw material from which the reform must be constructed.



V. What this chapter has done, and what the next thirteen chapters will do

I have tried, in this chapter, to establish three things. The first is that the National Education Policy of 2020 is a serious document that names AI-in-education as a matter of consequence without specifying what the specification must look like in practice. The second is that the gap between the policy's direction and its specification is the operational space the reform must occupy, and that this space is being occupied unevenly, without sufficient coordination, and without a coherent fiscal architecture. The third is that an institution head, reading the NEP today, is licensed to act within the spirit of the policy even in the absence of detailed central specification — and that the quality of her action will depend on whether she has the judgement and specificity this book tries to give her.

The chapters that follow take up, in turn, the components of that specification. Chapter Three addresses the curricular question: what should Indian children actually learn about AI at each grade from Class 3 to Class 12. Chapter Four addresses the content question: who writes the materials these children will encounter, in which languages, at what cost, to what standard. Chapter Five addresses the technology question: what devices and platforms

schools need, and how to procure them without being captured by vendors. Chapter Six – the longest and most important chapter of the book – addresses the teacher question, which is the question on which everything else depends. Chapter Seven addresses equity, which is the lens through which the whole reform must be held to account. Chapter Eight is written directly to parents. Chapter Nine is a practical playbook for the institution head. Chapters Ten through Twelve address the national machinery – the Mission, the money, and the delivery chain. Chapters Thirteen through Fifteen form the execution layer – outcomes and accountability, risks and failure modes, and the phased rollout. Chapters Sixteen through Seventeen address the legal framework and what India should and should not copy from abroad. Chapter Eighteen returns to Sunita.

By the time the reader closes this book, I hope she will feel equipped to do three things. To interpret, with judgement, the policy documents and circulars she will receive in the years to come. To act, within her own institution or office, with specificity informed by a national perspective. And to hold both the Centre and herself accountable to a standard higher than the one the current conversation is settling for.



Three decisions for the reader

1. Read the NEP yourself, or reread it.

The document is sixty-six pages long. It is available free in English and in Hindi on the Ministry of Education's website. No analyst's summary, including this chapter, is a substitute for the reader's own encounter with the primary text. If you are an institution head, reserve two evenings in the next two weeks to read it with a pen in hand. Mark what bears on your institution.

Mark what is silent. You will read the remaining chapters of this book with a sharper judgement if you have done so.

2. Name the single NEP commitment you believe your institution is most under-delivering against.

Every institution I have ever worked with has at least one NEP commitment on which its current practice falls conspicuously short of the policy's expectation. For some institutions it is foundational literacy. For others it is the integration of vocational education. For others it is teacher autonomy. Identify yours. Write it down. If you cannot name one, it is likely because you have not read the NEP closely enough recently, and decision one above is where *you should begin*.

3. Decide whether your institution's AI-education provision will be centre-led or institution-led.

You have a choice, and it is worth making consciously rather than by default. You can wait for the Centre or your state board to produce a detailed specification and implement it when it arrives, perhaps two to four years from now. Or you can develop your institution's own specification, consistent with the NEP's direction, informed by your own context, and refine it as central guidance emerges. Both paths are defensible. The second path takes more work and requires more judgement. It is also the path that produces the best institutions of the next decade, because it is the path that takes the NEP at its word when it invites institutional autonomy. Decide which path is yours, and say so to your senior leadership team within the next month. The rest of this book will be considerably more useful to an institution head who has made this decision than to one who has not.



END OF CHAPTER TWO

PART II



***The Classroom &
The Teacher***

CHAPTER THREE

The Grade-by-Grade Roadmap



What every Indian child should learn about artificial intelligence, from Class 3 to Class 12.



I want to begin this chapter with a concession, before I make the argument it contains. Curricular prescription at a national scale is the kind of work that should, properly, be done by committees of educators with deep pedagogical training, in consultation with children, teachers, subject specialists, and the institutions that will implement the result. This chapter is not such a document. It is one practitioner's considered view of what a reasonable grade-by-grade roadmap might look like, offered as a starting point for the serious curricular work that the NCERT, the state SCERTs, and the National AI Education Mission I propose in Chapter Ten will need to undertake. I write it because the absence of even a provisional roadmap in the current national conversation leaves institution heads, curriculum coordinators, and parents unable to make specific decisions. A provisional specification, openly offered for contestation, is more useful than no specification at all. I invite disagreement, particularly from educators with more classroom experience than I have. A better roadmap than this one will improve the book and, more importantly, improve the children it ultimately serves.

With that concession registered, let me make the argument the chapter is built on. The argument is that AI education, rightly understood, is not a new subject to be added to an already crowded school timetable. It is a set of cognitive capabilities — analogous to,

though distinct from, literacy and numeracy — that every child of the coming generation needs to develop, at grade-appropriate depth, across the span of her school career. A reform that treats AI as a new subject will produce the characteristic failure mode of the Indian school curriculum, which is to load the child with content she cannot integrate with her other learning and which she forgets within a year of her examination. A reform that treats AI as a capability to be developed across the curriculum, using dedicated curricular time where appropriate but primarily by reshaping how existing subjects are taught, will produce a child who is genuinely prepared for the world she is growing into.

With that framing, let me now offer the roadmap itself.

I want to note, before the roadmap, one further framing principle that shapes its design. The roadmap is cumulative, in the sense that each stage builds on the one before it, and the failures of an earlier stage compound into difficulties at the later stages. A child who leaves Class 5 without fluent literacy in her mother tongue will find the middle-stage work of reading charts and evaluating AI-generated text substantially harder than a fluent peer. A child who leaves Class 8 without comfort in the three strands of the middle stage will find the secondary-stage electives, and particularly the Applied AI track, inaccessible in ways the elective itself cannot remedy. The pedagogical discipline that follows from cumulative design is that each stage must genuinely deliver its outcomes for the substantial majority of children before the subsequent stage's ambitions can be pursued. A school that is meeting perhaps half of its Class 5 literacy outcomes cannot, with integrity, claim to be running a successful middle-stage AI curriculum; the failures at the earlier stage will, with near certainty, produce inadequate outcomes at the later one. The discipline is uncomfortable because it asks institutions to measure and report on the foundational outcomes honestly rather than moving quickly to the more visible and politically attractive middle-stage work. A reform that respects cumulative design is a reform that sometimes

says to an eager institution head — build the foundation first, we will come to the rest when it is ready — and that is prepared to hold the line against the institutional pressures that push in the opposite direction.




Figure 3.1 | The grade-by-grade roadmap — three stages, with focus, strands, and expected outcomes at each.



I. The three stages

I organise the roadmap around the three stages of school education that the NEP's 5+3+3+4 structure establishes, modified by the pragmatics of Indian classroom reality. The Foundational Stage — Classes 1 and 2, and the pre-primary years that precede them — is the stage at which, in my view, no direct AI instruction of any kind should take place. The cognitive development that matters at that stage is the development of language, of motor skills, of social play, and of the child's relationship with the physical and natural world around her. Introducing screen-based AI tools at this stage is a distraction from what the child most needs, and the evidence on the harms of early screen exposure, while not yet settled, is sufficient to counsel caution. The

Foundational Stage is the stage at which the reform takes the form of not doing AI education, and doing thoroughly the literacy, numeracy, and socio-emotional work that is already prescribed.

The first stage I address in detail, therefore, is what I will call the Foundational-to-Preparatory Stage — Classes 3 through 5, covering ages approximately eight to eleven. This is the stage at which, I believe, AI enters the child's educational experience for the first time, primarily as a tool that supports foundational literacy and numeracy, and secondarily as an object the child develops her first grounded understanding of.

I choose the name Foundational-to-Preparatory deliberately, because the NEP's taxonomy uses Foundational for Classes 1 and 2 only, and Preparatory for Classes 3 through 5. My framing bridges the two categories because the transition between them is, in the Indian reality, not as clean as the NEP's taxonomy suggests. Many children who are nominally in Class 3 are still, in terms of their actual competence, at a Class 1 level on the NEP's developmental scale, and a curricular specification that assumes they have completed the foundational stage cleanly will fail them. The bridging framing is meant to hold this reality in view — the reform addresses children across a genuine range of starting points, not children who have all reached the same developmental threshold on the schedule the taxonomy implies.

The second stage is the Middle Stage — Classes 6 through 8, covering ages approximately eleven to fourteen. This is the stage at which AI enters the curriculum as a subject in its own right, one period a week, through a dedicated curriculum organised around three strands: computational thinking, data literacy, and AI literacy proper. It is also the stage at which AI tools become deeply integrated into the teaching of every other subject.

The third stage is the Secondary Stage — Classes 9 through 12, covering ages approximately fourteen to eighteen. This is the stage at which AI education becomes specialised, with three optional

electives drawing different kinds of students toward different futures, and at which AI tools become genuinely a research partner for the child's own academic work.

I will take each stage in turn.



II. Classes 3 to 5 – the foundational stage of AI education

The child at the start of Class 3 is, in the best case, a fluent reader of her mother tongue, a confident performer of the four operations of arithmetic, and a curious observer of the world around her. In the Indian reality, the median child at the start of Class 3 is none of these things. The ASER surveys have established, over nearly two decades of careful measurement, that approximately half the children in India's government primary schools cannot read a Class 2 text at the start of Class 3, and a similar proportion cannot perform basic arithmetic reliably. Any serious proposal for AI education at this stage must begin by acknowledging that foundational literacy and numeracy are the preconditions for everything else, and that the child who lacks them will gain nothing from AI exposure except a reinforcement of her sense that school is a place where things happen that she does not understand.

The AI-education provision I propose for Classes 3 through 5, therefore, has foundational literacy and numeracy as its primary purpose. The AI in the classroom is there, above all, to help the child read fluently in her mother tongue and in English, and to become confident in the four operations of arithmetic. The device through which this happens is an AI tutor – a conversational agent, on a tablet or a shared school device, operating in the child's mother tongue, capable of reading with her, listening to her read, gently correcting her pronunciation, and giving her as much patient practice as she needs without ever running out of patience.

This is the use of AI in education for which the strongest evidence already exists, and the use most likely to produce measurable gains in learning outcomes at the stage where those gains matter most.

Alongside the use of AI as a tutor, Classes 3 through 5 should introduce the child to AI as an object — not a technical object, at this stage, but a recognisable presence in her world about which she forms her first accurate understandings. She should learn, by the end of Class 5, three specific things about the machine she is interacting with. She should learn that the machine is a tool made by people, not a being with its own intentions. She should learn that the machine learns from data, which means it knows only what it has been shown, and can be wrong about things it has not been shown. And she should learn that the machine sometimes makes mistakes, and that it is her responsibility to notice when it does, rather than to trust its output unconditionally. These three understandings — the machine as tool, the machine as data-shaped, the machine as fallible — are the cognitive foundation on which every subsequent stage of AI literacy will build. Developed at age eight, they become part of the child's mental furniture. Withheld until age fifteen, they are difficult to inculcate against habits of trust and dependence that have been forming for years.

I would add, at Classes 4 and 5, the completion of a first human-in-the-loop project. The project is not technical. It is the child working with the AI tutor on a task — writing a short story, solving a multi-step word problem, researching the habits of a local bird — and developing, through supervised practice, a sense of what the machine contributes to the task and what she contributes. The goal is not a polished output. The goal is a developed intuition for a form of collaboration with machines that will be part of her life forever, and for which the earlier she develops accurate intuitions the better.

The pedagogical footprint of this provision is approximately two hours a week at Class 3, rising to three hours a week at Class 5. The footprint does not appear as a new subject on the timetable;

it is woven into the existing language, mathematics, and environmental-studies periods, in the same way that the chalkboard and the textbook are not subjects but tools through which existing subjects are taught. The teacher's role is supervisory and scaffolding — she is not teaching AI, in any direct sense; she is teaching language and mathematics with AI's assistance, while occasionally stepping back to have brief, deliberate conversations with the children about the nature of the tool they are using. These conversations are the curricular element of AI literacy at this stage. They are brief. They accumulate.

The outcomes I would hope for, by the end of Class 5, are specific. Ninety percent of children should read fluently in their mother tongue and at a functional level in English. Ninety percent should perform the four operations of arithmetic confidently. Every child should have completed at least four human-in-the-loop projects of the kind I described, and should be able to describe, in her own words, what the machine contributed and what she contributed. Every child should be able to explain, in one or two sentences, what data is and why the machine needs it. The outcomes are ambitious compared to current baselines. They are not ambitious compared to what the technology and the pedagogy of this moment make possible, if the teacher capacity is built to match.

A word on what this stage does not contain, because I have been asked the question repeatedly by parents and by institution heads who expected a more ambitious provision. I do not recommend, at Classes 3 through 5, any introduction to coding, whether in block-based or text-based environments. I do not recommend the training of classifiers, the formal treatment of algorithms as a topic, or any introduction to the mathematics that underlies machine learning. The child of nine or ten can be taught these things, and some children will enjoy them, but teaching them at this stage displaces the foundational work that is the stage's genuine priority and produces a surface familiarity with

vocabulary that the child cannot anchor in understanding. A Class 4 child who can recite that a neural network has layers and weights, but who cannot read her Hindi textbook fluently, has been failed by the curriculum that gave her the vocabulary without the foundation. The discipline of restraint, at this stage, is itself a curricular design choice, and it is the one most likely to be eroded by well-meaning institution heads who want their primary school to appear ambitious. The ambition is in the foundation. The rest will come.

I also want to say a specific word about the pedagogical relationship between the AI tutor and the teacher at this stage, because the relationship is delicate and, if it is misunderstood, harmful. The tutor is not a replacement for the teacher's attention. It is an instrument through which the teacher can extend her attention to each child individually in a way that the classroom arithmetic of thirty children and one teacher does not otherwise permit. The teacher continues to teach — at the blackboard, in the conversation with the whole class, in the specific intervention with the child who is struggling. The tutor does the patient practice that the teacher does not have the hours to do at the individual level. The two work together, and the teacher's supervision of the tutor's work — noticing which child is progressing, which child is stuck, which child is engaging superficially — is itself a teaching activity that requires the teacher's full presence rather than her absence. A classroom in which the teacher has stepped away while the tutor works with the children is a classroom that has misunderstood the pedagogy. A classroom in which the teacher is actively circulating, reading each child's progress, and adjusting her own instruction on the basis of what she sees the tutor doing, is the classroom in which the tutor is earning its place.



III. Classes 6 to 8 – the middle stage

At Class 6 the child is, developmentally, at a threshold. She is beginning to think abstractly. She is forming the capacity to handle multi-step logical operations. She is capable, for the first time, of sustained attention to a formal discipline. The middle stage is the stage at which AI enters her curriculum as a subject in its own right – one period a week, forty minutes, across Classes 6, 7, and 8 – organised around three strands that run in parallel across the three years.

The first strand is computational thinking. This strand teaches the child to decompose problems into sub-problems, to identify patterns, to think about sequences and dependencies, to reason about cause and effect, and to design solutions that a machine could in principle execute. The strand does not require that the child write code, though some coding in a block-based environment like Scratch, and then in a simple text-based language like Python, is an effective way to teach the underlying habits. The key is that the child develops the cognitive discipline of algorithmic thinking, which is transferable across every discipline she will encounter, whether or not she ever becomes a programmer.

The second strand is data literacy. This strand teaches the child to read charts accurately, to notice when data has been presented misleadingly, to understand the difference between correlation and causation at a grade-appropriate level, to design simple surveys and interpret their results, and to become generally comfortable with quantitative arguments. Data literacy is arguably the single most important civic skill of the coming decade. It is currently, in most Indian classrooms, not taught at all. The middle stage is the latest point at which it should be introduced, and Class 6 is the right place to begin.

The third strand is AI literacy proper. This strand teaches the child what machine learning is, at a grade-appropriate level, through hands-on activities in which the child trains simple

classifiers — distinguishing between images of cats and dogs, between positive and negative movie reviews, between different kinds of plants. The point is not that the child masters the underlying mathematics, which she cannot at this stage. The point is that she develops, through practice, an accurate intuition for how a machine is taught, what kinds of mistakes it makes, and why those mistakes occur. By the end of Class 8, every child should have trained at least three simple machine-learning models, should have seen them fail, and should be able to explain why they failed.

The three strands run in parallel, not in sequence. A well-designed middle-stage AI curriculum would, in any given forty-minute period, include elements of all three. A period might begin with a computational-thinking exercise about sequences, move to a data-literacy activity involving a real chart from that week's newspaper, and end with a short demonstration of a classifier's behaviour on a new input. The integration of the strands is itself a pedagogical choice, because the three capabilities reinforce each other in ways that separating them would dilute.

Alongside the dedicated period, the middle stage sees the deepening of AI's presence as a tool across other subjects. The child of Class 7 uses the AI tutor to help her draft an essay on Premchand, and develops an awareness, guided by her teacher, of how to prompt the tutor well, how to evaluate its output critically, and when to disregard its suggestions. The child of Class 8 uses the tutor to work through a geometry problem, and learns to use the tutor as a check on her own reasoning rather than as a substitute for it. The teacher's pedagogical skill, at this stage, includes the capacity to model good AI-use, to notice bad AI-use, and to have the conversations with the class that turn habits of use into habits of judgement.

The outcomes I would hope for, by the end of Class 8, are these. Every child should be able to decompose a moderately complex problem into its sub-problems in writing. Every child should be able to read a chart, identify at least one way in which a

given chart is misleading, and produce a more honest version of the same data. Every child should have trained at least three machine-learning models and be able to explain, in her own words, what made each one succeed or fail. Every child should be able to use an AI tutor productively in at least one other subject, and should be able to describe the difference between using it well and using it badly. And every child should have had at least a dozen substantive classroom conversations about the ethical and social questions AI raises — which I will address as a fourth transversal theme in the next subsection.

A note on timetabling, because the question of where the forty minutes a week come from is the question that an institution head will have to answer when she implements this roadmap. I do not propose, at the middle stage, an expansion of the total school week, which is already full. I propose that the forty minutes come from a reallocation of existing periods — in most Indian state curricula, a period of general science, of work experience, or of art and craft that has lost its pedagogical anchor over the years can be repurposed with minimal loss to the displaced subject and substantial gain to the new one. The exact reallocation will vary by state, because state curricula differ, and the negotiation about which period makes way for AI is itself a curricular conversation that SCERTs and schools should have openly. What I would warn against is an addition of the AI period to an already overloaded week without displacement. A week that contains thirty-seven periods cannot absorb a thirty-eighth without something else silently collapsing, and the collapse is usually the period the institution head was least paying attention to — which may well be the one the child most needed.

A second note, on assessment at the middle stage, because what is not assessed is, in the Indian system, frequently not taught. I would argue for a light assessment framework at Classes 6 through 8 — not a board-style examination, which would introduce the rote pressures this stage should specifically avoid,

but a portfolio-based assessment in which each child accumulates, over the year, evidence of her work — the charts she has read and corrected, the classifiers she has trained, the computational-thinking exercises she has completed, the reflections she has written on the social questions she has discussed. The portfolio is reviewed by the teacher at the end of the year and given a simple pass-or-specific-remediation judgement. The judgement is not for ranking. It is for ensuring that the child has, in fact, done the work the stage requires. A school that does not assess the AI curriculum at all will, within three years, have teachers who do not teach it seriously, because in the Indian classroom the unassessed drifts toward the optional. A school that over-assesses it will kill the curiosity the stage is meant to build. The light portfolio-based assessment is the compromise I would recommend.



IV. The fourth theme — AI and society, running through all stages

Alongside the three strands of the middle stage, and running through every stage from Class 3 to Class 12, there is a fourth theme that I want to name separately because it is the theme most in danger of being squeezed out by the other, more technically specifiable content. The fourth theme is the social, ethical, and civic dimension of artificial intelligence.

This theme is not a strand in the same sense as computational thinking or data literacy. It does not require its own period on the timetable. It is a theme that enters the curriculum through every other subject and through deliberate, occasional classroom conversations facilitated by the teacher. It addresses questions that the child will encounter throughout her life, and for which an early grounding makes a substantial difference to the adult she becomes. What does it mean for a machine to be biased? Who decides what a machine is allowed to do? What kinds of decisions should not be

delegated to machines? What happens to workers whose jobs machines take over? What does privacy mean when machines know everything we tell them? What does honesty mean when a machine can produce plausible false statements at scale?

These are not easy questions, and no forty-minute period will resolve them. The pedagogical goal is not resolution but familiarity – the child should recognise these as questions she has the right to think about, has the capacity to form views on, and has the responsibility to revisit as she grows. The teacher's role is to introduce the questions at grade-appropriate levels, to facilitate the conversations that follow, and to resist the temptation to supply the child with the answers. A reform that produces children who can train classifiers but cannot think about whether a given classifier should be deployed is a reform that has produced technicians, not citizens. The fourth theme is the theme that distinguishes the two.

I want to offer, because the question of how to teach this fourth theme without preaching is the question the teacher chapter will not be able to answer in full, one specific pedagogical observation. The fourth theme is taught most effectively not through prescriptive discussion about what the child ought to think, but through the presentation of concrete cases in which reasonable people disagree. A classifier used by a bank to decide which loan applications to process first, trained on historical data in which women received fewer loans than men, now produces recommendations that disadvantage women applicants. Is this classifier biased? Who decides? If the bank's managers override the classifier's recommendations in favour of demographic balance, are they correcting bias or introducing a different kind of it? Children of Class 7 can engage with this question seriously. They will not resolve it. They are not meant to. The engagement itself – the practice of holding a complex question in view without collapsing it into a quick answer – is the disposition the theme is meant to cultivate. A teacher who brings two or three such cases

into the classroom each year, and who facilitates the discussion without steering it toward a pre-determined conclusion, is teaching the theme exactly as it needs to be taught.



V. Classes 9 to 12 – the secondary stage

At Class 9 the child, now a young person of fifteen, is capable of real specialisation. The secondary stage, in my recommendation, offers three optional AI-related electives, each of which is a substantial two-year programme culminating in a final examination and demonstrable projects. A student may take one of the three. A student who intends to pursue a scientific or technical university education may take two. No student is required to take any – the foundation laid in Classes 3 through 8 is, by design, sufficient for an intelligent adult life even without further specialisation. The electives are for the student who wants to go deeper, and they recognise that different students will go deeper in different directions.

The first elective is AI and Society. This elective draws on history, economics, political science, philosophy, and literature, and addresses the social and ethical questions the fourth theme has been raising since Class 3, now at the depth and seriousness an adult citizen requires. A student completing this elective should be able to write a competent essay on the political economy of AI regulation in India, to argue a defensible position on the displacement of labour by automation, and to engage critically with public discourse about AI in a way that is neither credulous nor reflexively alarmed. The elective is especially suited to students who will go on to study humanities or social sciences at university, and to students who will become the journalists, civil servants, and public intellectuals who shape the national conversation in the decades to come.

The second elective is Applied AI. This elective is the technical track. Students learn programming at serious depth, working in Python and perhaps one other language. They learn the mathematics of machine learning at introductory level, building on the statistics introduced in the middle stage. They train more sophisticated models. They deploy at least two working AI applications — which could be anything from a chatbot for a specific domain to a classifier that solves a specific problem — and submit them as their final projects. A student completing this elective is prepared to enter a computer-science or data-science undergraduate programme with a genuine head start, and is prepared, if she chooses, to enter the workforce directly as an AI-capable technical professional.

The third elective is Data and Design. This elective combines statistics, data visualisation, and design thinking. It teaches the student to ask good questions of data, to present findings clearly, and to design interactions between humans and AI systems that are usable, trustworthy, and accessible to users of varying backgrounds. This elective is especially suited to students who will go on to study business, design, public policy, or any field in which the capacity to work intelligently with data is an advantage without the technical specialisation of Applied AI being required.

Alongside the electives, AI at the secondary stage becomes genuinely a research partner for the student's own work. The student of Class 11 writing a history project on the First War of Independence uses AI to source contemporary accounts, to translate Urdu sources her teacher cannot read, to check her own writing for errors of fact and argument. The student of Class 12 preparing for her board examinations uses AI to identify her weakest topics and to practise them at the depth she needs. The sophistication of her use grows as she grows. By the time she leaves school, she should be capable of using AI tools as a professional adult would use them — with competence, with judgement, and with an awareness of the tools' limits.

I should say, because the question will be asked by every board-examination parent who reads this, how the board examinations themselves should treat AI literacy. My recommendation is that the three electives be examined at board level on the same basis as other electives — with a written examination for the AI and Society and Data and Design tracks, with a practical-project component for Applied AI, and with weightage comparable to other Class 12 subjects. The board-examination infrastructure, unloved as it is in much contemporary Indian educational writing, remains the single most consequential downstream signal that the school system responds to. Subjects that the board examines are subjects that schools resource and teachers teach seriously. An AI curriculum that the boards do not examine will, within five years, be taught with the perfunctory attention that unexamined subjects receive in most Indian secondary schools. The examination is therefore not a pedagogical afterthought but a structural decision that determines whether the curriculum is real. The CBSE, the state boards, and the council of boards must be brought into the mission's planning from the first year, not in year three, and the question of how the three electives are examined must be resolved before the first cohort of Class 11 students enters the stream.

A separate question, on which I hold a view but do not claim certainty, is whether the three electives should be examinable independently of the rest of the student's Class 11 and 12 stream. My view is that they should be. A student who is otherwise in a humanities stream, who takes the Applied AI elective out of intellectual interest, should be able to receive the elective's credential without its being tethered to a science-stream requirement she does not need. The elective-as-independent-qualification model is, in some ways, a departure from the Indian board tradition of tightly coupled subject combinations, but it is the model that most closely serves the educational purpose of the electives themselves, which is to give students the specific cognitive preparation that their own interests and trajectories call

for rather than to filter them into tracks the system has already decided they belong to.

The outcomes I would hope for, by the end of Class 12, are these. Every student should be capable of competent AI use as a working professional or university student would require. Every student should be able to argue a defensible position on at least one substantive question about AI in society, at a depth appropriate to a seventeen-year-old. Students who take the Applied AI elective should have deployed at least two working AI applications. Students who take the Data and Design elective should have completed at least two substantive data projects. Students who take the AI and Society elective should have produced written work at a depth that would be respected in a university admissions context.



VI. A note on pedagogy

The roadmap above is a specification of outcomes, not of classroom methods. The classroom methods that produce these outcomes are the subject of the teacher chapter that sits at the book's centre. I will not pre-empt that chapter here. I will observe only, in closing, that the roadmap rests on pedagogical assumptions that cannot be taken for granted in the current Indian classroom, and that without the investment in teacher capacity that Chapter Six specifies, the roadmap is a set of aspirations rather than a set of achievements. A curriculum without the teachers to deliver it is one of the most familiar failure modes of Indian education reform. I name this failure mode now, so that the reader carries it with her as she reads the chapters that follow, and so that she understands why the teacher chapter is where the book's weight eventually falls.

I want to add, at the risk of labouring the point, one further observation about pedagogy. The roadmap as specified is, in a sense, content-neutral with respect to the underlying pedagogical

philosophy that animates it. A constructivist teacher, a behaviourist teacher, a teacher trained in the activity-based learning methodologies that have spread through Tamil Nadu and Kerala over the past two decades, a teacher whose default is traditional chalk-and-talk — each of them can, in principle, deliver the outcomes the roadmap specifies, through her own pedagogical approach. I note this because the reform will be tempted, in its zeal, to mandate a specific pedagogy alongside the specific content. I would advise against this. A reform that tries to change both what teachers teach and how they teach, simultaneously, at the scale of nine million teachers, is a reform that will fail at both. The roadmap specifies outcomes. The pedagogical methods by which those outcomes are reached should, within broad professional standards, be left to the teacher's own judgement, informed by her training and adapted to her classroom. The discipline of not over-specifying the pedagogy is itself a form of professional respect, and it is the form without which teachers will experience the reform as an imposition rather than as a partnership.



Three decisions for the reader

1. If you are an institution head, commission your own institution's AI-education roadmap within six months.

Do not wait for the Centre. Using the framework above as a starting point, or using another framework you prefer, develop a specific plan for what AI education looks like at each grade in your institution. Specify the outcomes you will pursue at each grade. Specify the pedagogical methods you will use. Specify the teacher capacity required. Specify the timeline over which you will build to the final state. Share the plan with your teachers, your parents, and your board. Refine it on the basis of their input. Publish the final version. The plan is, in itself, a contribution to the national conversation, and it positions your institution as a leader rather than a follower in the reform.

2. If you are a parent, ask your child's school for its roadmap.

Most schools in India today do not have one. When you ask for it, and the school does not have one, you have done the school a service — you have put on the school's agenda a question it can no longer avoid. If the school does have one, read it carefully. Ask the school to explain how it developed the roadmap, who was consulted, how teacher capacity is being built, and what the school will do if a child falls behind the intended outcomes. The school's answers, or its lack of them, will tell you more about the school's seriousness on this question than any brochure ever will.

3. If you are an officer of the Centre or the state, consider whether a national roadmap should be issued — and if so, how.

The roadmap above is one practitioner's view. A better roadmap, produced by the National Curriculum Framework's custodians in consultation with educators across the country, would be of substantial value to the reform. Its value would lie not in producing a single standard every school is forced to implement, but in providing a credible reference point against which every school's own roadmap can be evaluated. I would urge the Centre to commission such a document within the next twelve months, to invite public consultation on its draft, and to publish it as non-binding guidance rather than as a regulatory requirement. Guidance that is good enough to be adopted widely by choice is more powerful than a requirement that is complied with resentfully.



END OF CHAPTER THREE

CHAPTER FOUR

The Content Question — Who Writes What Our Children Will Read



On the under-discussed problem that will determine whether the reform reaches the child.



I want to open this chapter with a small admission. When I first began thinking seriously about AI education in India, I, like most people who begin thinking about this, spent the first several months thinking about technology. What AI tutors can do. What they cannot yet do. Which vendors are credible. What the benchmarks say. What the research on learning outcomes shows. These are interesting questions, and I do not want to dismiss the time I spent on them. But I came, in the end, to an uncomfortable conclusion, which is that the technology question is not the binding constraint on whether this reform succeeds. The binding constraints are two, and they are the teacher and the content. This chapter is about the second.

By content I mean everything the child will read, hear, watch, and interact with in the course of her AI-education experience. The textbook she opens at the start of the year. The worksheet her teacher photocopies. The lesson plan on her teacher's phone. The short video that explains what a neural network is. The conversational prompt her AI tutor uses when it greets her at the start of the period. The classifier she trains in Class 7, and the images it is trained on. The essay her Class 11 sister reads about the

political economy of automation. All of this is content. All of it needs to be written, designed, produced, translated, reviewed, updated, and made freely available to every Indian school in every Indian language the children of the country actually speak. None of it exists today at the scale and quality the reform requires.

I include, in this definition of content, a category that is often overlooked because it is not obviously content in the textbook sense. The category is assessment content — the specific questions a child encounters on her classroom quizzes, her unit tests, her board examinations. A reform that produces excellent curricular content and excellent AI tutors but that leaves the assessment questions unchanged, written by the same question-bank contractors who have written them for the past two decades, is a reform whose downstream signal to teachers and students is that the assessment framework is what school is really about, and the new content is a parallel activity that does not connect to the real measure of performance. The assessment content must be rewritten alongside the curricular content, at comparable cost and with comparable quality standards, and the rewriting must be visibly tied to the specific outcomes the curriculum now prescribes. This is substantially more work than it may first appear, because the question-bank infrastructure in Indian education is large, old, and resistant to reform. Including assessment in the content scope is one of the design choices that makes the reform more expensive and more difficult than a narrower definition would suggest. It is also one of the design choices without which the reform will not, in practice, reshape classroom behaviour to the extent it aspires to.

The absence is not for lack of effort. NCERT has produced materials. CBSE has produced materials. The state boards have produced materials. Private publishers have produced materials. There is, in the aggregate, a great deal of material in circulation. What there is not, in the aggregate, is material that meets the three tests a serious reform must hold its content to: that it is pedagogically sound, that it is linguistically and culturally

grounded in the child's own context, and that it is freely available to every school without dependence on a vendor's commercial decisions. The argument of this chapter is that producing such content, at the scale of two hundred and fifty million children, in at least twenty-two Indian languages, across ten grade levels, across curricular and cross-curricular uses, is one of the three largest undertakings the reform requires – comparable in ambition to building the teacher capacity and comparable in fiscal demand to deploying the devices. It is also the undertaking that has received, in the current national conversation, the least serious attention.

I want to name the specific reason the content undertaking has received the least attention, because the reason is instructive and unflattering. Content is, in comparison to devices and teacher training, the least politically legible component of the reform. A device deployment produces visible objects that can be shown to visiting dignitaries. A teacher-training programme produces ceremonies, certificates, and photographs of workshops in progress. Content, by contrast, is largely invisible. A successful textbook is one the child reads without noticing it; a successful AI-tutor conversational corpus is one whose specific turns the child experiences as natural rather than as designed. The invisibility of content's success, combined with the invisibility of content's production – which happens in editorial rooms rather than in ceremonies – makes content a less attractive component of the reform for the political sponsors who need visible outputs on a political timeline. The reform's political leadership must therefore make a specific and sustained effort to champion the content work, against the structural invisibility that will otherwise consign it to the status of an administrative afterthought. A minister who insists, in her first year, on a formal update on the content programme's specific progress – commissions completed, languages served, artefacts produced – is a minister who has begun to counteract the structural disadvantage. A minister who does not is a minister whose reform's content component will,

predictably, fall behind the other components by the end of her first year.



I. What content means at each stage

Let me walk through, concretely, what content is needed at each stage of the grade-by-grade roadmap the previous chapter set out, because the abstraction of content is one of the reasons the problem is under-recognised. When we make the content specific, its weight becomes visible.

For Classes 3 through 5, the primary content requirement is the conversational corpus that the AI tutor uses to interact with the child. This is not a textbook in the conventional sense. It is the library of prompts, stories, questions, explanations, and corrections the tutor draws on as it works with the child in her mother tongue, on the foundational literacy and numeracy tasks the stage is built around. The corpus must be age-appropriate, culturally grounded, pedagogically scaffolded, and available in each of the major Indian languages the children actually speak at home. It must include children's stories the child will recognise — not translated stories from elsewhere, but stories from her own linguistic tradition, carefully selected and, where necessary, freshly written. It must include arithmetic word problems set in contexts the child knows — her own village, her own market, her own festival calendar — rather than abstracted or imported contexts that distance her from the problem. It must include explanations that respect the child's intelligence without presuming knowledge she does not yet have. The corpus does not exist, at the scale and quality required, in any Indian language. It must be built, from scratch, by teams that include writers for children, pedagogical experts, linguistic specialists, and classroom practitioners, working alongside the technical teams that will deploy the corpus into the AI tutors.

For Classes 6 through 8, the content requirement broadens. There is a textbook for the dedicated AI-literacy period, which must cover computational thinking, data literacy, and AI literacy proper, across three years, in twenty-two languages, with worked examples and exercises that a teacher whose own background in these subjects is limited can use with confidence. There are lesson plans that help the teacher structure each forty-minute period around the three strands. There are teacher guides that explain not only what to teach but the common misconceptions children have and how to address them. There are short videos, between two and five minutes each, that can be used in class to illustrate concepts that a written explanation handles less well. There are simple software environments — in-browser, ideally, so that no installation is required — where children can train their first machine-learning models under the teacher's supervision. There are reading lists for children who want to go further on their own. None of this, in the current landscape, exists in even the major Indian languages, let alone in the less-served ones.

For Classes 9 through 12, the content requirement becomes more specialised and, in some ways, more manageable. The three electives I proposed in Chapter Three each require their own curriculum, textbooks, and assessment materials. The Applied AI elective requires coding environments, problem sets, and project templates. The AI and Society elective requires a carefully curated set of readings — some translated from English into Indian languages, some originally written in Indian languages — that offer the student multiple perspectives on the questions she is being asked to think about. The Data and Design elective requires real datasets drawn from Indian contexts, covering subjects from agriculture to public health to election results, on which students can practise the analytical work the elective trains them for. At this stage the child's English competence is typically stronger, which reduces though does not eliminate the translation burden. The bulk of Class 11 and Class 12 content in the electives can be in English with translations into the major regional languages, and this is

probably the right trade-off given the scale of the overall content challenge.

Alongside the curricular content, there is a much larger cross-curricular content requirement. The mathematics textbook of Class 8, if AI is genuinely to reshape how mathematics is taught, needs to be rewritten with AI integration in mind. The history textbook of Class 10, if the student is to use AI as a research partner in a disciplined way, needs to include worked examples of how to use AI to source primary material and how to evaluate what the AI produces. The Hindi textbook of Class 7, if AI is to improve the student's writing rather than substitute for it, needs to include guidance on when to use the tutor and when not to. Each of these cross-curricular reshapings is, in itself, a substantial content project. The scale of the overall content requirement is genuinely very large, and there is no honest way to make it sound small.

Let me try to put an arithmetic on it, because the scale is hard to hold in mind without specific numbers. The core AI-literacy textbook for the middle stage, at quality, runs to approximately two hundred and fifty pages per grade, across three grades, in each of twenty-two scheduled languages — sixty-six volumes of approximately two hundred and fifty pages, each requiring commissioning, writing, review, illustration, typesetting, and ongoing revision. The teacher guides that accompany each volume add a further three hundred pages per grade per language — another sixty-six volumes. The lesson-plan bank for the middle stage, if it covers all one hundred and twenty periods across the three grades with two or three alternative plans per period, runs to perhaps four hundred plans per language, or nine thousand plans across the scheduled languages. The video library, if it covers approximately sixty concepts at the middle stage with between three and five minutes of carefully produced content per concept, runs to sixty videos per language — thirteen hundred videos in scheduled languages, plus the originating English and Hindi versions. The conversational corpora for the AI tutors at the

foundational stage, if each corpus is to support genuinely adaptive individualised practice across the literacy and numeracy curricula, run to several hundred thousand carefully authored and reviewed prompts per language — a volume that has no real precedent in any Indian content-production effort to date. At the secondary stage, each elective requires its own full textbook sequence, problem bank, project specification library, assessment materials, and teacher preparation package, across the two-year duration. The aggregate, summed across stages and languages, is of the order of a hundred and fifty thousand discrete substantial content artefacts, each requiring its own authorial judgement, editorial review, and maintenance commitment. This is the content footprint of the reform. A country that understands the scale of what it is committing to is a country that has a chance of resourcing it. A country that does not is a country that will, four years into the reform, discover that it has deployed devices into classrooms for which the content to populate them has not been written, and will scramble to fill the gap with whichever vendor's catalogue is closest to hand.



II. The languages we must serve

I want to say a word about language, because the Indian content challenge is, in a way that the Chinese or the Singaporean or the American content challenge is not, fundamentally a multilingual challenge. India has twenty-two scheduled languages under the Eighth Schedule of the Constitution. Its 2011 census recorded one hundred and twenty-one languages spoken by ten thousand or more people each. The Indian schoolchild, depending on where she is born and what her family speaks at home, may receive her foundational education in any of perhaps thirty different languages of instruction. A reform that produces excellent content in Hindi and English, and nothing else, is a reform that has

served approximately half the country and abandoned the other half.

The linguistic scope of the reform must, at minimum, cover the twenty-two scheduled languages. That is the floor. The floor is itself ambitious, but it is non-negotiable. Serving fewer than twenty-two languages means, in practice, serving fewer than the number of states whose government schools the reform is claimed to address. A state whose children are taught in Odia or Assamese or Bodo, and who receive AI-education content only in Hindi or English, is a state whose children have been given the name of the reform and denied its substance. The Centre cannot pretend otherwise.

Within each of the scheduled languages, there is further variation the content must accommodate. The Hindi of Lucknow is not the Hindi of Bhopal is not the Hindi of Patna. The Tamil of Chennai is different in register from the Tamil of Madurai. A content programme that insists on a single standardised variant within each language will produce material that reads as foreign to some of the children who encounter it, even in their own official language. The best Indian children's literature has, over decades, navigated this problem through careful editorial practice and regional adaptation. The content for AI education must learn from that tradition.

Bhashini — the India Language Technology Mission under MeitY — has, as I noted in Chapter Two, produced infrastructure for Indian-language AI that is, in global perspective, genuinely impressive. Its translation models are among the best available for most major Indian languages. Its speech recognition and speech synthesis systems are workable for several languages and improving rapidly for others. Bhashini's existence is one of the most important preconditions for the reform, and it is one of the specific assets that makes a national AI-education rollout more feasible in India now than it would have been five years ago or than it is in most comparably complex countries today.

But Bhashini is infrastructure. It is not content. The distinction matters. Bhashini can translate a paragraph from Hindi to Marathi with reasonable quality. It cannot, and was not designed to, originate a paragraph of pedagogically sound, age-appropriate, culturally grounded educational material in either Hindi or Marathi. The origination is human work, done by human writers and educators, informed by Bhashini's tools but not substitutable by them. A content programme that assumes Bhashini will solve the language problem on its own is a programme that has misunderstood what Bhashini is. It is a translation tool, not a content-origination tool, and the distinction between those two is the distinction between a useful infrastructural asset and the magic wand that is sometimes implied in the national conversation.

There is, moreover, a more specific problem that translation tools do not solve. A great deal of the content the reform needs is content that simply does not exist in the major Indian languages at sufficient quality or quantity to translate from. The pedagogically sound AI-literacy textbook in English is itself not yet written at the quality required. The conversational corpus for an AI tutor that is pedagogically excellent does not exist in English in the form that translation could then adapt to Indian languages. The content origination work must happen at the source, in parallel in multiple languages, with translation serving as a supplement rather than as the primary engine. I emphasise this because a cheaper, faster model is available — write everything in English, translate it into twenty-two languages via Bhashini, deploy — and this model is, I believe, the path to a failed reform. The children who encounter the translated material will know, within weeks, that they are receiving second-class content written for somebody else.

I want to say a specific word about a category of languages that the scheduled-languages framing hides, which is the tribal and regional languages that are not on the Eighth Schedule but that are the home languages of several tens of millions of Indian children. Santali is on the Schedule; Mundari, spoken in the same districts

by neighbouring communities, is not. Bhili is spoken by fifteen million people across four states and is not on the Schedule. Gondi, Kurukh, Ho, Khasi's many dialects, Garo, the Konyak family of Nagaland's languages, the Bodo-related languages of the northeast — each of these is the mother tongue of a population of children the reform is claimed to serve, and each presents a content challenge that the scheduled-languages framing does not address. I do not propose that the Repository, in its first decade, produce full-stage content in every one of these languages — that is beyond what can reasonably be delivered. I propose, instead, that the Repository identify a specific tier of priority languages below the Scheduled twenty-two — perhaps thirty additional languages whose speaker populations exceed one million children of school age — and commit to a reduced but genuine content provision in each of them, at least at the foundational stage. The provision will be slower, the investment per child higher, the political visibility lower. But a reform that has nothing to offer the child whose home language is Bhili or Gondi is a reform that has defined her, from the beginning, as secondary to its purpose, and the definition will be read accordingly by her community.

I should also acknowledge, in the interest of honesty, that the ideal of parallel origination in multiple languages will in practice be compromised by realities of talent, institutional capacity, and timelines. There are not, today, enough children's-literature-trained writers and pedagogically-literate editors working in Odia or Assamese or Bodo to produce, from scratch, the full content volume I have described on the timeline the reform requires. Kerala has the depth for Malayalam; Tamil Nadu has it for Tamil; West Bengal has it for Bangla. Several other states have institutional strength in their own languages. But the weaker state ecosystems will need to import talent, train new talent, and rely more heavily on adapted translation in the early years of the reform than the principle would ideally demand. The honest planning acknowledges this. It specifies the target state — parallel origination at quality, in every scheduled language — and accepts

the transition state, during which some languages are served primarily through careful adaptation of material that originated elsewhere. The transition state should be time-bounded, publicly reported, and managed with the explicit intention of reducing dependence on adaptation over the first decade. A transition that becomes permanent is a transition that has replicated the second-class content problem it claimed to be solving.



III. Who should write this content, and how

The question of who writes the content is a question about institutional design, not a question about individual authors. I want to sketch, in this section, the institutional arrangement I believe the content programme requires.

At the centre, there should be a National AI Education Content Repository, operated by NCERT with the support of the National AI Education Mission I propose in Chapter Ten. The Repository is the authoritative home for curricular materials, teacher guides, videos, software environments, and the conversational corpora used in AI tutors deployed in government schools. The Repository's materials are produced under the Creative Commons Attribution licence, which means they are free for any school, public or private, to use, modify, and redistribute, subject only to attribution. The licence choice matters. A Repository that charges for its materials, or restricts their reuse, is a Repository whose contents do not circulate widely enough to define the standard of the reform. The Creative Commons model is the model that has worked, in other countries and other fields, for content that is intended to be genuinely public infrastructure.




Figure 4.1 | The National AI Education Content Repository – commissioning, origination, adaptation, delivery.

The Repository's materials are developed through a distributed process. NCERT does not produce most of them directly. NCERT, through the Repository, commissions them – from state SCERTs, from universities, from civil society organisations, from individual authors of established reputation, from coalitions of the above. Each commission is for a specific deliverable in a specific language at a specific grade level, with a specific pedagogical specification and a specific delivery deadline. The commissions are competitive – NCERT publishes the specifications, receives proposals, selects the commissioners on the basis of their track record, their proposed approach, and their capacity to deliver. The cost of each commission is transparently declared. The timeline is public. The draft is published for public consultation before it is finalised, and the finalised version is published alongside a record of what was changed in response to consultation and what was not.

A specific note on the commissioning-process design, because the quality of the commissioning is the single factor that most

determines the quality of what the Repository ultimately produces. Indian procurement culture, in educational content as in other fields, has tended to produce a specific pathology — the commissioning criteria are written to be defensible on audit, the selection criteria weight lowest-price bids disproportionately, the selected commissioners are often the organisations that have learned to navigate the process rather than the organisations best equipped to produce excellent content. The Repository must escape this pathology. Its commissioning process should privilege demonstrated track record over administrative compliance, should weight quality at least as heavily as price, and should build in a specific discipline of commissioning across a diverse range of authors and organisations rather than allowing a small number of well-networked providers to capture the bulk of the work. The process should also allow for small, exploratory commissions — at the scale of a few lakh rupees for a single pilot artefact, with the expectation that most will not proceed to larger commissions — so that the Repository can discover new voices and new approaches alongside the established ones.

The quality review process inside the Repository deserves specific attention. Every piece of content the Repository publishes should pass through at least three layers of review — a subject-matter review by a specialist in the domain, a pedagogical review by an educator with classroom experience in the relevant grade, and a cultural-fit review by a reader grounded in the linguistic and cultural context of the intended audience. The three layers are not redundancies. Each catches a specific class of problem that the others will miss. A textbook chapter may be factually accurate and pedagogically sound and yet culturally discordant in ways that a subject-matter expert and an educator, both reading in their second language, do not notice. Only the reader grounded in the target community will catch the specific unfamiliarity, the specific absence, the specific image that does not land in the way the writer intended. Building a cadre of cultural-fit reviewers across the twenty-two languages, at the scale the Repository requires, is itself

a substantial undertaking, and it is an undertaking that conventional content-production pipelines in India have systematically underinvested in.

This model has important virtues over the alternatives. It draws on the substantial content-development capacity that exists across India's universities, civil society organisations, and professional writers, rather than trying to concentrate all production inside NCERT. It is transparent to parents, teachers, and the public about what is being produced, by whom, at what cost, and to what quality. It provides a mechanism for continuous improvement, because the Repository is not a one-time production but an evolving library that improves its own materials in response to classroom feedback. And it creates a visible, auditable trail that protects the reform against the capture by any single vendor or institutional interest that has killed earlier content programmes in other countries.

The Repository needs a dedicated fiscal envelope. My estimate, which is necessarily rough, is that producing and maintaining a Repository of the scale the reform requires – covering all three stages, at least twenty-two languages, with continuous updating over the first ten years – will cost approximately one thousand crore rupees a year, or ten thousand crore over the decade. This estimate is built up from assumptions about the cost of individual commissions, the number of commissions required, the cost of the ongoing editorial and technical infrastructure to maintain the Repository, and the cost of the translation, review, and updating processes that keep the materials current. The estimate is a starting point for a serious budgeting exercise, not a final figure. It is, in any case, less than one percent of the annual central education budget, which is a defensible allocation for an undertaking of this consequence.

Alongside the Repository's curricular materials, there is a specific category of content that requires separate institutional handling. The conversational corpora used in AI tutors are, in

effect, the things that actually speak to the child when she sits with the tutor on a Tuesday morning in her classroom. The pedagogical quality of these corpora is the single most consequential determinant of whether the tutor works or fails. I believe the Repository should include, as a core component, reference conversational corpora for each stage and each language — corpora that vendor-operated AI tutors must either adopt directly or demonstrate pedagogical parity with, as a condition of being qualified to sell into government schools. Without this reference, the market will fragment across vendor-specific corpora of highly variable quality, each optimised for its vendor's commercial imperatives rather than for the child's learning outcomes. With the reference, there is a floor below which no vendor may go, and a body of shared pedagogical infrastructure that the market can then compete to exceed.



IV. The translation and adaptation layer

I have argued that content origination must happen in parallel in multiple languages, rather than in English first and then translated. This is true as a principle. In practice, the reform cannot wait for perfect parallel origination in twenty-two languages from day one. The Repository's initial content will necessarily be richer in some languages than in others, and the gaps must be filled through a combination of fresh origination and adaptation of existing material.

The adaptation process is itself a serious undertaking. Adaptation is not translation. A translated paragraph conveys the same information in a different language. An adapted paragraph replaces examples, references, and cultural assumptions with those appropriate to a new linguistic context. The arithmetic word problem set in a Gujarati child's market is different from the same problem set in a Manipuri child's market, even if the mathematical

content is identical. The adaptation is the difference between content that respects the child and content that pastes her language onto somebody else's world. The Repository's commissioning process should specify adaptation, not translation, as the standard for material that is developed in one language and then made available in others.

There is also a category of content that should be developed by each state independently, within a common framework provided by the Repository. The local examples — the local festivals, the local geography, the local heroes — that make content feel like it is from the child's own place rather than from some generic India are best developed by educators from that place, not by the Centre. The Repository's role, for this category of content, is to provide the framework and the quality standards, and to accredit the state-produced materials when they meet the standards. The state SCERTs are well placed to run these state-specific content programmes, with their DIETs contributing the district-specific localisation where that is appropriate.

The translation and adaptation layer is also the layer where Bhashini's infrastructure makes its most useful contribution. Bhashini's translation tools can produce a first draft of a translation from one Indian language to another, which a human editor then revises for accuracy, register, and pedagogical appropriateness. The human editor's time is substantially reduced, compared to translating from scratch, and the quality of the final output is comparable. The Repository should build its translation workflow explicitly around this human-in-the-loop model, combining Bhashini's efficiency with the editorial judgement that the reform's quality demands.



V. The maintenance problem

Educational content, once produced, is not finished. It must be maintained — updated for factual changes, refined on the basis of classroom evidence of what works and what does not, revised in response to curricular updates, and occasionally rewritten when the technology or the pedagogy underlying it advances. A reform that treats content production as a one-time project is a reform that will find, three years after its launch, that its content is stale, inaccurate in places, and unloved by the teachers who are meant to use it.

The maintenance problem is specific to content for a rapidly evolving field. The AI of Class 6 in 2028 is, in the current trajectory of the technology, going to look meaningfully different from the AI of Class 6 in 2025. A textbook written in 2025 and not touched again will, by 2030, describe systems that no longer exist, ignore systems that have become central, and teach pedagogical approaches that the field has moved past. The Repository must therefore be a living institution, with an editorial team whose permanent job is to maintain currency, rather than a one-time production effort that disbands after delivery.

The editorial team's size, I estimate, should be of the order of two hundred people, spread across the languages and subjects the Repository covers. This is a substantial but not unreasonable permanent establishment. Its cost is part of the one thousand crore annual estimate I offered earlier. Its composition should include senior writers and editors, subject-matter specialists, pedagogical experts, classroom practitioners on secondment from schools, and technical staff for the software and infrastructure. It should be physically distributed across multiple Indian cities, not concentrated in Delhi — Bengaluru, Chennai, Kolkata, Hyderabad, and Guwahati are among the cities where its editorial nodes might sensibly be placed. A distributed editorial organisation is more closely connected to the regional linguistic contexts it serves than a Delhi-centralised one is, and the distribution itself is a form of

institutional discipline against the Delhi-centric habits that have weakened earlier content efforts.



VI. Two risks to name honestly

Before I close the chapter, I want to name two risks that the content programme must guard against, because if they are not named now they will be discovered late.

The first risk is capture by commercial interests. There is substantial money at stake in the content market, and substantial lobbying pressure will be applied, by both Indian and foreign vendors, to have their own content adopted as the Repository's standard, or to have the Repository's standards written in ways that advantage their existing products. Some of this pressure is legitimate — vendors who have invested seriously in content have expertise worth drawing on, and there is no reason to exclude their contributions if they meet the Repository's quality standards under its licence terms. But the line between a vendor contributing within the Repository's standards and a vendor shaping the standards to fit its product is a line that must be held explicitly. The Repository's commissioning process, its quality review, and its editorial decisions must be insulated from commercial influence by procedural and institutional design. The mechanisms that have protected similar public content infrastructures in other countries — transparent processes, independent review committees, public consultation, conflict-of-interest disclosures, whistleblower protections — are the mechanisms that must be built in from the beginning, not added after a scandal.

I want to be slightly more specific about what capture by commercial interests can look like in practice, because the general statement of the risk does not give the reader enough to recognise the pattern when it occurs. Consider three forms the capture can take. The first is the specification-shaping capture, in which a

vendor works with the Repository's standards-setting process to ensure that the technical specifications for, say, the AI tutor interaction format correspond closely to what their existing product already produces — and the subsequent commissioning calls are, accordingly, written in ways that make their product maximally competitive. The second is the editorial-board capture, in which a vendor's founders, board members, or advisors come to occupy a disproportionate share of the Repository's review and advisory committees, and the editorial judgements those committees make tilt, over time, in directions that serve the vendor's interests. The third is the data-standard capture, in which the Repository's interoperability standards are set in ways that conform to one vendor's existing data architecture, with the consequence that other vendors face substantial migration costs while the favoured vendor incurs none. Each of these forms is subtle enough that it can be defended, by the participants, as normal engagement. Each of them, accumulated, produces an institution whose public character has been compromised without any single decision being visibly improper. The specific procedural safeguards — rotating committee memberships, public disclosure of participating affiliations, published rationales for specification choices, periodic external audit of the standard-setting process — are the mechanisms through which the incremental capture is resisted. No single safeguard is decisive. The aggregate discipline of applying them all, consistently, is the only protection that works.

The second risk is capture by ideological interests. Content for children is content about which people disagree, and some of the disagreements are, in the Indian context, politically charged. The history of the introduction of computers into Indian schools over the past three decades includes episodes in which content was shaped, in ways that were sometimes obvious and sometimes subtle, by the ideological preferences of the government of the day. The AI-education content programme is particularly vulnerable to this pattern because so much of its material — the examples, the case studies, the readings in the AI and Society elective, the corpora

the tutors draw on — is narratively and culturally loaded. A programme that produces content which children across the political spectrum can recognise as fair is a programme that will outlast changes of government. A programme that produces content which reads as partisan will be revised with each change of government, and the children will be the casualties of the revisions.

The safeguard, again, is institutional. The Repository's governance should include representatives from across the political spectrum, and its editorial processes should include review mechanisms that catch ideological imbalance before publication. No safeguard is perfect, but the attempt to build one is itself a form of discipline, and the attempt is preferable to its absence.

I want to close the chapter with an observation about the cultural politics of content production that the preceding paragraphs have not fully addressed. Indian educational content, over the past several decades, has often been produced under the shadow of the question of what Indian education should be about — a question on which thoughtful Indians disagree strongly, and which the content they produce inevitably reflects in some measure. The AI-education content programme will not be exempt from this. The examples chosen to illustrate a machine-learning classifier, the historical case studies used in the AI-and-Society elective, the reading list compiled for the Class 11 essay on automation and labour, the representative images in the foundational-stage reader — every one of these choices is also, in some degree, a choice about which India the content is for. A content programme that pretends these choices are neutral is a programme that will be making them unconsciously. A programme that names the choices openly, discusses them in its commissioning process, and produces content that is recognisably inclusive of the country's genuine diversity — rural and urban, north and south, dominant-caste and marginalised-caste, Hindu

and Muslim and Sikh and Christian and tribal, boy and girl, abled and disabled — is a programme that has at least tried to serve the country it is ostensibly serving. The attempt will be imperfect. Any attempt at this scale will produce specific missteps that individual readers will, rightly, object to. The discipline is not the achievement of perfect representation but the sustained commitment to the question, and the willingness to revise content when specific objections are found to be well-grounded. A Repository that revises on the basis of good-faith feedback is a Repository that will, over time, become better than any single commissioning cycle could make it.



Three decisions for the reader

1. If you are an institution head, audit your current AI-education content within three months.

Make a list of every piece of content — textbook, worksheet, video, software — that your institution is currently using to deliver any element of AI education. Review each item against the three tests I named at the start of this chapter: is it pedagogically sound, is it linguistically and culturally grounded for your children, is it freely available without vendor lock-in. Note the gaps. Where material is missing entirely, note that. Where material is present but fails one of the tests, note that. The resulting audit is, in itself, a substantial contribution to your institution's AI-education planning, and is a document your senior leadership team should work with over the following twelve months.

2. If you are a content producer, consider commissioning your work into the Repository.

If the National AI Education Content Repository exists by the time you read this book, study its commissioning process and

consider whether the content your organisation produces, or could produce, fits within its specifications. If the Repository does not yet exist, produce your content anyway, under a Creative Commons licence, in the format the Repository will eventually adopt, and publish it in a form that other educators can find and use. Early producers of quality open-licensed content set the de facto standard that the eventual Repository will inherit, and your work today will, if it is good, become part of the national infrastructure tomorrow.

3. If you are a policymaker, commission the Repository within the next twelve months.

The Repository is the single most consequential institutional artefact the reform requires, alongside the teacher-capacity infrastructure I address in Chapter Six. If it does not exist, the reform will produce devices and policies that have nothing consistent to carry. The commissioning decision should be taken at the level of the Ministry of Education, with NCERT as the lead implementing body, with a dedicated fiscal line in the next Union Budget, and with the governance safeguards I have sketched built in from the start. The precise design is a matter for the Cabinet note that follows such a decision. The decision itself can, and should, be taken in the current financial year.



END OF CHAPTER FOUR

CHAPTER FIVE

The Technology Question — Devices, Platforms, and Not Being Captured



*What schools actually need, what they do not need, and how to procure
without losing control.*



I have said, more than once in this book already, that the technology question is not the binding constraint on whether the reform succeeds. I stand by that. But the technology question is not trivial either, and an institution head or a policymaker who underestimates the complexity of getting the technology right will find herself, three years into her deployment, trapped in arrangements she did not intend and cannot easily escape. This chapter is about what to think about when thinking about the technology, and particularly about how to avoid the specific patterns of vendor capture that have characterised educational technology procurement in India and elsewhere.

I will observe, before the specific arguments begin, that the technology question tends to attract disproportionate attention from reformers precisely because it is the most tractable of the reform's components. The technology has specifications that can be written, vendors that can be evaluated, deployments that can be measured. By contrast, the teacher question and the content question resist the same clean specification. A reformer whose comfort is with specifications and measurements will tend, without necessarily intending to, to shift the centre of gravity of her work

toward the component where her comfort is highest and away from the components where it is lowest. The resulting reform is often well-resourced on technology and under-resourced on teachers and content, and then fails for the reasons the earlier chapters have specified. I name this pattern not because it is anyone's fault in particular but because it is a specific cognitive tendency that institution heads and policymakers should guard against in their own work. The technology chapter is, in this book, deliberately placed after the content and teacher chapters, with the intention that it be read as the third most important rather than the first.

I want to name the pattern explicitly, because it is so common that it is often invisible to the people inside it. The pattern works as follows. A school or a state signs a contract with a technology vendor for an AI-education product — tablets, a platform, a content library, and the integration of these into a working system. The contract is for three years. At the end of year one, the school has invested substantial teacher training in the vendor's specific platform. At the end of year two, the school has accumulated learning records for all its students in the vendor's proprietary data format, and has adapted its pedagogical routines around the vendor's specific features. At the end of year three, when the contract comes up for renewal, the school discovers that switching vendors would require retraining its teachers, reformatting or abandoning its learning records, and disrupting the pedagogical routines it has built. The cost of switching is, in practice, prohibitive. The vendor knows this. The renewal terms reflect the knowledge. The school renews. And a decade later, the school and the vendor are so entangled that the question of whether the vendor is still the best available choice has become unaskable.

This pattern is not unique to education. It has occurred in every sector where complex software has been procured by institutions that did not, at the time of procurement, understand the long-term dynamics of software dependence. What is specific

to education is the consequence of being captured — not a stranded investment, which is bad enough, but a generation of children whose learning has been shaped by a vendor's commercial priorities rather than by the educational community's judgement. The stakes are high enough that the procurement question deserves serious attention, and this chapter is my attempt to provide it.

I want to add, because the asymmetry is worth naming plainly, that the vendors who operate in this market are not, in most cases, acting in bad faith when they construct the capture arrangements I have described. They are acting in accordance with the commercial logic of their industry, which rewards customer retention and penalises customer loss, and they are using the specific tools of platform design and contractual engineering that their industry has developed for this purpose. The asymmetry is not that vendors are unusually predatory; it is that institutions, particularly educational institutions operating outside the commercial-software world, are unusually unprepared for the dynamics of the relationships they are entering into. The institution head who is negotiating her first substantial technology contract, without the procurement experience and the legal support her vendor counterpart has accumulated over many comparable transactions, is systematically outmatched in the negotiation. The remedy is not moral indignation at vendors but institutional preparation — legal support, procurement expertise, shared institutional learning — that brings the institution's capacity closer to the vendor's. Much of what the mission should offer to institutions is precisely this preparation, in the form of standard contract templates, model procurement specifications, shared vendor-evaluation resources, and peer networks through which institution heads can learn from one another's negotiations.



I. What schools actually need

Let me start with a simpler question than procurement strategy, which is: what does a school actually need, technologically, to deliver the roadmap I set out in Chapter Three? The answer is less than is commonly assumed, and specifying it clearly reduces the space in which vendors can sell schools things they do not need.

At the Foundational-to-Preparatory Stage — Classes 3 through 5 — the school needs, per classroom of approximately thirty children, a small number of tablet-class devices — perhaps six to ten — that can run the AI tutor. The ratio of children to devices is, at this stage, not one-to-one. The children take turns with the tutor under the teacher's supervision, working individually or in pairs, while the rest of the class is engaged in other activities. The devices do not need to be high-end. They need to be reliable, durable enough for daily handling by small hands, and capable of running the tutor's software, which is generally not demanding. They need to be rechargeable in the classroom, which requires a small amount of electrical infrastructure that many government primary schools currently lack and will need as part of the reform's deployment. They need basic network connectivity, ideally wireless, to sync the child's learning records with the school's server and the district's dashboard. The specifications are not exotic. A well-built Indian-manufactured tablet in the four to six thousand rupee range is entirely sufficient for this stage, and a reform that insists on higher specifications is a reform that has allowed vendor marketing to drive its requirements.

At the Middle Stage — Classes 6 through 8 — the device ratio moves toward one-to-one, and the device's capability requirement rises. The student is now training simple machine-learning models, working in a block-based or text-based coding environment, and interacting with richer content. A Chromebook-class device, or a comparable Indian-manufactured laptop, in the fifteen to twenty-five thousand rupee range, is suitable. The devices are shared

between classes within the school rather than assigned individually to each student, with a system for distribution and collection that the school manages. The network requirements rise — the student's work, including her code and her datasets, now needs to be reliably synced, and the latency of the interactions matters for the quality of her experience. Many schools will need, at this stage, upgraded connectivity — a stable broadband connection, in most places, rather than the mobile-network fallbacks that have sufficed at the Foundational Stage.

At the Secondary Stage — Classes 9 through 12 — the student who is taking the Applied AI elective in particular needs a device with enough computational power to train small-to-medium models and run a development environment, which places her in the thirty to fifty thousand rupee device category. The device-to-student ratio for the electives approaches one-to-one during elective periods, though schools can manage the shared-device approach across subjects outside the elective. The connectivity requirement is firmly broadband-class, and the school's network infrastructure must support the simultaneous work of fifty or sixty students without the latency and dropout that would make serious work impossible.

Summed across the three stages for a typical government school of approximately one thousand students, the device requirement is of the order of two hundred to three hundred devices, of varying specifications, with a replacement cycle of approximately four years. At current Indian-manufactured prices, the capital cost per school is of the order of fifteen to twenty-five lakh rupees, plus infrastructure costs of three to five lakh for connectivity and power, plus ongoing costs of two to three lakh per year for maintenance, replacement, and connectivity charges. These numbers are large at the school level and very large at the national level — aggregated across approximately 1.5 lakh government schools, the capital cost is of the order of thirty to forty

thousand crore rupees, which stands as the largest single fiscal line of the reform.

I give these numbers not because they are final but because the reader — particularly the institution head — is entitled to a specific picture of what the reform's technology footprint looks like. Many institution heads, in conversations I have had, have been quoted figures by vendors that are substantially larger than these, and have then concluded that the reform is unaffordable for their institutions. The reform is, at the numbers above, difficult but affordable. It is unaffordable at the numbers vendors quote, because vendors quote for the product they would like to sell rather than for the product the child actually needs. The institution head's first discipline, in technology procurement, is to specify what she needs independently of what vendors are offering, and then to evaluate vendor offerings against her specification rather than against each other.

A specific note on the electrical and network preconditions, because the device question is often treated as the whole technology question when, in practice, the device is the smallest part of the problem. Approximately one in four government primary schools in India does not have a reliable electrical connection — reliable, in this context, meaning a connection that delivers power for more than six hours of the school day, five days a week, across the school year. A smaller but still substantial proportion lacks any electrical infrastructure at all, or has a nominal connection that is not functional because of unpaid bills or disconnected meters. The reform's device deployments cannot proceed in these schools without a specific infrastructure investment that precedes the device — a solar installation, a battery backup, or, where the grid is present but unreliable, a hybrid arrangement that combines grid and solar supply. The cost of the electrical infrastructure, per school, is of the order of one to three lakh rupees in the schools that need a full solar deployment, and somewhat less where the existing grid supply is improved rather

than replaced. Multiplied across the schools that need it, the aggregate electrical-infrastructure cost is substantial – somewhere between eight and twelve thousand crore rupees across the five-year deployment – and a reform that does not budget for it will find itself deploying devices into schools where the devices cannot be charged. The same arithmetic applies to the network. Approximately forty percent of government schools do not have a broadband connection of any kind. A further twenty percent have a nominal connection whose effective bandwidth is inadequate for the concurrent use of the devices the reform envisions. The network upgrade is its own substantial component of the fiscal envelope, and the procurement should treat the network, the power, and the devices as a bundled deployment rather than as three separately procured items, because the coordination of the three is what actually produces a functioning classroom.


A separate observation, for the institution head in a state where the government will not move quickly on the infrastructure side. The temptation, in such a state, is to proceed with the device procurement and to hope that the infrastructure will catch up. I would strongly advise against this sequence. A device procured but unable to be charged or unable to connect is a device that, within six months, has been quietly written off by the school that received it. The classroom returns to its pre-deployment state, with the added damage that the teachers have been taught, through direct experience, that the reform is incapable of delivering what it promised. An institution head in a weaker-infrastructure context should treat the electrical and network preconditions as the first phase of her own implementation plan, securing them through her own efforts or through targeted approaches to the state's infrastructure departments, and should only move to device procurement once the preconditions are demonstrably in place. The sequence is slower. It is, in every significant respect, the sequence that produces a working classroom.



II. The platform question

Alongside the devices, the reform requires a software platform through which the AI tutors, the content, the assessment tools, the teacher dashboards, and the administrative functions all operate. The platform is where the capture pattern I described at the start of this chapter most frequently occurs, because platforms are where the switching costs accumulate most steeply.

The principle I would argue for, following the Digital Public Infrastructure approach that has shaped Aadhaar, UPI, and DigiLocker, is that the platform layer for AI education should be an open public infrastructure, not a vendor-specific proprietary system. Specifically, the Centre — through the National AI Education Mission I propose in Chapter Ten — should operate a thin public platform layer, open-sourced under an Indian government licence, that any vendor must integrate with to sell into government schools. The public layer handles the things that must be standardised — learner-record format, teacher-dashboard specification, inter-operability between content from different sources, privacy and safety audit logging, and the core authentication and access control that schools require. On top of the public layer, private vendors build the specific tools, tutors, and content offerings that compete in the market. A school or a state chooses, among approved vendors, which tools it prefers, knowing that its data, its records, and its pedagogical routines remain portable across vendors because they conform to the open public standards.



Switching costs eliminated. Data portable by mandate. Vendors compete on merit, not lock-in.

Figure 5.1 | The DPI architecture — a thin public layer, private innovation on top, no vendor lock-in.

This architecture has specific virtues that proprietary alternatives do not. The switching costs that trap schools today are eliminated, because a school can change vendors without losing its learning records or retraining its teachers on an entirely new system. The privacy and safety standards are enforced by the public layer rather than depending on each vendor's voluntary compliance. The pricing competition among vendors becomes genuine, because vendors can no longer extract rents from captured customers. The innovation in specific tools is preserved, because vendors compete on what they build on top of the public layer, and the best tools rise to the top through market selection rather than through procurement politics.

The DPI approach is not new. It is what India has already done, with considerable success, in other domains. The argument of this chapter is not for a novel architecture but for the application of a proven Indian architectural pattern to the specific domain of AI education, where it has not yet been applied and where its absence is producing the same vendor-capture patterns that it has resolved elsewhere.

The platform's development is itself a substantial undertaking — perhaps one to two thousand crore rupees of initial investment and of the order of three to five hundred crore per year of ongoing operations. The investment should be undertaken by the Mission, in partnership with NCERT for the pedagogical specifications and with one or more of the technical institutions the Centre has capacity in — the CDAC, the NIC, or the consortia around the IITs — for the engineering implementation. The platform's development is not the kind of work that should be contracted to a single private vendor for construction, because the platform's value depends on its being a public good rather than a commercial asset, and the governance arrangements that protect its public-good character are best constructed from the outset.

I want to say a specific word about the platform's long-term operations, because the construction of the platform is the easier half of the problem. The harder half is the ongoing operation, maintenance, and evolution of the platform across the decades of its useful life. The DPI platforms that India has built well — Aadhaar's authentication infrastructure, UPI's payment rails, DigiLocker's document store — each depend, for their continued functioning, on a specific operational institution with the technical capacity, the budget, and the organisational continuity to sustain them. The AI-education platform will require the same. The specific question is which institution operates it. The options are the Mission itself, through an internal technology directorate; an arms-length technology body within the government, along the lines of the NIC or CDAC; a section 8 not-for-profit company along the lines of NPCI, created specifically for this purpose; or, as a fourth option, an operation split between a public-sector technology anchor and a rotating set of private contractors. My preference, offered tentatively because the question is contested, is the section 8 company model — the entity that has produced UPI's sustained operational excellence and that has proved most resilient to the political and bureaucratic pressures that the alternative models struggle with. The Mission would be the sole public-

interest shareholder, the section 8 company would operate the platform under a long-term service agreement, and the governance would include representation from state governments, from the educational community, and from the private sector in roles that are advisory rather than directive. The institutional design of this operating entity is itself a substantial piece of work, and it should begin within the Mission's first year rather than being deferred to year two or year three.



III. The AI tutor question

At the heart of the reform's classroom experience sits the AI tutor. The tutor is the software, typically running on the child's device, that directly interacts with her — that listens when she reads aloud, that answers her questions, that offers her practice problems calibrated to her current level, that notices when she is struggling and when she is bored. The quality of the tutor is the single most consequential determinant of whether the reform's classroom experience is excellent or ordinary.

I want to say a few honest things about AI tutors as they currently exist, because the marketing around them has run somewhat ahead of the reality. The best AI tutors available today — including several built by Indian companies — are genuinely useful tools. They can maintain a child's engagement over extended practice sessions. They can adapt to the child's pace in ways that a teacher handling thirty-five children at once cannot. They can provide, at the Foundational Stage, the repetition and individualised practice that drill children to fluency in ways classrooms rarely achieve at scale. They make real differences, measurably, to learning outcomes in the specific areas — foundational literacy, arithmetic practice, language learning — where the evidence base is strongest.

They are also not, in their current state, the transformational educational intelligences that some of the marketing material implies. They do not replace the teacher. They do not replace the curriculum. They do not, in most cases, handle the full complexity of a subject's pedagogy — they handle specific components of it well, leaving the broader pedagogical work to the teacher. They sometimes make mistakes, some of which are subtle enough that a child will not notice them and a teacher will need to catch them. They can be gamed by children who figure out how to produce the responses the tutor rewards without doing the underlying work. They raise privacy and safety questions that the current regulatory framework only partially addresses. They work better in some languages than in others, with the gap between English and Indian-language performance narrowing but still real.

The honest summary is that AI tutors, used well, are a substantial positive contribution to the classroom, and the reform is right to deploy them at scale. The honest summary also includes the observation that they work best when integrated into a broader pedagogical design rather than deployed as self-sufficient solutions, and that the quality of the integration matters as much as the quality of the tutor itself. A mediocre tutor well-integrated into a thoughtful pedagogy produces better outcomes than an excellent tutor deployed without such integration. This is the pedagogical layer at which the teacher's skill becomes decisive, which is among the reasons the teacher chapter is the book's centre.

I should also note a specific class of concerns about AI tutors that sits at the intersection of child safety and product design, and that the reform's safety code will need to address explicitly. The current generation of consumer-facing AI products is, in its behavioural design, substantially shaped by the attention-economy incentives that shape consumer technology more broadly — engagement metrics, session length, daily active user counts, and the various behavioural affordances that extend these metrics. Some of these affordances are innocuous. Some are harmful when

applied to children. An AI tutor that, as a design choice, sends a notification to the child when she has not opened the app for three days, or that offers gamified rewards for consecutive-day usage, or that engages the child in open-ended conversational extensions beyond the pedagogical purpose, is an AI tutor that has been designed, in part, to produce engagement for its own sake rather than to serve the child's learning. The reform must be explicit that these affordances are not permitted in AI tutors deployed into government schools, and that the vendor's product design — not merely its pedagogical content — is subject to review. The specific boundary between engagement that serves learning and engagement that serves the product's own metrics is not always easy to draw, but the attempt to draw it, rigorously, is part of what distinguishes a child-centred AI tutor from a consumer product that has been repackaged for the classroom.

For procurement, the specific discipline I would recommend is that the Mission issue a single national framework agreement, after rigorous technical and pedagogical evaluation, under which five to eight approved vendors are qualified to supply AI tutors that meet the standards in each of the major Indian language groups. States and schools procure from this pool, with prices pre-negotiated at the national level. This model provides school choice while preventing single-vendor lock-in, drives down unit prices through scale, and ensures every deployed product meets the same quality and safety threshold. An institution head who is deciding, today, which AI tutor to pilot should look for vendors whose offerings are compatible with the kind of national framework I am describing, because those are the vendors whose products will survive the coming procurement consolidation. An institution that commits to a vendor whose architecture is incompatible with the DPI approach is an institution that may find itself, in three years, re-platforming its pilot onto a compliant alternative.



IV. A word on large language models

I want to touch, briefly, on a question that sits somewhat underneath the procurement question I have been discussing — the question of which large language models the AI tutors the reform deploys will actually be built on. This is not a question the institution head needs to resolve herself, but it is a question the Centre must think about, and one the reader is entitled to understand because it will shape the tutors she encounters.

The large language models that underpin most of the best AI tutors available today are, overwhelmingly, produced by a small number of American and Chinese companies. The European and Indian industries are playing catch-up, in different ways and with different levels of state support. India's own foundation-model ecosystem has grown substantially in the past two years, with the IITs' work, Bhashini's developments, and several private Indian labs each contributing. Indian models are, in most specific benchmarks, still behind the best American models, but the gap is narrower than it was eighteen months ago and is, on current trajectories, closing further.

The choice facing the Centre is whether to build the reform's AI tutors on the best available models, wherever they come from, or to privilege domestic models and accept some capability gap in exchange for sovereign control of the infrastructure. My view, which I offer with the humility a contested question deserves, is that the first path better serves Indian children in the short run, provided three conditions are met: that foreign models operate under strict data-residency and audit requirements when deployed in Indian schools, that the Centre invests with genuine seriousness and over a sustained decade in domestic model development, and that the reform retains the optionality to transition to domestic models when they become competitive. The strict data-residency condition is non-negotiable. No Indian child's interactions with an AI tutor should leave Indian jurisdiction in a form from which the child can be re-identified, regardless of the technical convenience

of offshore processing. The audit condition is non-negotiable. The Mission, or the statutory authority that enforces the Children's AI Safety Code I propose in Chapter Sixteen, must have the right to audit the models for the specific behaviours and biases that matter in the educational context. The investment condition is non-negotiable. A Centre that uses foreign models today while failing to develop its own has made a strategic choice it will regret by the 2030s.

I want to add one further consideration, which is the specific question of how the foreign-domestic transition should be managed operationally, because the principle is easier to state than the arrangements that implement it. If the reform begins, as I expect it must, with foreign foundation models as the technical backbone of many of the AI tutors deployed in the first wave, the transition to domestic models over the subsequent decade will need to be an explicit programme, not a hope. The Mission should, from its first year, commission an ongoing capability-benchmark exercise, in which the leading domestic foundation models are evaluated against the leading foreign models on the specific tasks that matter for AI tutors in Indian classrooms — reading comprehension in each of the major Indian languages, arithmetic reasoning at age-appropriate level, conversational responsiveness within pedagogical scripts, factual accuracy on the specific subject-matter curricula the tutors support. The benchmark is published. It tracks, quarter by quarter, the closing or widening of the capability gap. When a domestic model reaches parity on a specific task, the Mission signals to vendors that tutors serving that task should, within a defined transition period, migrate to the domestic model. The transition is managed rather than coerced — vendors are given time, technical assistance, and, where appropriate, financial support for the migration work — but it is treated as a commitment rather than as an aspiration. Without this specific programme, the transition will not happen. The incumbent foreign-model arrangements will become entrenched by the sheer inertia of deployment, and the domestic industry will struggle to build at

scale in the absence of a clear demand signal from the reform. With the programme, the transition has a realistic trajectory. The details of the programme's design are matters for the Mission and for MeitY to work out together. What matters, for the reader of this chapter, is that the transition is treated as a first-order design question from the beginning, rather than as a matter to be addressed when the time comes.

The reform should not, in my view, be delayed while the domestic model ecosystem catches up. The children who are in Class 3 today will not wait. The right path is to proceed now with the best models available under the disciplines I have described, while building, in parallel, the domestic capability that will replace the foreign dependence when it is ready to do so.



V. What the institution head should actually do

I want to close the chapter with a specific set of suggestions for the institution head who is making technology procurement decisions for her own school or college, before the national framework I have been describing exists. The Mission I am proposing does not yet exist. The national framework agreements do not yet exist. The Repository does not yet exist. The institution head, today, must make decisions without the infrastructure that an adequate reform would provide, and she is entitled to specific guidance on how to make those decisions well.

First, do not commit to any vendor for a period longer than one year, no matter how attractive the multi-year pricing appears. The field is moving rapidly enough that a three-year lock-in signed today is a lock-in you will regret within eighteen months. One-year contracts renewable at the institution's discretion, with explicit exit provisions and data portability requirements, are the right structure for this stage of the reform's development.

Second, insist, in writing, on your ownership of your students' learning data, and on your right to export that data in a standard format at any time. If a vendor's contract does not grant you this ownership and this export right, reject the contract. The vendor's reluctance to grant the right is a reliable indicator of their intent to capture you.

Third, pilot with multiple vendors in parallel where possible, rather than committing institutionally to one. A pilot with three vendors across three classrooms for one year is more expensive than a pilot with one vendor across three classrooms, but the comparative learning it produces is worth the cost several times over, and the diversity of experience protects the institution against a single vendor's specific weaknesses.

Fourth, resist the temptation to move first on the most visible technology components — the tablets, the platform — and underinvest in the less visible components that will determine whether the visible ones work. The teacher training, the content localisation, the network infrastructure, the technical support the school needs to keep the deployment running — these are the unglamorous components that determine the pilot's success or failure. A pilot that looks good in a minister's visit but falls apart after six months is worse than no pilot at all, because it teaches the school that this kind of reform does not work. The unglamorous components should receive as much attention, and in most cases more budget, than the glamorous ones.

Fifth, document your pilot's successes and failures honestly, and share your documentation with other institutions. The institutional learning that the national reform needs will, in large part, come from institutions like yours that piloted early and reported back on what they learned. Your documentation is a public good, and is worth more to the reform than any specific technology choice you make.

A specific note on procurement documents, because several institution heads have asked me for something more concrete than the principles above. I will not, in this book, offer a model RFP — a model document requires the specific institutional context in which it will be used, and a generic template encourages the lazy adoption that serves no one. What I will offer is the short list of clauses the institution head should ensure her RFP contains. First, a clause specifying that the institution owns all student-generated data and retains the right to export it in a specified standard format at any time. Second, a clause specifying that the vendor's source code for any custom-built components is held in escrow, accessible to the institution if the vendor's business ceases or the contract is terminated for cause. Third, a clause specifying that the vendor will comply with the Children's AI Safety Code and with any subsequent regulations issued under the statutory safety authority I propose in Chapter Thirteen. Fourth, a clause specifying the specific service-level metrics the institution will hold the vendor to — uptime, response times, content-update frequency — with financial consequences for non-performance. Fifth, a clause specifying the termination and transition arrangements, including the vendor's obligation to support a transition to an alternative provider for a specified period after termination. These five clauses, in my observation, are the five that vendors most commonly resist including and that institutions most commonly fail to demand. Their inclusion is the minimum contractual discipline the reform requires.

And, on the specific question of piloting in partnership with peer institutions, I would offer one more suggestion. Find, in your city or in your region, two or three other institutions whose leaders are equally serious about this work, and consider running a coordinated pilot across your institutions. The coordination does not need to be heavy — a monthly meeting, a shared documentation template, an agreement to benchmark against each other honestly. The discipline of comparison with peers is the single most effective institutional corrective for the self-flattering

tendencies that afflict most school-level evaluations. An institution that reports to itself alone will, over time, develop a version of its own performance that is more comfortable than accurate. An institution that reports, monthly, alongside two or three peer institutions will find that the comparison surfaces problems it would otherwise have missed, and will push its own standards higher than it would have set for itself. The coordinated pilot is, in effect, a small-scale version of the accountability discipline the mission aspires to build at national scale. Building it at your own institutional scale is good practice for what is coming, and it produces immediate institutional benefit.



Three decisions for the reader

1. If you are an institution head, write your technology specification independently of any vendor's offering.

Before you take the next vendor's call, write down what you actually need — specific device types, specific numbers, specific platform capabilities, specific content requirements, specific data-portability standards. Base the specification on the child's learning requirements, not on the vendor's product catalogue. When vendors pitch to you subsequently, evaluate each pitch against your specification. The specification is yours, and the vendor's job is to meet it, not to change it.

2. If you are a policymaker, commission the DPI platform within the next eighteen months.

The public platform layer is the single most important infrastructural commitment the Centre can make on the technology side of the reform. Without it, the market will fragment across incompatible vendor platforms and the capture patterns I described at the start of this chapter will become endemic. The

commissioning decision involves the Ministry of Education, MeitY, the Mission, and whichever technical institutions will implement the platform. The decision can be taken quickly if the political will is present. I would urge that it be taken within eighteen months of this book's publication, with a target of initial deployment within three years.

3. If you are a parent, ask your child's school for its technology procurement policy.

Most schools do not have one. When you ask, and the school cannot produce one, you have identified a gap the school needs to close. When the school can produce one, read it carefully. Ask specifically about the data your child is generating through her AI-education tools, where that data is stored, who has access to it, what it is used for, and what happens to it if the school changes vendors. A school that cannot answer these questions is a school whose technology procurement is not being managed with sufficient care, and your question, asked publicly, is what will cause the answers to be developed.



END OF CHAPTER FIVE

CHAPTER SIX

The Teacher Question



On the nine and a half million Indians on whom this reform will succeed or fail.



Three teachers. Three states. Three versions of the same question, which is the question at the centre of this book.

Kusum Devi teaches Class 3 at the government primary school in Semra village, Barabanki district, Uttar Pradesh. She is forty-seven years old. She qualified through a two-year diploma at a District Institute of Education and Training in 2011, and she has taught at this same school for nine years. Her monthly salary, after nine years of service, is forty-eight thousand rupees. She teaches all subjects to Classes 3 and 4, with a combined class strength that varies between sixty and seventy-two children depending on the time of year and the condition of the classroom roof. She has never used an AI tutor. She has heard the term artificial intelligence on the radio and in conversations at the weekly Block Resource Centre meetings, and her impression of what AI is comes primarily from a short demonstration by a vendor who visited her school last April and left within twenty minutes without finishing the demonstration because the internet connection failed. When I asked her, last November, what she would need in order to teach her children about artificial intelligence, she answered, after a long pause: I would need someone to teach me first. And I would need to understand why you are asking me to do this. I have quoted her reply in every conversation I have since had with officials about this reform.

Renu Balakrishnan teaches Class 7 mathematics at a government-aided school in Kozhikode district, Kerala. She is thirty-nine years old. She qualified through a B.Ed. from the University of Calicut in 2009 and has taught at her current school for twelve years. Her monthly salary is sixty-two thousand rupees. Her school has, since 2019, received computers under the state's KITE programme, and she has attended three week-long professional development programmes on the use of digital tools in mathematics teaching, two of them organised by the SCERT and one by KITE itself. When I asked her what she needed in order to integrate AI tools into her teaching, she said: I need time to experiment. I need two colleagues I can compare notes with. And I need the principal to trust that we will make mistakes. The combination of specificity and self-respect in her answer made the distance between Kerala's teacher ecosystem and Uttar Pradesh's more concrete for me than any statistical comparison has.

Moirangthem Brojeswor teaches Physics at a higher-secondary school in Imphal West, Manipur. He is fifty-three years old. He qualified through an M.Sc. at Manipur University and a B.Ed. that followed two years later. He has taught at various schools in the state for twenty-seven years, with an interruption of sixteen months between 2023 and 2024 when the ethnic violence that displaced hundreds of thousands of his fellow Manipuris also closed his school. When I asked him what he needed in order to teach his students about AI, his answer had a bleakness the other two lacked: I need my school to be a school again. And I need to know whether the future my students are preparing for is a future that will include them. A reform that does not understand, in a specific sense, why Brojeswor's question is the right question to ask — and why the answer cannot be supplied by a curriculum document or a device-deployment plan — is a reform that has failed the Manipuri child it is ostensibly also for.

Three teachers. Three answers. This chapter is about all three of them and about the nine and a half million Indians, of whom

they are three, on whom the reform this book proposes will succeed or fail. The argument of the chapter is that the current national conversation about AI education in India does not take teachers seriously as the load-bearing agents of the reform — treats them, rather, as a constraint to be managed, a cost to be minimised, or a tool to be equipped — and that this failure of framing is the single most consequential risk the reform faces. The argument proceeds in eight sections. I warn the reader that this chapter is long. It is long because it is the chapter on which the rest of the book depends.

I want to name, before the sections begin, the specific register in which the Indian policy conversation has historically discussed teachers, because the register shapes what the reform can become. The register is, at its most common, the register of grievance management. Teachers are discussed as people who have complaints about pay, about transfers, about working conditions, about the administrative burdens placed on them, and the policy response has been, for decades, a succession of accommodations that address the complaints piecemeal without reconceiving the profession's place in the broader educational architecture. A reform that enters this register — that sees the teacher-capacity agenda as a set of accommodations to be made to a restive workforce — is a reform that will produce incremental improvements to the profession's working conditions without producing the substantive partnership the AI-education work requires. The alternative register, which this chapter tries to establish, is the register of professional respect. The teacher is the institution's primary professional. Her pedagogical judgement is the specific asset the institution is employing her for. Her engagement with the reform is not a concession to be extracted but a professional contribution to be invited, and the terms of the invitation are the specific design question this chapter addresses. The registers are not fully separable in practice — any serious engagement with the profession will involve addressing the grievances that have accumulated over decades of institutional neglect — but the

primary framing matters, because it shapes which questions the reform asks and which answers it can accept.



I. The scale of the teaching workforce, and what it implies

The Indian school system employs approximately nine and a half million teachers across its government and government-aided institutions, with a further approximately two and a half million in the recognised private sector. The exact number varies by which source one consults and which categories one includes — contract teachers, para-teachers, honorarium workers — but the order of magnitude is settled. This is the third-largest civilian workforce in the country, after agriculture and construction, and it is larger than the entire working population of Singapore, Finland, and New Zealand combined. Any reform that depends on this workforce's action at scale is a reform whose scale matters as a first-order constraint, not as a practical detail to be handled in implementation.

I want to sit with this scale for a moment, because the statistical framing makes it easy to move past the specific consequence of the number. Nine and a half million teachers is, in rough terms, one teacher for every twenty-six Indian children of school age. It is also, distributed across the country, approximately two hundred and seventy teachers per average parliamentary constituency, or fifteen thousand teachers per average state — numbers that correspond to substantial professional communities in every part of the country, each with its own internal dynamics, its own informal hierarchies, its own professional conversations. The reform's engagement with this workforce is not, therefore, an engagement with a centrally-addressable category but with a set of thousands of local professional communities that must each, in its own context, be brought into the reform's work. The implication for design is that

the reform must have, at every level of its delivery chain, specific mechanisms for engagement with the local professional community that corresponds to that level – the SCERT must have specific mechanisms for engaging with the state's teaching community as a whole, the DIET must have specific mechanisms for engaging with the district's, the BRC with the block's, and the school-level coordinator with her own school. Without this layered engagement, the reform is communicating with an abstract category of teachers rather than with the specific professional communities whose cooperation the reform requires. The design of the engagement mechanisms is itself substantive work, and the mission should treat it as such rather than defaulting to the one-way communication that has characterised most earlier reforms.

The scale has specific implications the reform must respect. The first implication is that no centrally-designed training programme, however excellent, can reach nine and a half million teachers with quality in less than several years. A training programme that reaches a million teachers a year is running at an extraordinary pace by any comparative standard. At that pace, reaching the full workforce takes a decade, which is the realistic horizon over which serious teacher capacity can be built. A reform that expects to equip its teachers in eighteen months is a reform whose pace and whose method are both at odds with the scale it is operating within.

The second implication is that training-of-trainers is the only workable architecture. A central institution cannot train nine and a half million teachers directly. It can, at best, train a cadre of master trainers who then train state-level trainers, who then train district-level trainers, who then work with block-level resource persons to train the teachers in schools. Each level of this cascade introduces loss – concepts that were clear at the top become muddled at the bottom, nuances that required expert judgement at the centre become mechanical rules by the time they reach the classroom. The loss is not eliminable; it is reducible, through careful design of each

cascade level, but never eliminable. The reform must be designed to be robust against the loss, which means specifying the core content in a form that can survive the cascade and accepting that some nuance will be lost in exchange for reach.

The third implication is that the in-service capacity development of the existing workforce, while essential, is not sufficient. The reform must also address the pre-service preparation of the next generation of teachers, because the teachers who will be in classrooms in 2035 are, in substantial part, those entering B.Ed. programmes now. A reform that focuses only on upgrading the existing workforce will find, in a decade, that it has produced a workforce that is AI-competent in its veteran segment and AI-illiterate in its junior segment, with the latter growing year on year. The B.Ed. curriculum reform is, for this reason, as important as the in-service training programme, even though it is less politically salient in the short term and shows no visible impact until four to six years after it is initiated.

The fourth implication, and the most consequential, is that the reform must respect the teaching workforce's diversity. The nine and a half million are not a homogeneous category. They range from the Class 3 teacher in a remote village who has limited English and limited prior exposure to digital tools, to the higher-secondary physics teacher in a metropolitan private school with a doctoral degree and decades of international professional engagement. A single training programme designed for the median teacher will underserve both ends of this distribution – wasting the veteran's time with content she already knows, overwhelming the remote teacher with content she cannot yet absorb. The reform requires differentiated pathways, at multiple levels of entry, with clear progression routes between them, and this differentiation is a design discipline that the current one-size-fits-all training tradition does not readily accommodate.

There is a fifth implication, less often stated, that deserves to be placed alongside the four I have named. The teaching workforce

is aging, in the specific sense that a substantial proportion of current teachers will retire within the ten-year horizon of the reform. In some states — Uttar Pradesh, Bihar, Madhya Pradesh among them — the proportion of teachers within ten years of superannuation exceeds a third of the total workforce. A reform that invests the bulk of its training resources in the veteran teachers who are near the end of their careers will produce a short-term capability that retires alongside its bearers. A reform that invests primarily in the younger teachers, who will still be in classrooms in 2040, will produce a slower initial deployment but a durable generational shift. The right balance is a matter of judgement, and the balance should differ by state according to the age profile of its workforce, but the question must be asked explicitly in each state's implementation pathway rather than being ignored. A state that does not know its teacher age profile in any usable detail — and most states do not, today, hold this data in a form that can inform policy — is a state that is operating its teacher-capacity programme without one of the basic planning inputs it requires.

I want to add, in the same spirit, a specific observation about contract teachers — the approximately one million teachers who, across the Indian school system, work on short-term or renewable contracts rather than as permanent government employees. These teachers bear, in many states, a disproportionate share of the classroom teaching load in the schools most distant from state capitals, with pay substantially below their permanent colleagues' and with job security that is, in most cases, uncertain from year to year. The reform's teacher-capacity agenda, if it is to be credible on equity, cannot exclude them. But the reform also cannot fully resolve, within its own mandate, the structural question of why contract teaching exists as a separate category in the first place, which is a question about state-level teacher-recruitment policy that sits beyond the reform's remit. The compromise I would propose is that contract teachers receive the same training and certification entitlements as permanent teachers, without discrimination, and that the state's own implementation pathway

include, as a specific commitment, a time-bound programme for regularising contract teachers whose training is completed to the required standard. The compromise is imperfect. It is also the form in which the reform can engage with the contract-teacher question without resolving the larger policy debate that must be resolved elsewhere.



II. What teacher preparation actually requires for this reform

Let me now specify, as concretely as I can, what serious teacher preparation for AI education looks like. The specification has six components. It is not the kind of specification that can be delivered in a two-day workshop.

The first component is pedagogical competence with AI as a classroom tool. The teacher needs to know, at a working level, how to integrate an AI tutor into her classroom routines. This includes practical matters — how to distribute devices, manage screen time, handle technical failures — and pedagogical matters — when to let the tutor work one-on-one with a child, when to pull the class together for a collective discussion, how to use the tutor's output as a teaching moment rather than a conclusion, how to build on what the tutor has done in the minutes after the child closes it. The competence is not technical. It is pedagogical. It is the competence of a skilled teacher using a new tool well, which is a specific professional skill that takes focused development to build.

The second component is subject-matter literacy in AI itself. The teacher needs to know, at a level appropriate to the grade she teaches, what artificial intelligence actually is. She needs to be able to explain, in her own language, what machine learning means, how a classifier is trained, why a generative model sometimes produces plausible nonsense, what data-shaped bias looks like, and what the difference is between the AI tutor's genuine

competence and its surface performance. She does not need to be able to write Python code or train a neural network from scratch — that is a requirement for the Applied AI elective teacher at the secondary stage, not for the foundational teacher. But she needs enough genuine understanding of the technology that she can answer her students' questions without deflecting, can recognise when the tutor is wrong, and can teach the social and ethical themes the fourth transversal theme of the curriculum requires. Most Indian teachers do not currently have this subject-matter literacy, and building it is a substantial undertaking separate from the pedagogical training.

The third component is data literacy, parallel to what the middle-stage students are being taught but at the level appropriate to adults who must teach it. The teacher who is herself uncomfortable with charts, averages, and probabilistic reasoning cannot effectively teach these to her students. Data literacy development for teachers has rarely been done at scale in India, and the reform provides both an opportunity and a necessity to address this gap.

The fourth component is facility with the social and ethical dimensions of AI. This is perhaps the most difficult component to develop through training, because it requires the teacher to have formed her own considered views on contested questions and to be comfortable facilitating classroom discussions on topics where she does not have, and does not claim to have, the final answer. Indian pedagogical tradition, with important exceptions, has not strongly emphasised open-ended discussion-based teaching, and many teachers find it uncomfortable. The training programme must create both the skill and the space for this form of teaching, and the reform's success on the fourth theme depends on whether it succeeds.

The fifth component is comfort with ongoing professional learning. The AI of 2028 will not be the AI of 2025. A teacher who completes her AI-education training in 2027 and does nothing

further for a decade will, by 2037, be teaching concepts and tools that have become dated. The reform must build, from the outset, the expectation that AI-education teaching is a continuously evolving practice, supported by ongoing peer learning, periodic refresher training, and access to current resources. This is a cultural shift — Indian teaching practice has historically treated training as a one-time event rather than a continuous practice — and the shift is itself a component of the teacher preparation the reform requires.

The sixth component is the development of the teacher's confidence and agency. This is the component hardest to specify and most often neglected. A teacher who has attended training but does not feel authorised to adapt what she learned to her classroom is a teacher who will, in practice, do little with her training. A teacher who feels her agency is respected — who believes her judgement is trusted, that her adaptations are welcomed, that her failures will be treated as learning opportunities rather than occasions for censure — is a teacher who will use her training to produce outcomes the training itself did not specify. The confidence and agency are built not through the training programme but through the surrounding system of how teachers are treated, evaluated, paid, and protected. I return to this in Sections V and VI.

Six components. Each is a substantial undertaking. The total teacher-preparation requirement for the reform is, when laid out at this level of specificity, one of the largest professional-development undertakings any country has attempted in any field. The cost is the cost of doing this preparation seriously. A reform that proposes to do it for less is a reform that is not planning to do it.

I want to sketch, very briefly, what a serious training cadence looks like for a teacher in her first year of the reform, because the six components are easier to understand as categories than as an actual week-by-week experience. The teacher begins with a residential programme of eight to ten days at her BRC or DIET, in which the pedagogical and subject-matter components are

introduced, the classroom tools are demonstrated, and the teacher has her first supervised hands-on practice with the AI tools she will use in her own classroom. She returns to school, begins using the tools in her own classroom under the support of the school-level AI-education coordinator and with weekly visits from a master trainer for the first month. After approximately six weeks, she attends a three-day follow-up programme at her BRC, at which the problems she has encountered are discussed, the corrections are taught, and the peer-learning networks are strengthened. Over the next six months, she participates in monthly peer-learning sessions at her BRC, with online supplementary modules available through DIKSHA and the Institute's digital platform. At the end of the first year, she takes the certification examination – a combination of written assessment, practical demonstration, and portfolio review – and, if she passes, receives the formal AI-education certification and the associated compensation increment. The first year is deliberately heavy. The subsequent years, maintenance of certification through ongoing professional-learning engagement, are substantially lighter. The cadence I have described is not, within Indian training practice, unprecedented – the master-teacher programmes that several states have run for mathematics and science teaching have adopted similar cadences at smaller scale – but it is a substantially heavier engagement than most teachers have experienced in any previous national training initiative, and the reform must budget, in money and in teacher time released from classroom duties, for what it actually involves.




Figure 6.1 | Six components of teacher preparation – each substantial, none reducible to a workshop.



III. The National AI Teacher Institute – a specific proposal

The scale of the teacher preparation requirement implies, I believe, a specific institutional response. I propose that the Centre establish, through the National AI Education Mission, a dedicated National AI Teacher Institute with a single institutional mandate: the development of AI-education teaching capacity across the nine and a half million teachers of the Indian school system.

The Institute is not a new training organisation competing with NCTE, NIEPA, NCERT, and the state SCERTs. It is a coordinating body, anchored at the national level, that works through the existing teacher-education infrastructure to achieve coherence at scale. Its legal home is a statutory body – ideally, a body created by the Parliament Act that also creates the Mission, though an autonomous society under the Ministry of Education

would be an acceptable second-best — with an independent budget, an independent hiring authority, and the capacity to issue certifications that are recognised nationally.

The Institute's core functions are five. First, the development of the national training curriculum for AI-education teaching, at each stage and for each subject specialisation, in consultation with NCERT and the state SCERTs. Second, the accreditation of training providers — public, private, and civil-society — who wish to deliver the training under the Institute's standards. Third, the operation of a master-trainer cadre — perhaps two thousand individuals at full strength — who serve as the top of the cascade and as the editorial conscience of the training programmes delivered down the chain. Fourth, the administration of the certification process, through which teachers who complete the training demonstrate competence and receive the formal certification that attaches to career benefits I will describe in Section V. Fifth, the production and curation of ongoing professional-learning resources for certified teachers, delivered through a digital platform that complements rather than replaces in-person training.

The Institute's governance includes, importantly, a teachers' council — a body of practising teachers, elected from each state, that reviews the training programmes, advises on curriculum, and flags issues from the field. A training organisation that does not include the voice of the teachers it serves is an organisation that will, over time, develop training that is more responsive to institutional politics than to classroom reality. The teachers' council is the structural protection against this drift.

I want to be specific about where the Institute should not be located. It should not be based in Delhi. The temptation, as with all new central institutions, will be to site it in the NCR, where the bureaucratic convenience is greatest. The convenience is precisely the reason to site it elsewhere. Bengaluru, Pune, or Hyderabad — cities with strong technology ecosystems and proximity to

educational-research institutions — would each serve. A city in central India such as Bhopal or Nagpur would serve from the perspective of geographic accessibility for teachers from multiple regions. What matters is that the Institute is close enough to classrooms, culturally, that its staff routinely visit schools rather than summoning school representatives to Delhi for committee meetings. The physical location is itself a form of institutional discipline.

The Institute's initial establishment budget, I estimate, is of the order of a thousand crore rupees over three years, with an ongoing operational budget of perhaps five hundred to seven hundred crore per year. The budget is, within the reform's overall envelope, modest. The return on the budget, if the Institute is well-designed and well-led, is decisive for the reform's success.

The Institute's first director is perhaps the single most consequential individual appointment of the reform's early years, and I want to say a word about what her qualifications should be. The default inclination, in Indian appointments of this kind, is to look for a senior civil servant with an education or technology portfolio, or for a distinguished academic administrator from one of the IITs or the IIMs. Both are plausible candidates, and both have occasionally produced effective institutional leadership in comparable roles. Neither is, in my view, the right default for this particular appointment. The right candidate is a career educator who has actually taught children, who has subsequently led a teacher-education institution or a substantial educational programme at scale, who has working familiarity with the technology question without being a technologist, and who has demonstrated, across a career, the capacity to earn the trust of teachers as a peer rather than as a superior. Candidates of this profile exist — I can name half a dozen, without thinking hard, from among my own acquaintance — but they are not the candidates who emerge most readily from the standard Delhi search process. The Mission's founding leadership must make a

specific effort to identify and recruit them, and the Governing Council must protect the director, once appointed, from the political and bureaucratic pressures that will inevitably attempt to push the Institute in more technology-centric or more administratively-conventional directions. The Institute's success depends, more than on any other single factor, on whether its founding director is the educator the appointment requires or the administrator the appointment default produces.



IV. The delivery chain for teacher capacity

Even with a first-rate Institute at the centre, the reform's teacher preparation must flow through a delivery chain that reaches every one of nine and a half million teachers. The chain I propose builds on existing Indian teacher-education infrastructure rather than attempting to replace it, because replacement would take longer than the reform has, and because the existing infrastructure, for all its well-documented limitations, represents a substantial accumulated capacity that would be wasteful to discard.

At the top of the chain sits the National AI Teacher Institute, producing the national curriculum, training the master trainers, and operating the certification framework. The master trainers — approximately two thousand at full strength, drawn from the most capable teacher-educators across the country — are the direct extension of the Institute's reach.

At the state level, each State Council of Educational Research and Training receives the Institute's curriculum and adapts it to the state's specific linguistic, pedagogical, and cultural context. The adaptation is not optional; a programme that does not accommodate Kerala's strong existing teacher ecosystem will waste Kerala's time, and a programme that does not accommodate Bihar's weaker starting baseline will overwhelm Bihar's teachers.

The adaptation is done under the Institute's supervision, to ensure that quality standards are maintained, but the substantive work is the state's. Each SCERT runs, as a result, its own state-specific training programme that is nationally coherent in its standards and regionally grounded in its delivery.

At the district level, each District Institute of Education and Training becomes the operational anchor for teacher training in its catchment. The DIETs – approximately six hundred and fifty in the country, one per district – are uneven in capacity today, and one of the reform's early deliverables must be a substantial capacity-building programme for the DIETs themselves. A DIET that is not functioning well cannot deliver training well, and the reform cannot proceed evenly if the DIETs in some districts are strong while in others they are barely present. The DIET capacity-building programme is, in effect, a precondition for the teacher-training programme it supports.

At the block level, Block Resource Centres serve as the immediate point of contact for teachers in schools. The BRCs, currently used primarily for administrative communication and occasional pedagogical coordination, become the sites of most of the direct teacher training under the reform. A teacher typically attends training at her BRC, two or three times a year, for periods of several days each, with interspersed peer learning sessions with colleagues from her block.

At the school level, each school designates an AI-education coordinator – typically a teacher who has completed advanced certification and who serves as the first port of call for her colleagues when questions arise in daily practice. The coordinator's role is pedagogical rather than administrative. She is a peer mentor, not a supervisor, and her position is structural rather than a formal promotion. She receives a modest honorarium for the additional responsibility, recognised within the reform's overall compensation structure.

I want to add a specific note on peer learning, because it is the layer of the delivery chain that is most often underinvested in and whose return on investment is, in my observation, the highest. The scheduled training programmes — at the BRC, at the DIET, at the SCERT — do necessary work, but they cannot by themselves produce a teacher who is comfortable and confident with the new tools and pedagogies. The additional work is done in the ordinary weeks of the year, in the informal conversations between colleagues, in the specific problem one teacher brings to another over chai in the staffroom, in the WhatsApp group where a question posed at 8 PM receives three answers by 9 PM from colleagues who have encountered the same issue. This informal peer-learning layer is the layer at which most of the ongoing capability development actually happens, and it is the layer most easily extinguished by an administrative culture that treats training as an event rather than as a continuous practice. The reform must explicitly nurture the peer-learning layer — through the school-level coordinator's role, through facilitated peer-learning circles at the block level, through digital platforms that connect teachers across geographies, through the small time and budget allocations that make sustained peer engagement possible. The reform's annual reports should, alongside the certification numbers and the training-programme completion rates, publish indicators of peer-learning health — the number of active peer circles, the topics they are engaging, the specific problems they are solving. The indicators are harder to measure than the training-completion numbers, but they are more predictive of whether the reform's teacher capacity is actually being built.

This chain — Institute, SCERT, DIET, BRC, school — is a long chain with multiple points of potential failure. The reform must invest in the chain's integrity as seriously as it invests in the chain's components. This means specific mechanisms for quality assurance at each link, for feedback flowing from the classroom back up the chain, for mid-course correction when problems are identified, and for the protection of the chain against the ordinary erosion of

quality that afflicts all long chains in Indian public administration. The mechanisms are not novel — they have been developed, with varying degrees of success, in Indian reform after reform — but their application to the AI-education reform requires specific design choices that I will leave for implementers to make.




Figure 6.2 | *The teacher-capacity delivery chain — from national institute to the classroom.*



V. *The compensation question*

I want to turn now to a question that, in most discussions of teacher capacity in India, is either ignored or mentioned as an afterthought. The question is what teachers are paid, and how the reform's demands on them relate to that pay.

The median monthly salary of a government primary-school teacher in India, across states, is in the range of thirty-five to fifty-five thousand rupees, varying substantially with the state and the length of service. The median for a secondary teacher is somewhat higher, typically in the range of fifty to eighty thousand. A contract teacher, of whom there are approximately a million in the system, typically earns between twelve and twenty-five thousand rupees a month. These figures do not, in most states, represent a middle-

class wage. They represent a lower-middle-class wage that places the teaching profession, in most Indian contexts, below many other graduate-level occupations and well below what is required to attract the most capable graduates into teaching as a primary career choice.

The reform's demands on teachers are, as the preceding sections have specified, substantial. A teacher who engages seriously with the training, integrates AI tools into her classroom, keeps current with the evolving technology, and takes on the additional pedagogical weight the reform implies is a teacher who is doing more demanding work than she was before. A reform that expects this additional work at no additional compensation is a reform that is, in effect, extracting value from teachers whose remuneration is already below what the profession's importance would justify.

I propose, therefore, that the reform include a specific compensation increment tied to AI-education certification. A teacher who completes the Institute's certification and is actively teaching in an AI-education context receives a ten percent increment over her base salary. An AI-education coordinator at the school level receives a further five percent increment for the additional responsibility. The increments are portable across schools within the government system, recognised nationally, and not subject to individual-school discretion. The fiscal cost of this commitment, at full scale across a certified workforce of nine million, is of the order of fifty to sixty thousand crore rupees a year at steady state, which is the dominant component of the ongoing compensation commitment the reform must bear.

I should acknowledge, in the interest of honesty, that a ten-percent increment is at the lower end of what the profession's comparable advanced-practice premium looks like in other countries, and that a stronger case can be made for a higher increment – fifteen percent, or even twenty. I have held to ten because the fiscal and political arithmetic of a higher figure

becomes prohibitive quickly, and because a ten-percent increment that is actually delivered is more consequential than a fifteen-percent increment that is announced and not delivered. But the reader should know that the figure is a compromise with political economy rather than a first-principles specification. A reform that, in its early years, successfully delivers the ten-percent increment on the schedule it promises, and that demonstrates through its own operation the value of the teacher-capacity investment, will be in a stronger position, five years in, to argue for an additional increment tied to advanced certification or to demonstrated classroom outcomes. The incremental path is the path through which the compensation structure of Indian teaching is most likely, over a decade, to move toward what the profession's importance warrants.

I understand this proposal is politically difficult. Teacher pay is one of the most contested fiscal questions in Indian public administration, and a proposal to increase the base compensation of millions of teachers will face resistance from finance ministries, from state governments whose budgets bear the bulk of the cost, and from a public discourse that often treats teacher pay increases with scepticism. I offer the proposal anyway, because the alternative — asking the reform's demands of teachers at their current compensation — is either unrealistic or exploitative, and neither is an acceptable basis for a national programme. The reform can, if the political will exists, make the case that this increment is the single most consequential investment the country can make in the generation of children now in its classrooms.

There is a second compensation question, less visible than the first, that also requires attention. The career progression of teachers in India — the promotion pathway from junior teacher to headmaster to administrative positions — currently rewards seniority and, in some states, qualifications accrued in areas unrelated to teaching quality. A reform that seeks to encourage teachers to invest in AI-education certification should reform the

progression structure to recognise this certification as a factor in promotion decisions. A teacher who has invested three years in serious AI-education training and practice, and whose classroom demonstrates the results, should be a stronger candidate for promotion than a teacher who has coasted on years of service alone.

The broader incentive question extends beyond compensation and promotion. Indian teaching, in most state systems, has a set of informal incentives that shape teacher behaviour at least as powerfully as the formal ones. Among the informal incentives are: the non-monetary recognition that comes from colleagues and the community, the autonomy a respected teacher is granted by her headmaster, the specific problems she is assigned and the specific opportunities she is invited into, the professional-development travel that takes her outside the school for training or exchange, and the assurance that her work will not be undone by the next change of administrative leadership. A reform that touches only the formal compensation structure, while leaving the informal incentive architecture unchanged, is a reform that will produce substantially less teacher engagement than its designers expected. A reform that deliberately reshapes the informal incentives — recognising certified teachers publicly at the school and district level, granting them specific autonomy over classroom choices, inviting them into the mission's own advisory structures, protecting their work from arbitrary administrative revision — is a reform that engages the profession at the level where its behaviour is actually shaped. The formal compensation increment I proposed earlier is the visible part of the incentive reform. The informal architecture is the less visible, and in some ways more consequential, part.



VI. Protections the reform must build in

I want to name, explicitly, a set of protections the reform must provide to the teaching profession, because without them the teacher-capacity agenda will run aground on professional resistance that is entirely rational given the risks teachers face.

The first protection is against deskilling. A common pattern in educational-technology deployments globally has been the use of the technology as a justification for reducing the status, autonomy, and eventually the compensation of teachers. The argument, rarely stated openly but repeatedly enacted, is that if AI can deliver much of what teachers used to deliver, then teachers require less training, less pay, and less job security. This argument is wrong on its merits — what AI delivers well is the routine practice component of teaching, not the pedagogical judgement, the relational work, or the subject expertise that good teaching also requires — but it is a persistent argument that will be made repeatedly as the reform progresses. The reform must pre-emptively reject it. Specifically, the reform's foundational documents should state, and the Mission's statutes should codify, that AI tools in classrooms are complements to teachers and not substitutes, that the reform's deployment will not be used to reduce teacher positions or pay, and that any subsequent attempt to use the reform to justify workforce reductions will constitute a breach of the reform's foundational commitments.

The second protection is against surveillance. AI tools in classrooms generate substantial data not only about students but about teachers — what pedagogical choices they make, how effectively they use the tools, how their classes perform on AI-administered assessments, how much they deviate from prescribed lesson plans. This data, in the hands of a school administrator, a district officer, or a state government with particular priorities, becomes a powerful instrument of control over teachers' professional autonomy. A reform that allows classroom-interaction data to be used for teacher evaluation, without careful protection,

will produce a teaching workforce that teaches defensively – optimising for the metrics rather than for the children – and that loses precisely the professional judgement the reform depends on. The protection I recommend is categorical: no classroom-interaction data generated by AI tools in a teacher's classroom may be used for that teacher's evaluation, without the teacher's explicit, informed, uncoerced, individual consent, which may be withdrawn at any time.

The third protection is against algorithmic management. Related to the surveillance concern but distinct from it is the risk that AI-generated recommendations – about which students need more attention, about which pedagogical approaches work best, about which teachers are underperforming – come to dominate decisions that should remain in human judgement. The algorithm's recommendations are often useful. They are not infallible, they are not neutral, and their adoption without careful human review produces outcomes the algorithm's designers did not anticipate and the teachers subject to them did not consent to. The protection is that decisions affecting a teacher's career, pay, placement, or evaluation must always be made by human decision-makers, with algorithmic inputs treated as advice rather than as direction.

The fourth protection is against rushed implementation. Pilot programmes and first-wave deployments of the reform will, inevitably, encounter problems – software that does not work as promised, training that was inadequate, content that proves inappropriate, infrastructure that was not ready. The teachers caught in these problems will bear the operational costs – the frustrated children, the angry parents, the administrative reprimands for performance against metrics they could not meet under the actual conditions. The reform must explicitly protect teachers against censure for problems attributable to reform-level failures rather than to individual teaching deficiencies. A specific mechanism – an ombudsman function within the Mission, with the authority to review cases in which teachers are being penalised

for deployment failures, and to protect the teachers when their complaints are justified — should be established from the outset. Without this mechanism, the first wave of teachers bearing the reform's problems will become the reform's most effective opponents, and the second wave will be harder to mobilise.



VII. The B.Ed. reform

I said, in Section I, that the pre-service teacher education system must be reformed alongside the in-service capacity development. I want now to specify what that pre-service reform looks like, because it is a component of the teacher agenda that receives considerably less attention than the in-service component and is, in the medium run, arguably more consequential.

The B.Ed. programme — the two-year Bachelor of Education programme that is the standard pre-service preparation for secondary teachers in India, with variations at different levels — is regulated by the National Council for Teacher Education. The NCTE sets the norms that B.Ed. programmes must meet, recognises new programmes, and oversees the existing ones. There are, at present, approximately sixteen thousand institutions in India offering teacher-education programmes, of wildly varying quality. A substantial number are, by any serious evaluation, poor-quality institutions whose graduates require substantial additional training before they are ready for classroom work. The NCTE's own reviews, and the various commissions that have examined teacher education over the past two decades, have consistently noted this quality problem without fundamentally resolving it.

The reform this book proposes requires, at minimum, that the B.Ed. curriculum be updated to include AI education as a core competence for all graduates from the next cohort onwards. The update involves specific changes. A dedicated course on AI in education, covering the pedagogical, technical, and ethical

dimensions at the depth appropriate to a pre-service teacher. Integration of AI-tool-use exercises across the pedagogical methods courses for specific subjects — mathematics, languages, sciences, social sciences — so that each subject-specialised graduate enters her first classroom having practised the pedagogical integration she will need. Training on the social and ethical discussion-facilitation skills that the fourth transversal theme requires. Practical classroom experience with AI tools during the B.Ed. programme's internship, in schools where the tools are actually deployed, supervised by teachers who are themselves capable of mentoring the pre-service teacher on this dimension.

The curriculum update is the easier part. The harder part is the faculty of the B.Ed. institutions themselves. A B.Ed. faculty member whose own preparation predates the AI-education era, and who is now asked to teach pre-service teachers about AI in education, is not in a position to do so without substantial preparation of her own. The NCTE must, therefore, run a parallel faculty-development programme for the B.Ed. institutions, as a precondition for the curriculum update being meaningful. This programme is substantial — approximately fifty thousand B.Ed. faculty across the country, of varying backgrounds and capacities — and its cost should be estimated and budgeted separately from the in-service programme. I estimate it at approximately two thousand crore rupees over four years.

The B.Ed. reform's payoff is delayed. The first cohort that graduates under the reformed curriculum enters classrooms four to six years after the reform is initiated. The cohorts that follow continue to enter over the subsequent decades. The reform's visible impact on the classroom workforce is, therefore, a decade out — which is why it does not receive the political attention it deserves. But a reform that is serious about the long-term teacher capacity of Indian education cannot skip this component. The in-service programme alone will produce a workforce that ages out and is

replaced by unprepared newcomers if the pre-service pipeline is not fixed in parallel.

There is, additionally, a category of B.Ed. institutions that the reform's quality agenda must confront more seriously than the current regulatory framework does. Several thousand of the sixteen thousand institutions are, by any honest evaluation, institutions that exist primarily as credential-issuance mechanisms rather than as genuine teacher-education establishments. Their faculty are minimal or absent, their infrastructure is nominal, their students rarely attend classes, and their graduates enter the profession with the certificate but without the training the certificate is supposed to represent. The NCTE has, at intervals, moved against some of these institutions, and some have been de-recognised. But the pattern persists, because the institutions serve a specific economic demand — for the teaching credential as a qualification for government employment — that has not been met by the genuine institutions in the numbers the demand requires. The reform's B.Ed. agenda should include a specific commitment to the progressive elimination of the weakest institutions, with a time-bounded programme that identifies them, offers them a structured pathway to remediation with a specified quality threshold, and, if the threshold is not met within the specified period, de-recognises them formally. The de-recognition will displace students and faculty, and will produce short-term costs. The costs are worth bearing, because the alternative is the continued production of nominally certified teachers who then require substantially more intensive in-service remediation than the reform's budget can sustain. The quality of the pre-service pipeline is, in the medium term, the most consequential leverage point on the overall quality of the teaching workforce, and a reform that does not act decisively on the institutional quality problem is a reform that has chosen not to exercise its own leverage.



VIII. What trust would look like

I want to close this chapter with a passage I have been delaying, because it is the passage the chapter has been building toward and because I want it to land with the weight the argument deserves.

The reform this book proposes asks a great deal of India's teachers. It asks them to learn new tools, new content, and new pedagogies. It asks them to engage with a field that is moving faster than any educational reform they have previously encountered. It asks them to hold, in their classrooms, the difficult conversations with children about what this technology is and what it means for them. It asks them to do this, for the most part, with inadequate infrastructure, inadequate support, and compensation that remains below what the profession's importance would justify. It asks them to do this in a policy environment where they are more often the objects of reform than the authors of it, and in an administrative culture that has, for decades, treated them as implementers of others' decisions rather than as professionals whose judgement is trusted.

The reform will not succeed if it continues this treatment. It will not succeed if it produces, at its end, a teaching workforce that is technically updated but professionally demoralised, that delivers the curriculum but does not own it, that tolerates the reform but does not believe in it. The reform's technical components — the curriculum, the content, the devices, the platforms — can all be put in place without the teachers' genuine engagement. They will, in that case, produce a reform that looks successful in the Ministry's progress reports and that fails in the classroom.

What trust would look like is not difficult to specify. It looks like teachers being consulted, seriously, on the curriculum before it is finalised, and having their feedback visibly incorporated or visibly rebutted with reasons. It looks like training programmes that treat them as professionals whose time is valuable, rather than

as recipients of instructions. It looks like compensation that acknowledges the additional work the reform demands of them. It looks like career progression that rewards investment in their practice. It looks like classrooms where their judgement is respected, their autonomy protected, their mistakes treated as learning, and their successes recognised. It looks like a policy environment in which they can speak, without fear, about what the reform is getting right and what it is getting wrong.

None of this is novel. All of it is what any professional workforce, in any country, requires in order to deliver at its best. The specificity is only that Indian teaching has, for most of its recent history, received less of this than the profession's consequence has merited, and that the AI-education reform — because of its scale, its ambition, and its demands — is the specific reform at which the pattern must change.

I want to say one further thing, which is a personal observation offered as such. In the classrooms I have spent time in, across the states I have visited in preparation for this book, I have met teachers whose commitment, in circumstances I would find disabling, has left me, repeatedly, shaken. Kusum Devi in Semra, with sixty-eight children in one room and a leaking roof. Renu Balakrishnan in Kozhikode, with her two colleagues whose trust she values more than the principal's. Moirangthem Brojeswor in Imphal, returning to a school in a state that is not yet at peace. Each of them, and tens of thousands like them across the country, is the actual agent through whom this reform will succeed. The Ministry will not educate the next generation of Indian children. The states will not educate them. The Mission I propose in Chapter Ten will not educate them. The vendors whose platforms and tutors we have been discussing will not educate them. These teachers will. If we ask them to do this, we owe them — as the minimum — the respect that honest professional partnership implies.

I do not know, as I write this, whether the reform will actually treat them this way. I know that it must. I know that every specific

recommendation in this book, from the Mission's statutory design to the Repository's commissioning process to the platform's DPI architecture, is either in service of this requirement or peripheral to the reform's real purpose. I know that a reader who takes one argument from this book should take this one. If you are a principal reading this, build your institution around the trust of your teachers. If you are a policymaker, build your reform around the trust of the workforce you are asking to deliver it. If you are a parent, understand that the single most important fact about your child's education is the teacher in whose classroom she sits, and honour her accordingly when you next meet her at a parent-teacher meeting.

The rest of the reform is implementation detail. This is the heart of it.

I will not, by the way, apologise for the length of this chapter or for the passages at its close that may read as more insistent than the book's ordinary register. The teacher chapter is insistent because the book's argument depends on its being so. A reform that asks nine and a half million people to undertake substantial new work cannot be designed, analysed, or advocated for without keeping those nine and a half million at the centre of the design, and a book about the reform cannot do its work without placing them at its centre too. I have placed them there as directly as I could. The subsequent chapters return to the institutional machinery — the head's playbook, the mission, the money, the delivery chain — and the reader may occasionally experience the return as a cooling of the book's register after the heat of Chapter Six. I hope the reader carries some of the heat forward, and reads the subsequent chapters with the teachers still in view, because everything that follows is, ultimately, in service of the partnership this chapter has tried to describe.



Three decisions for the reader

1. If you are an institution head, ask your teachers what they need, and listen.

Within the next month, hold a conversation with your teaching staff about what AI education in your institution should look like, and what they would need from you in order to deliver it. Do not arrive at the conversation with a plan. Arrive with questions. Listen more than you speak. Take notes. Bring the notes back in the subsequent meeting and report what you heard, with specific attribution where the teachers are comfortable, and say what you will do about it. The conversation is, in itself, the beginning of the professional partnership the reform requires at your institutional level.

2. If you are a policymaker, make the teacher the explicit centre of your reform's design.

When the Cabinet note that establishes the National AI Education Mission is drafted, ensure that its opening paragraphs name the teacher as the reform's central actor, and that its subsequent sections demonstrate the structural commitment this naming implies — through the Institute, through the compensation framework, through the protection provisions, through the teachers' council in the Mission's governance. A Cabinet note that lists teachers as one of many stakeholders is a Cabinet note that will produce a reform with the teacher at its periphery.

3. If you are a teacher reading this, organise.

The reform that this book proposes will be shaped, in important ways, by what the teaching profession demands of it. The profession's voice, in national reform, has historically been channelled through teacher unions and associations whose engagement has been strong on compensation and working conditions and, for understandable reasons, less developed on the

pedagogical design of reforms. The AI-education reform invites a specific form of professional engagement — at the level of curriculum, of training, of classroom practice, of protection provisions — that the profession's existing organisations are equipped to take on if they choose. Join the conversation through your union, your association, or a new professional body if the existing ones do not serve.



END OF CHAPTER SIX

CHAPTER SEVEN

The Equity Question



The lens through which every other argument in this book must be held.



Equity is not a separate chapter of the reform. Equity is the whole reform, viewed from the perspective of the children furthest from the current national conversation. I said this in Chapter One and I repeat it here because the temptation, in a book that has multiple chapters and a structured argument, is to treat equity as one chapter among many — a chapter that addresses the equity dimension and then allows the other chapters to proceed as if equity were not their concern. This treatment is a mistake. Every chapter of this book — on curriculum, on content, on technology, on teachers, on the machinery of the state — must be read, and in the final reckoning judged, against its consequences for the Indian child who starts with the least and whom the reform most urgently claims to serve.

The structural compromise of placing equity in a single chapter, while insisting throughout the book that equity is not a single-chapter matter, is one I have struggled with in the writing. An earlier draft threaded the equity arguments through every chapter, without a chapter dedicated to them. Readers of that draft told me, consistently, that the threaded approach diluted the equity claim — the specific readers most likely to care about equity missed the cumulative argument because it was never stated in one place, and the specific readers least likely to care about it found the threaded references easy to read past. The current structure, with a dedicated chapter, makes the claim visible and allows the reader to

encounter it as a specific argument. Its cost is that the reader who reads only the dedicated chapter, and not the earlier chapters, misses the extent to which equity shapes the recommendations those chapters contain. I have accepted the cost because it is, on balance, less harmful than the cost of the alternative. The reader who takes one argument from this book, I hope, takes the equity argument, and the structural concession of placing it in a chapter of its own is meant to increase rather than reduce its legibility.

The argument of this short chapter is that AI education in India can go in one of two directions, and that the direction is not yet decided. In one direction, the reform becomes the single most powerful instrument for equalising educational opportunity in Indian history – the instrument that finally gives the child in Semra a functional tutor that her parents could never afford, content in her own language that honours her context, and entry into a cognitive preparation that was previously available only to the children of the urban middle class. In the other direction, the reform becomes the single most powerful instrument for aggravating the inequality that Indian education has produced and perpetuated – the instrument through which the children of the already-served receive sophisticated AI tools that compound their existing advantages, while the children of the un-served receive degraded versions, broken infrastructure, and the symbolic rather than the substantive benefits of the reform. Both outcomes are possible. The determination between them rests on the specific design choices the reform makes – on who receives what, when, at what quality – and on whether the equity question is held as a live concern throughout the reform's design and implementation or treated as a problem to be addressed after the reform has otherwise been designed.

I want to note, before I enumerate the inequities, that the bifurcation I describe in this paragraph is not a matter of technical inevitability but of specific institutional choices, and the choices are being made, in the national policy conversation, at the moment this

book is being written. The choices include: whether the Repository's commissioning schedule privileges the languages of weaker state systems or defers them; whether the device procurement specifications reflect the infrastructure of stronger schools or the infrastructure of weaker ones; whether the teacher-training programme's default pace assumes the capacity of Kerala's SCERT or of Jharkhand's; whether the Equity Fund is ring-fenced or permitted to be reallocated to competing priorities when they arise. Each of these choices, taken in isolation, may seem technical. In aggregate, they determine which of the two directions the reform will go in. A reader who wishes to influence the direction should attend specifically to these choices as they emerge over the next two to three years, rather than waiting for the reform's aggregate outcomes to reveal, five years later, which direction was in fact taken.

Let me enumerate, as specifically as I can, the inequities the reform must address.

The rural-urban gap is the most visible. Indian rural schools receive, on almost every dimension — infrastructure, teacher availability, learning outcomes, administrative attention — less than urban schools do. The gap is entrenched, it is substantial, and it compounds over the years of a child's schooling. A reform that deploys AI education evenly across rural and urban schools will, in practice, deploy it unevenly, because the infrastructure preconditions — electrical, network, administrative — differ so substantially that the same nominal deployment produces different realities. The reform must either invest disproportionately in the rural preconditions to achieve equal realities, or accept that equal deployments will produce unequal results. The first option is expensive. The second is what will happen by default unless specific design choices are made otherwise.

Let me make the rural-urban asymmetry more concrete, because abstraction conceals it. Consider two schools. The first is a

government primary school in a village in Balrampur district of Uttar Pradesh — the kind of school Kusum Devi teaches in. It has four classrooms for five grades, one of them usable only when it does not rain. It has electrical wiring to three of the classrooms, of which the connection is active for perhaps four hours of the school day on average. It has no broadband connection; the headmaster's mobile phone, on a 4G tariff, is the school's entire digital infrastructure. It has 192 children on the rolls, of whom approximately 130 attend on a typical day. It has five permanent teachers and two para-teachers on honorarium. The nearest secondary school is seven kilometres away, reachable by an unpaved road that becomes impassable during monsoon. The second school is a private school in a middle-class neighbourhood of Gurugram. It has a dedicated computer laboratory with thirty-two machines, a fibre-optic broadband connection, air-conditioned classrooms with projectors, a library with ten thousand volumes, and a staff of sixty including eight who hold advanced degrees in their teaching subjects. Its fees are approximately eighteen thousand rupees a month per child. A reform that deploys, into both schools, the same AI tutor and the same curriculum and the same teacher-training programme, and then reports its success against an average metric that sums across both, is a reform that has misrepresented its own outcomes. The first school, without the infrastructure investment that precedes the deployment, will barely register the reform's arrival. The second, which already had most of what the reform provides, will gain marginally. The gap between them, which was enormous to begin with, will not narrow — in some dimensions it will widen, because the marginal benefit to the second school is immediate and the marginal benefit to the first school requires a sequence of infrastructure investments that may take three or four years to complete. The reform's honest arithmetic must recognise this asymmetry from the start.

The government-private gap is the second structural inequity. Private schools in India — the recognised private sector as distinct from the unrecognised — educate approximately thirty percent of

Indian children, concentrated in urban and semi-urban contexts, and serve, in their best institutions, children from families with the resources and aspirations to pursue high-quality education aggressively. A reform that leaves the government-private deployment pace to market forces will see private schools adopt AI education quickly, with well-trained teachers and current content, while government schools proceed slowly. The gap widens during the years of asynchronous adoption. A child who enters Class 3 in 2026 in a well-resourced private school and a child who enters Class 3 in 2026 in a government school in a weak-SCERT state may, by the time they reach Class 8, have had educational experiences so different that their life trajectories have diverged in ways they did not before the reform began. The reform's equity question includes the question of whether this divergence is managed or accepted.

I want to say a word about the specific regulatory instruments through which the government-private question can be addressed, because the rhetoric of equity in this domain is often unaccompanied by the specific mechanisms that would make it real. The Right to Education Act's Section 12(1)(c), which reserves twenty-five percent of seats in private unaided schools for children from economically weaker sections, is the one substantive instrument currently on the books, and its implementation has been uneven across states and contested in the courts at various junctures. The AI-education reform presents an opportunity to strengthen the instrument, not by expanding its reach but by ensuring that the children who enter private schools under the provision receive, within those schools, the full AI-education experience that the private-fee-paying children receive. The commitment should be codified in the Children's AI Safety Code I propose in Chapter Sixteen: every child in every school, regardless of the fee the child's family pays, receives the same access to the AI tutor, the same training from teachers certified to the same standard, the same opportunity to take the secondary-stage electives. This is a modest provision. It is also the kind of provision

that, absent specific codification, is frequently not honoured in practice, with the EWS-quota children in private schools receiving, in some institutions, a visibly differentiated experience from their fee-paying classmates. Equity at the level of specific provision, within the institutions that already exist, is a form of equity that the reform can actually deliver, and that it should not neglect in favour of the broader and more politically contested questions of inter-institutional equity.

The linguistic gap is the third. I have addressed this in Chapter Four. The child whose mother tongue is one of the less-served Indian languages will, unless the reform specifically ensures otherwise, receive content and tools of lower quality than the child whose mother tongue is English or Hindi. The inequity is particularly acute for children whose mother tongues are not among the twenty-two scheduled languages — the tribal languages of central and north-eastern India, the regional languages that have not received state sponsorship, the minority languages within states whose official language is different. The reform cannot serve every Indian language to the same depth in the first decade. It can, and must, ensure that the scheduled languages are served equitably, and that the linguistic inequity between them and English is progressively narrowed.

The specific mechanism through which the narrowing happens is, in my view, the Repository's commissioning schedule. If the Repository commissions its English and Hindi content in the first eighteen months and its content in Odia and Assamese and Konkani in the fourth and fifth years, the three-year delay is itself a specific harm done to the children whose mother tongue is the later language. They encounter the reform, in its early years, through inferior translated material, and by the time the properly commissioned content in their language arrives, their schooling in its foundational years has already been shaped by the second-class provision. The Repository's commissioning schedule should therefore be simultaneous across the major language groups from

the first year, even if this means a slower overall pace than a sequential approach would produce. Slower but equitable is the right trade-off. The slower pace is a price the reform pays for being the kind of reform it claims to be.

The gender gap is the fourth. Indian education has, over the past two decades, made substantial progress in closing the gender gap in enrolment, particularly at the primary stage. The gaps that remain are in the later stages of schooling — where adolescent girls are disproportionately likely to leave school for reasons ranging from marriage to safety to family economic pressure — and in the subject and stream choices that shape post-school trajectories. A reform that deploys AI education evenly across genders at the foundational stage, but that produces Applied AI electives at the secondary stage that are disproportionately taken by boys, has replicated the existing gender-and-technology pattern of Indian higher education and the Indian workforce. The reform must, as a specific design concern, ensure that its secondary-stage electives are genuinely accessible to girls — through facilities, through scheduling, through the messaging that shapes how girls and their families understand what the electives are for — and must measure and report the gender composition of elective enrolment as one of its public accountability metrics.

I want to be specific about one sub-dimension of the gender question that the broader framing often obscures, which is the question of who the infrastructure reaches first in households where it does not reach everyone. When a rural household receives a mobile phone under a state scheme, or when a middle-class family acquires a tablet that is nominally for the children's educational use, the specific allocation of access within the household is not always equal between the boys and the girls of the family. Cultural patterns and household hierarchies produce, in some proportion of cases, a de facto allocation that gives the boys more hours, more privacy, and more trust in their use of the device than the girls receive. The reform cannot address intra-household

allocation directly — this is beyond its remit and beyond the effective reach of any schooling policy — but it can reduce its salience through specific design choices. School-based device provision, rather than household-based, reduces the intra-household allocation problem because the child's access to the tool is secured within the school's own hours and under the school's supervision. Teacher training that includes specific attention to the gendered patterns of AI-tool use within classrooms — to noticing when girls are being crowded out of shared devices, when their contributions in AI-related classroom discussions are being undervalued — is the specific classroom practice that protects the gender agenda within the reform's own operational space. The design attention is not costless and is not automatic. It requires explicit inclusion in the training programme, in the classroom-observation rubrics, in the specific metrics the mission reports on. Inclusion by design is the specific mechanism through which the gender commitment survives the ordinary institutional pressures that dilute it.

The gender dimension of the AI-education reform has a further specificity that deserves attention. The global pattern of technology employment is sharply gendered, with women substantially underrepresented in the technical fields that AI education most directly prepares students for, and the Indian pattern is no different — arguably more pronounced than the global average. The reform will either replicate this pattern at school level or disrupt it. Disruption requires, among other things, that the role models presented in the Applied AI and Data and Design electives be substantially female; that the teacher cadre delivering these electives be consciously recruited with gender balance as a specific goal; that the career guidance presented to students at the elective-choice stage actively counter the cultural defaults that direct girls away from technical tracks; and that the classroom experience itself be designed to be hospitable to girls in mixed-gender settings, which is a specific pedagogical challenge the current Indian classroom often does not handle well. The

design work required here is substantial. It is also work for which there is genuine international experience to draw on — the gender-balance interventions in computer-science education that have, in countries like Estonia and Rwanda, produced measurable shifts in elective uptake over a decade-long horizon. The Indian context is different, and the interventions will need to be adapted to it, but the underlying design principles translate. A reform that does not make this design investment, and that simply hopes the existing gender patterns will change on their own, is a reform whose gender equity commitment is rhetorical rather than substantive.

The caste and community dimensions of Indian educational inequity are the fifth. Dalit, Adivasi, and Muslim children, on average, complete fewer years of schooling than children from other communities, perform less well on standardised assessments, and are less likely to continue into higher education. The causes are multiple — economic, geographic, social, and sometimes discriminatory. A reform that deploys AI education across Indian classrooms without specific attention to these community patterns will, in its aggregate outcomes, reinforce them. The specific attention required includes content that is representative rather than exclusionary in its examples, teacher training that addresses the specific forms of discrimination that affect these communities' children, and assessment frameworks that do not systematically under-recognise the achievements of children whose home contexts are different from the dominant-community child the framework implicitly assumes.

I want to be specific about what community representation in content looks like in practice, because the principle is often stated without the details that distinguish genuine representation from its cosmetic imitation. Genuine representation means, in the content the child encounters at the foundational stage, stories whose protagonists come from a range of communities and castes, whose settings include the village and the mohalla and the forest as well as the city, whose economic circumstances include the farmer and

the artisan and the wage-labourer as well as the salaried professional. It means, in the middle-stage AI-literacy examples, case studies that address the specific concerns of marginalised communities — the use of machine learning in agricultural advisories for smallholder farmers, the use of AI in identifying and responding to caste-based discrimination online, the use of voice technologies for communities whose literacy rates are low. It means, in the secondary-stage AI-and-Society elective, readings that engage Dalit, Adivasi, and Muslim perspectives on the questions the elective addresses, not as tokens but as central intellectual contributions. The commissioning work required for this representation is specific and substantial. It cannot be delegated to the default content-production networks, which are, in their own composition, often insufficiently diverse to produce the representation naturally. The Repository's commissioning process should explicitly recruit authors and editors from the communities the content is meant to represent, should audit its aggregate commissions for representational balance, and should publish the audit results. Representation as audited commitment is what survives the inevitable pressures toward the default; representation as aspiration does not.

The disability dimension is the sixth. Indian education, despite legislative commitments over the past three decades, remains largely inaccessible to children with significant physical, sensory, or cognitive disabilities. AI tools have, in principle, the capacity to transform accessibility — through speech-to-text that serves children with visual impairments, through text-to-speech for children with reading difficulties, through simplified interfaces for children with cognitive disabilities, through interpretation for children who communicate in sign languages. The capacity is there. Realising it requires specific design investment — accessible interfaces, universal design principles, specialised content for different forms of disability — that will not happen by default in a market-driven deployment. A reform that leaves accessibility to the vendors who make the tools will produce accessibility only to

the extent that the vendors' commercial priorities reach, which in most cases is insufficient. The reform must require accessibility as a core standard, and must invest directly in the components — specialised content, assistive-tool integrations — that are non-commercial but essential.

There is a specific sub-population within the disability category that deserves separate attention, because it is the sub-population for which the reform's potential benefit is largest and the current market provision is smallest. These are children with cognitive and learning disabilities — children with dyslexia, children on the autism spectrum, children with ADHD, children with intellectual disabilities at various levels. Their total number in Indian classrooms, if the usual epidemiological estimates apply, runs to between ten and fifteen million. Most of them receive, in the current system, no specifically tailored educational provision. Many are mis-identified as behaviourally difficult or as academically weak, and are passed through the system without diagnosis or accommodation. The AI tools the reform deploys, if they are designed with these children specifically in mind, can produce the kind of individualised provision — patient, paced to the child's own trajectory, responsive to the specific pattern of her difficulties — that the current classroom cannot produce for her. The investment required is specific: clinical expertise built into the content design, teacher training on the recognition of learning differences, assessment frameworks that can distinguish cognitive disability from academic under-preparation, and the integration of the education system with the paediatric and clinical psychology services that can support diagnosis where it is needed. The investment is modest in fiscal terms but substantial in institutional complexity, because it requires coordination across the education and health systems, which do not currently coordinate well. The Mission should include, within its Evaluation and Equity Directorate, a specific team addressing cognitive and learning disabilities, and the reform's commitments should include a concrete target — that within five years, no Indian child with an

identifiable learning disability passes through the school system without the specific accommodation her condition requires. The target is ambitious. The technology, finally, makes it achievable. A reform that does not pursue it has defined, implicitly, which children it is prepared to leave behind.

These inequities are not independent. They intersect. A girl with a disability from a Dalit family in a Scheduled Tribe-concentrated district in eastern Madhya Pradesh, studying in a government school with chronic teacher vacancies and intermittent electricity, in a language that is regionally dominant but not officially scheduled, is a child whom the reform must specifically, and centrally, consider as it is designed. The list of her disadvantages is not a rhetorical flourish. It is the specification against which the reform's equity commitment is tested. A reform that claims to serve her must be able to say, concretely, what it does for her — what infrastructure reaches her school, what teacher capacity is built in her district, what content is produced in her language, what accessibility is designed into the tools she will use, what specific barriers are dismantled so that her trajectory is shaped by her capabilities rather than by her circumstances.

I want to name this child, because naming is a form of discipline. Let me call her Lakshmi. She is seven years old. She is in Class 1 at a small government primary school in Umaria district, eastern Madhya Pradesh. Her mother speaks Gondi at home. She has mild cerebral palsy that affects her fine motor skills, which the local health system has not diagnosed formally and which her school has interpreted as the child being slow. Her father is a daily-wage agricultural labourer who is unsure whether she will continue to school past Class 8 because the school is four kilometres from the village and the secondary school is thirteen kilometres. The nearest DIET is in the district headquarters, seventy kilometres away. The nearest Block Resource Centre coordinator visits her school twice a year, with about ninety minutes to spend. Her teacher has never used a computer. When, in 2030, the reform

claims that India has deployed AI education at national scale, the question is whether Lakshmi's specific schooling experience has, by that date, changed materially in her favour. If it has, the reform has done the work it claimed to do. If the reform has reached every other school in the country but has not reached Lakshmi's, it has not done the work, whatever its aggregate numbers may show. The specification of equity is not a rhetoric. It is the specific accumulation of specific provisions that, added together, change the specific experience of the specific child the reform was meant to serve.

I have proposed, in Chapter Eleven of this book, an Equity Fund – a dedicated fiscal line within the Mission's budget, of the order of thirty thousand crore rupees over five years, earmarked explicitly for the states, districts, and institutions that begin the reform from weaker baselines. The Fund is the specific fiscal mechanism through which the equity commitment becomes real rather than rhetorical. It is a visible, audited, politically accountable commitment that the reform will spend more where the need is greater.

The Fund's design principles deserve to be made specific, because the generic concept of an equity fund has been attached, in other Indian programmes, to arrangements that have sometimes been substantively different from what the name implies. The principles I would propose are five. First, the Fund's allocation to states, districts, and institutions must be formula-based, with the formula published and open to public scrutiny. The formula should include objective indicators of disadvantage – baseline learning outcomes, infrastructure availability, teacher density, linguistic minority status – weighted in ways that are defensible on the record. Discretionary allocations outside the formula should be rare and specifically justified. Second, the Fund's uses must be tracked at the level of the specific district or institution receiving the allocation, with published annual reports that show where the money went and what it produced. Third, the Fund's governance

should include specific representation from the communities the Fund is meant to serve — not as a token, but with voting authority over Fund decisions at a level proportionate to the Fund's aggregate commitment to each community. Fourth, the Fund's administrative costs should be capped at a low proportion — perhaps three percent — of its total envelope, to ensure that the substantive benefit reaches the children rather than the administrators. Fifth, the Fund should be subject to an independent external evaluation every two years, with the evaluation's findings published in full and with the Governing Council formally responding to each recommendation. The five principles together convert the Fund from a fiscal line with an equity label into an equity instrument whose performance can actually be assessed.

But fiscal mechanisms are necessary and not sufficient. The equity commitment is also a design commitment — a commitment that specific design choices at every layer of the reform are made with the most disadvantaged child in view. The curriculum that is age-appropriate for the urban middle-class child may not be age-appropriate for the rural child whose literacy baseline is lower. The training programme that works for the Kerala teacher may not work for the Jharkhand teacher. The AI tutor that performs well in Hindi may not perform at all in Santali. The equity commitment requires that these differences are designed for, not ignored and addressed as exceptions after the fact.

I want to close this chapter with an observation that is, in its way, both hopeful and difficult. The AI-education reform, done with genuine equity commitment, has the potential to be the single most powerful instrument for closing the Indian educational divide in our lifetimes. The technology is, at its core, an instrument of individualised attention at low marginal cost. The traditional barrier to educational equity — that quality individualised attention is expensive, and therefore rationed toward the children who can pay for it — is, for the first time in the history of mass schooling, potentially breachable. An AI tutor that reaches the

child in Semra is, in principle, the same tutor that reaches the child in Bengaluru's most expensive school.

The hope is real. Whether it is realised is, again, a matter of design choices. The same technology that could equalise can also, if deployed carelessly, further entrench the inequities it could have closed. The bifurcation between these two outcomes is not determined by the technology. It is determined by the specific commitments the reform's designers, funders, and implementers make — commitments to equity as a first-order constraint, to disaggregated measurement of outcomes by community and geography, to a willingness to invest disproportionately where disadvantage is deepest, and to a discipline of honesty when the reform is falling short.



Three decisions for the reader

1. Disaggregate every metric.

Whether you are running an institution, a district, a state, or the national reform, report every metric you track in disaggregated form — by gender, by community, by urban-rural location, by disability status, by linguistic group. Aggregate metrics conceal the inequities the reform must close. Disaggregated metrics expose them. The simple act of disaggregation is, itself, one of the most powerful equity disciplines available to you.

2. Name the child the reform must most urgently serve, and test every decision against her.

The child against whom your decisions should be tested is not the median child of your institution or your district. It is the child furthest from the current baseline — the child with the most compounding disadvantages. If your decision does not improve her education, your decision is inadequate regardless of what it

does for the median child. The test is severe. It is also, in the aggregate, the test that produces reforms that genuinely change the country rather than reforms that change the top half of the distribution and leave the bottom half where it was.

3. Allocate the Equity Fund, or its equivalent at your scale, transparently.

If you are a national or state policymaker, ensure that the Mission's Equity Fund is allocated through a transparent, publicly-reported process, with specific commitments to the states, districts, and populations that most need it. If you are an institution head, the equivalent at your scale is the commitment of your own institutional resources disproportionately toward the children in your school who are starting from the weakest baseline. The transparency of the allocation is the guarantor of its integrity.



END OF CHAPTER SEVEN

CHAPTER EIGHT

A Letter to Parents



The conversation the school cannot have for you.



This chapter is written differently from the others. It is written to you, directly, in the second person, because it is a chapter you may read on your own without the rest of the book, and because the matters it addresses are matters about which you are the decision-maker in your own home. The principal of your child's school can do many things. She cannot have the conversation with your child that you can have. She cannot model for your child the adult posture toward artificial intelligence that you, by what you do at home, will model whether you intend to or not. She cannot replace you, and she is not trying to. What she can do, she will do. The work this chapter is about is work that falls to you.

I will be direct with you, because I think that is what a parent reading a chapter of this kind wants. You have been hearing about artificial intelligence, constantly, for two or three years. You have heard that it will transform your child's education, her career prospects, her life. You have heard that if she does not learn it she will be left behind. You have also heard that it is dangerous, that it hallucinates, that it replaces human thought, that it will take jobs and reshape societies in ways we do not yet understand. Both sets of things are, in varying degrees, true. Neither is the whole truth. You are, reasonably, uncertain about what to do.

The uncertainty itself is worth acknowledging rather than being dismissed as something the remainder of this chapter will resolve for you. The uncertainty is justified. The people who speak

most confidently about what AI will do to your child's life — whether they are the evangelists who promise transformation or the alarmists who predict catastrophe — are, in most cases, speaking with more certainty than the available evidence warrants. The honest posture, at this moment, is a posture of thoughtful engagement with genuine uncertainty, which is the posture the best teachers, the best researchers, and the best parents I have met have adopted. I will not, in what follows, pretend to resolve the uncertainty. I will try to give you, within the bounds of what is actually known, the specific guidance that seems to me most likely to serve your child well. The guidance may prove wrong in specific ways as the technology develops. It is the best guidance I can offer from where I stand. You will, necessarily, refine it in light of what you observe in your own household and what you learn as the years pass.

This chapter tries to give you enough specificity to act with confidence, while being honest that no one — not the author of this book, not the designers of the tools, not the government, not your child's school — knows exactly how this will unfold over the next ten years. What I offer you is the best judgement of a practitioner who has spent years thinking about this and who, in his own life, is making decisions about the children he cares for that are informed by the judgement below. You may disagree with any of it. Disagreement is part of your work as a parent. But you are entitled to know what someone who has thought carefully about this would do in your position.

I want to add, because the situation is common and deserves specific acknowledgement, a word addressed to the parent who feels she is the less technologically fluent adult in her own household — who, when the conversation turns to AI, often defers to her spouse or to an older child, and who has absorbed the impression that she is not qualified to hold opinions on these matters. I want to say, directly: the impression is wrong. The judgement this chapter asks you to exercise is not technological

judgement. It is parental judgement — the judgement about what your child needs at her specific age, about what norms you want to establish in your home, about when to intervene and when to trust. This judgement is not a function of how comfortable you are with technology. It is a function of how well you know your child, how carefully you think about her development, and how seriously you take the parental role. A parent who has never opened an AI tutor in her life but who knows her daughter well is, on the questions this chapter is about, better equipped than a parent who uses AI tools fluently but whose attention to the child is secondary to her own professional preoccupations. The specific conversations this chapter asks you to have with your child are conversations you are fully qualified to have, and the fact that you may feel you are not is itself among the reasons you are reading this book.

Let me walk through what I would suggest, item by item, from the youngest child to the oldest, with the specific caveats that apply at each age.



I. If your child is under eight

If your child is under the age of eight, my suggestion is simple. She does not need AI tools. Not in any form. Not for educational purposes, not for entertainment, not as a novelty. The cognitive and developmental work of the years before eight is the work of language, of social interaction, of play, of the slow acquisition of literacy and numeracy through methods that have worked for a long time. An AI tutor or chatbot does not improve on this work, and the screen time it involves risks crowding out the activities that do the developmental work best. There are exceptions — children with specific learning needs for whom assistive AI tools are genuinely useful — but for the typical child, the answer at this age is that she does not need this now, and the time for introducing it will come.

What I would do, at this age, instead: read to her, in whichever language your household speaks most naturally. Let her see you reading, occasionally. Talk with her about what is around you — the birds, the weather, the people she meets. Let her play without constant adult direction. Keep screens, of any kind, modest in her day. If she goes to a pre-school or a primary school that is using AI tools before Class 3, ask the school, directly, why. The school may have a good answer. If it does not, you have raised a question that is worth raising.

What I would not do, at this age: buy any of the educational AI toys or apps that are marketed to parents of young children. Most of them do not help, some of them actively hurt, and almost all of them distract from the developmental work that matters at this stage. If a relative gives one as a gift, you are not required to use it. Your child's early years are short. Protect them.

I want to say a word, because the question is pressing for many parents of young children, about the voice-assistant devices — the Alexa, the Google Home, the Siri on the kitchen counter — that have become ambient in many Indian middle-class homes. These are AI products. Your young child interacts with them whether you have thought about it or not. The question is not whether she will encounter them, since she already has, but what norms you establish around her interaction. The norms I would suggest, at this age, are simple. The voice assistant is not her friend. It is not her teacher. It is a tool that belongs to the adults of the household, which she can use, under supervision, for specific purposes — asking what the weather is, playing a particular song, setting a timer. She does not have extended conversations with it. She is not left alone with it for long periods. When she asks it a question, an adult, occasionally, follows up with her on what it said and whether it was right. The norms are not because the devices are dangerous, in any strong sense; they are because the relationship between a young child and a machine that responds to her

questions is a relationship that is better shaped deliberately by you than left to develop on its own.



II. If your child is between eight and eleven

If your child is between eight and eleven — roughly, Class 3 through Class 5 — she is at the age where AI tools can begin to play a useful role, supervised and modest, in her education. What the role looks like depends substantially on her school's provision and on what you do at home, and both matter.

At school, if her school is using an AI tutor for foundational literacy or numeracy practice, and the tutor is well-chosen and well-supervised, this is a good thing. The AI tutor at this stage is, at its best, patient individualised practice — the kind of practice that a teacher managing a class of fifty cannot provide in the quantity the child needs, and that the child benefits from substantially. The best current AI tutors in Indian contexts, operating in the child's mother tongue, on the simple tasks of reading practice and arithmetic practice, are genuinely useful, and the evidence base for this is reasonably strong. If your child's school is using such a tutor and your child is enjoying her time with it, this is evidence that the school is doing something right, not evidence that the school is pushing her too hard or too fast.

If her school is not using an AI tutor, this is not a gap you need urgently to fill. You can, if you wish, introduce her to a tutor at home — there are free versions available in most major Indian languages now, and the paid versions are, for most middle-class families, affordable — but this is optional, not required. The child who does not have AI-tutor access at this stage is not at a significant disadvantage. Her more important work, at this age, remains foundational literacy, numeracy, curiosity about the world, and the development of sustained attention that screens of any kind tend to undermine.

What I would suggest you do, regardless of what the school is doing, is begin the conversations about AI with your child now, at this age. The conversations are short, and they are not lectures. When your child asks a question to a voice assistant — what is the weather in Delhi today — you can say, afterward: did you notice how it knew that? It knew because people made it and gave it a lot of information. But it can be wrong sometimes. Do you know how we could check? The conversation is thirty seconds long. It introduces the three things I named in Chapter Three as the foundational understandings — the machine as tool, the machine as data-shaped, the machine as fallible. Repeated, casually, over months and years, these conversations produce a child who has accurate intuitions about the technology before she has any need to formalise them.

A specific warning for this age: limit the total screen time your child has, from all sources, including the AI tools I am discussing. The evidence on the effects of screen time on young children is not fully settled, but the prudent parent restricts it, and this restriction does not conflict with the use of AI tools for short, focused, educational purposes. An hour of AI-tutor time on literacy practice is fine. Four hours of general screen time that includes some AI-tutor time alongside entertainment apps is not. The discipline, as with many things at this age, is yours to maintain.

A specific question that parents have asked me at this age range concerns the fine line between what the school provides and what you yourself provide at home. The pattern I have seen in many Indian middle-class households is a reflex toward supplementation — if the school is using an AI tutor for thirty minutes a week, the family adds another thirty minutes at home; if the school is doing coding in Scratch, the family enrolls the child in a parallel coding class. The instinct is well-intentioned. It is also, in most cases, counterproductive. A child of nine who has three hours a week of AI-tutor time across school and home is not, in general, progressing substantially faster than a child who has one hour, and

she is definitely losing hours that she could have spent on the slower, less measurable, more developmentally valuable activities of her age — the outdoor play, the conversation at the dinner table, the uninstructed reading. If you are confident that your child's school is doing the AI work of this stage competently, the supplementation you add at home should be modest and chosen with care. If you are not confident, the remedy is usually not more AI time but the specific conversation with the school about what it is doing and why.



III. If your child is between eleven and fourteen

If your child is between eleven and fourteen — roughly, Class 6 through Class 8 — she is at the age where AI becomes a subject she learns about, not only a tool she uses. This is also the age where the risks of unsupervised AI use begin to become significant, because she is capable of using the tools with sophistication she did not have at nine, but without the judgement she will have at sixteen.

At school, her exposure to AI should be broadening. She should, in a well-run school following the roadmap this book suggests, be having a dedicated period each week on computational thinking, data literacy, and AI literacy. She should be using AI tools across other subjects, with the teacher's supervision and guidance. She should be having, occasionally, classroom conversations about the social and ethical dimensions of AI. If she is in a school that is doing these things, she is in a good place. If she is in a school that is not, the gap is beginning to matter more than it did at eight, and you have some work to do at home to fill it.

At home, three specific things. First, let her use AI tools, supervised. The supervision is not constant — she is old enough to have some independent use — but the tools she uses should be

known to you, the accounts should be her own with your oversight, and the ground rules for use should have been discussed. The ground rules I would suggest: AI tools are for learning and exploration, not for doing her homework for her. If she uses a tool to help her draft something, the teacher should know, and the draft should be a starting point for her own work rather than its substitute. If she uses a tool to answer a factual question, she checks the answer — the AI is not a final source, it is a starting point. These rules are not rigid. They are the beginnings of a professional relationship with AI tools that she will refine throughout her life.

Second, have the social and ethical conversations with her, not only as they come up but proactively. At this age, she is forming views — about fairness, about privacy, about work, about the future — and the conversations about AI are an opportunity to engage with those views seriously. When a news story appears about AI and jobs, read it together and discuss it. When she tells you about something she encountered online, ask her what she thought the AI behind it was trying to do, and whether she agrees with what it was doing. When she comes home from school with a classroom discussion about machine bias, ask her what she thought, and share, in your own voice, what you think. These conversations do not require you to be an AI expert. They require you to be a thoughtful adult engaging seriously with a young person who is becoming capable of serious thought.

Third, protect her, actively, from the AI uses that are specifically designed to exploit young adolescents. These include certain social-media AI systems whose engagement-optimisation is harmful to adolescent mental health, AI-generated content that normalises patterns of relationship or self-presentation that are damaging, and, at the darker end, AI-facilitated predatory behaviour that targets children of this age. The protection is not through surveillance alone — though some supervision is appropriate at this age — but through conversation, through trust, through the establishment in your home of norms around the

technology's use that your child understands are for her protection rather than for her control.

A specific warning at this age: the temptation, for parents whose children are academically high-performing, is to add AI-based enrichment that pushes the child harder. AI tools for advanced mathematics. AI-based test-preparation. AI-intensive coding tutorials. Some of this, in modest doses, is fine. In large doses, it is counterproductive — the adolescent who is academically pushed too hard through AI-optimised preparation often burns out, loses intrinsic motivation, or develops a relationship with learning that is transactional rather than genuine. The goal of these years, academically, is not maximum acceleration; it is the development of a child who continues to love learning into her adulthood.



IV. If your child is between fourteen and eighteen

If your child is between fourteen and eighteen — Class 9 through Class 12 — she is now capable of using AI tools as a professional adult would use them. The parenting task shifts, accordingly, from direct supervision to the kind of consultative presence a young adult requires. She will make her own decisions, increasingly, about how she uses the technology. Your role is to be available as a sounding board, to offer perspective when it is invited, and to hold the larger questions — about her studies, her path after school, her well-being — in view when she is too immersed in the immediate to hold them herself.

At school, she is, if her school is following the roadmap, choosing among the three electives — AI and Society, Applied AI, Data and Design. Her choice of elective, or her decision to take none, says something about what she is interested in and where she wants to go. Engage with the choice, but do not make it for her. Her own judgement about what she finds compelling is, at this age,

more reliable than your judgement about what will serve her future. She may take the AI and Society elective and go on to study engineering; she may take Applied AI and go on to study literature. The electives at this stage are formative, not determinative. Trust her to choose and to learn from her choice.

Outside school, her use of AI tools is likely intensifying. She may be using them for serious academic work — research, writing, problem-solving — in ways that would have been implausible a few years ago. Some of this use is excellent. Some is concerning. The line between them depends on whether the AI is augmenting her capabilities or substituting for them. If she uses an AI tool to help her understand a difficult physics concept, this is augmentation. If she uses an AI tool to write her English essay for her, submitting the AI's work as her own, this is substitution. The distinction is not always obvious. It requires judgement. Help her develop the judgement by discussing specific cases with her, without lecturing. The question to ask is not is this cheating — that framing produces defensiveness — but what is the work of your mind that this AI is doing, and what is the work of your mind that you are doing, and are you comfortable with the balance.

I want to add a specific note on academic honesty at this age, because the AI-facilitated forms of dishonesty are different from the forms parents themselves grew up with, and the old frameworks do not always map well onto the new situations. A student who copies a friend's homework is doing something every parent recognises as dishonest. A student who uses an AI to help her understand a concept before she writes her own essay is doing something that is not obviously dishonest, and that many adult knowledge workers now routinely do in their own professional work. The line between the two, in the middle of that spectrum, is where most adolescents are spending most of their time. The conversations you have with your child should engage this spectrum rather than insisting on a bright line that does not exist. The question to ask is whether, at the end of her engagement with

the AI on a particular piece of work, she has learned what the assignment was meant to teach her. If she has, the AI has helped her learn, and this is what the AI is for. If she has not — if she has produced the output without the understanding — then she has, in a meaningful sense, skipped the work the assignment was designed to do, and this is the harm to guard against. The framing is not about what is allowed and what is forbidden. It is about what serves her education, which is the question she will be asking herself, about her own AI use, for the rest of her life.

The career conversations with a teenager in this age range are difficult, because the careers of the next decade are genuinely uncertain in ways that careers of the past were not. I would not, if I were advising a teenager today, tell her to pursue any specific career because of its AI-proof status. I am not confident I know which careers will be AI-proof, and neither is anyone else. What I would tell her — and what I do tell young people who ask me for advice — is that the most durable preparation for an uncertain future is not a specific technical skill but a combination of deep expertise in at least one field, good judgement, the capacity to learn continuously, and the interpersonal and ethical sensibilities that machines do not substitute for.

A specific question at this age is the question of AI and mental health. The adolescent years are, for many young people, the hardest years of their lives psychologically, and the interaction with AI tools — including certain AI companions and chatbot systems that are designed for emotional engagement — can be beneficial or harmful depending on the specific product, the specific use, and the specific child. I will not pretend to give you general guidance here, because the field is moving too quickly and the individual variation is too large. What I will say is that if your teenager is spending substantial time in conversation with an AI companion, you should know what she is doing, and you should assess — perhaps with the help of a professional — whether the

interaction is supporting her well-being or substituting for human connection in ways that are damaging.

I want to offer one further observation about AI and adolescent mental health that is, I am aware, more contestable than the advice in the preceding paragraphs. A subset of AI products currently available is specifically designed to create the experience of emotional intimacy with the machine – companion apps, roleplay-chatbot services, AI girlfriends and boyfriends, AI therapists that are not actually therapists. These products are engineered for engagement, and the engagement they produce, at the adolescent stage, can be intense. I have spoken with parents whose teenage children have, over several months, developed relationships with AI companions that occupied six, seven, eight hours of their daily waking time. In the specific cases I have knowledge of, the teenagers in question were already experiencing loneliness, depression, or social anxiety, and the AI companion became a substitute for the human connection their circumstances had not provided. The short-term effect, in these cases, was a reduction in the adolescent's immediate distress. The medium-term effect, across the cases I can assess, was a deepening of social withdrawal and an erosion of the skills of human connection that the adolescent would eventually need to rebuild. I am not suggesting that every adolescent using an AI companion is moving toward this outcome. I am suggesting that the risk is specific, that the products are designed in ways that intensify it, and that a parent who notices her teenager withdrawing into prolonged engagement with such a product should treat this as a signal to engage professionally with what is happening, rather than waiting to see whether it resolves on its own. This is, perhaps, the single most important piece of parenting advice in this chapter. I offer it as such.



V. What to ask the school

Regardless of your child's age, there is a set of questions you can ask her school that will tell you, in a short conversation, whether the school is thinking seriously about AI education or whether it is not. The questions are not trick questions. A school that is thinking well about this will have good answers. A school that is not will have vague answers, defensive answers, or no answers.

Ask: what is your school's policy on AI tools in the classroom, across each grade? The answer should be specific, written, and available for parents to read. If the school does not have a written policy, that is information. Ask: who at the school is responsible for AI education, and what is her background? A school that is serious about this has designated someone, and that person has some relevant preparation. Ask: what are the learning outcomes you are pursuing for my child at her current grade, and how are you measuring whether they are being achieved? The answer should be specific.

Ask: what data about my child are being collected by the AI tools the school uses, where is it stored, who has access to it, and what is it used for? The answer should be clear and reassuring. If the school does not know the answer, ask when they will find out. Data about children is a serious matter, and a school that has not thought carefully about it is a school that has outsourced the question to vendors whose incentives are not aligned with your child's privacy.

A question I have not seen asked often enough in these conversations, and that I would encourage you to add to your list, is whether the school has a specific plan for what it will do if the AI tool your child depends on becomes unavailable. Vendors fail. Products are withdrawn from the Indian market. Licences expire and are not renewed. The school your child attends has made a specific set of technology commitments that carry a specific set of

dependencies, and a school that has not thought seriously about what happens when a dependency fails is a school whose technology deployment is fragile in ways the parents have not been told about. The answer to the question does not need to be elaborate. It needs to be specific — we maintain, as a matter of policy, the pedagogical capacity to continue the curriculum without this tool for up to six weeks, during which time we would either transition to an approved alternative or, if no alternative is available, return to the pre-deployment approach. A school that can answer in roughly this form is a school that has done the contingency work. A school that cannot answer is a school whose children are, in a small but specific way, at the mercy of decisions made in offices they have no line of sight to.

Ask: how are the teachers in my child's grade prepared to use these tools? Who trained them, on what, for how long? A school that is using AI tools without having invested in teacher preparation is a school that is likely to produce disappointing outcomes regardless of the quality of the tools. Teacher preparation is the single most reliable indicator of whether a school's AI-education programme is serious.

Ask: what do you do if my child falls behind the intended outcomes? The answer should involve supplementary support, not simply more time with the AI tutor. An AI tutor is not a remediation strategy for a child who is struggling; it is a component of a remediation strategy that also includes teacher attention, peer support, and, in some cases, professional evaluation.

These questions are reasonable. You are entitled to ask them. A school that treats your questions as impertinent or as a sign that you are a difficult parent is revealing something about itself that is worth noticing. A school that welcomes the questions and engages with them seriously is a school that is likely to be a good partner in your child's education, on this issue and on others.



VI. *A closing note*

I want to say one last thing, directly, because it is the thing I most want you to take from this chapter.

The reform this book describes will happen, in one form or another. Your child's education will be shaped by artificial intelligence whether or not you or her school or the government is ready for it. The question is not whether she encounters this technology. She will. The question is whether the adults in her life — you, her teachers, her school, her government — have prepared her to encounter it with judgement, with agency, and with the habits of mind that will allow her to make this technology serve her rather than the other way around.

You are one of those adults. The specific contribution you make — the conversations you have, the norms you establish, the questions you ask, the examples you set — is not replaceable by any institution. The school will do what it can. The government will do what it can. You will do the rest. You do not need to be an expert in artificial intelligence to do this well. You need to be a thoughtful adult engaging seriously with a child who is becoming her own person, and you need to bring to the task the same judgement, patience, and love you have been bringing to every other aspect of her upbringing. If you are doing that, you are doing enough. If you are uncertain whether you are doing it, the uncertainty itself is evidence that you are taking the work seriously.

Your child is going to live a long life in a world that will continue to change. What you give her now — her foundational sense of who she is, her capacity to think, her relationships, her values — is what she will carry into that world. The artificial intelligence is a tool she will learn to use. The person who learns to use it is the person you are helping to raise.



Three decisions for the reader

1. Have the first conversation with your child this week.

Whatever her age, within the next seven days, have a single specific conversation with your child about artificial intelligence. Do not make it a formal event. Do not sit her down. Find the occasion — a question she asks, a news item you see together, a tool one of you uses — and take the thirty seconds to engage with it rather than letting it pass. The first conversation is the hardest. Every subsequent one is easier. The sooner you start, the more natural the conversations become, and the more they accumulate into the shared understanding you will need with her over the next decade.

2. Ask your child's school the questions in Section V.

Within the next month, at the next parent-teacher meeting or through a written communication to the principal, ask the questions about the school's AI-education policy, the designated person, the learning outcomes, the data, the teacher preparation, and the remediation strategy. Do not ask as a complaint. Ask as an interested parent. Note the answers. Compare notes with other parents you know at the school. A few parents asking these questions, in the same month, produces changes at the school that any one parent alone cannot.

3. Notice, honestly, the AI use in your own life, and what it is modelling.

Your child is watching how you, the adult in her home, use the technology. She is watching whether you check what the AI tells you or accept it uncritically. She is watching whether you use it for tasks that benefit from it or to avoid doing thinking you should do yourself. She is watching whether your phone is present in conversations or whether you put it aside. The model you set, every day, shapes her relationship with these tools more than any

conversation you have with her. This week, pay honest attention to what your own use of AI tools is modelling for her, and consider whether the model is what you would wish her to develop.



END OF CHAPTER EIGHT

PART III



Making it Real at Scale

CHAPTER NINE

The Institution Head's Playbook



A hundred-day plan, a two-year path, and the specific discipline of leading this change.



This chapter is written for the primary reader of this book — the institution head, the principal, the dean, the VC, the head of a training institute — who has, by this point in the book, formed her own view of what the reform requires and is ready to translate that view into specific action in her own institution. Everything that has come before has been argument, analysis, and specification. This chapter is operational. It assumes you have authority over an institution, that you are serious about acting, and that you need a structured plan rather than further persuasion.

I want to acknowledge, at the outset, that the plan this chapter sets out is ambitious for an institution head whose attention is already claimed by a dozen competing priorities — the academic calendar, the examination results, the governing body's expectations, the parent community's concerns, the teacher grievances that require her personal intervention, the regulatory compliances that consume the administrative team's weeks. The plan is not a modest addition to an existing workload. It is a substantial new commitment that will, if taken seriously, reshape how you and your senior team allocate your attention across the first two years. I offer two honest observations about this. The first is that the institutions that will succeed with this reform are the institutions whose heads have decided, at the outset, that the reform is a priority comparable to the academic calendar rather

than a peripheral initiative subordinate to it. If you are not in a position to make that decision — because your governing body has other priorities, because your own bandwidth is fully committed elsewhere, because the institutional conditions are not right — the honest response is to defer the work rather than to attempt it in a form too diluted to succeed. The second observation is that the effort the plan requires is front-loaded. The hundred-day foundational work is the most demanding period. The subsequent work, once the foundations are in place, distributes more naturally across the senior team's ongoing responsibilities, and the demand on your personal attention decreases substantially after the first six months. Knowing that the peak is early, and planning for it accordingly, is among the specific disciplines that distinguish institutions that complete the work from institutions that abandon it midway.

I will organise the chapter around three horizons. The first is the first hundred days, during which you establish the foundations for serious action without committing yourself to specific implementations that later proved to be wrong. The second is the first year, during which you begin to pilot specific deployments and build the institutional capacity that scale will require. The third is the two-year path, at the end of which your institution is, in a specific sense, an AI-education institution — not because the reform is finished, but because the institutional transformations required to sustain it are in place.

The three-horizon framing is not original to this book, and I want to acknowledge the lineage because it will be recognised by institution heads who have worked on earlier reforms. The horizons-of-change language was developed in management consulting in the 1990s and has been adapted in various institutional settings since. My use of it here is the specific adaptation I have found most useful for educational institutions approaching substantial reform — the first horizon focused on foundations, the second on piloting, the third on

institutionalisation. Other adaptations are possible, and an institution head who has developed her own version of this framing through earlier change-management work should use what she is comfortable with rather than adopting my version just because it appears in this book. What matters is the discipline of horizon-based planning itself – the explicit recognition that different kinds of work belong in different time frames, and that the attempt to compress them all into a single push is among the most common failure modes of institutional reform.




Figure 9.1 | The institution head's three horizons – foundations, building, scale.



I. The first hundred days

The first hundred days are about foundations. No AI tool is deployed. No curriculum is finalised. No vendor is selected. What you do in these hundred days is establish the conditions under which the subsequent work will succeed, and this establishment is neither glamorous nor immediately visible, which is why most institutions skip it and then regret the absence later.

The first task, to be completed in the first two weeks, is to designate a senior lead for the AI-education initiative. This is not a teacher on top of her existing responsibilities. This is a member of your senior team — a vice-principal, a dean, an academic head — whose remit for the next eighteen months includes substantial responsibility for this initiative. The role is part-time if your institution is small; full-time if it is large. The person must have genuine authority, access to your desk, and enough seniority that when she speaks to the rest of the institution she is understood to be speaking with the institution's weight behind her. If you do not have such a person on staff, the first two weeks are the time to hire one or to promote one. A reform led by someone without authority is a reform that will produce meetings but not outcomes.

The second task, in weeks three and four, is to hold a structured listening exercise with your teaching staff. This is the conversation I named as a decision at the end of Chapter Six, and I repeat it here as a specific hundred-day task. You hold the conversation. You listen. You take notes. You report back, at the subsequent faculty meeting, on what you heard and what you plan to do about it. The listening is not symbolic. It surfaces the specific concerns, aspirations, and capacities of the workforce you will be asking to deliver the reform. It also establishes, at the beginning of the process, a precedent of consultation that will shape how the rest of the initiative is received.

The third task, in weeks four through six, is to audit your institution's current AI-education provision honestly. The audit covers content, teacher capacity, technology, student experience, data practices, and parent engagement. Each dimension receives a candid assessment. The audit is not for the Education Department or for publication. It is for you and your senior team, and its value is in its honesty. An audit that concludes you are doing reasonably well across all dimensions is, in most Indian institutional contexts, an audit that is flattering itself. An honest audit almost always

identifies substantial gaps, and the identification of the gaps is the beginning of the plan to close them.

The fourth task, in weeks six through ten, is to develop, with your senior lead and a small working group that includes teachers, parents, and a senior student representative where appropriate, an initial vision for what AI education in your institution will look like over the next three years. The vision is not a detailed plan. It is a short document — five to ten pages — that specifies the principles, the grade-level outcomes, the broad phasing, and the commitments the institution is making to its teachers, its students, and its parents. The vision is drafted, consulted on, revised, and then formally adopted by the institution's governing body. The formal adoption matters because it gives the initiative institutional legitimacy that no memorandum from the principal alone can provide.

The fifth task, in weeks ten through twelve, is to communicate the vision to all relevant constituencies. The teaching staff receives a full briefing and an invitation to ask questions. The parents receive a written communication and, ideally, a meeting. The students at the secondary stage receive an age-appropriate presentation. The governing body, the alumni, and the school's broader community receive the summary that suits their engagement. The communication is not a sales pitch. It is a serious statement of intent, acknowledging what the institution does not yet know, committing to what it will do, and inviting the community's participation in the work to come.

The sixth task, in weeks twelve through fourteen, is to identify one specific pilot that will begin in the next academic term. The pilot is modest. It involves one grade, one subject, one carefully chosen AI tool, and a small number of teachers who have volunteered to be part of the pilot. The pilot is heavily documented — baselines are established, methods are recorded, weekly reflections are captured. The pilot is intended to produce institutional learning, not to produce a generalisable result. Its primary function is to develop, within your institution, a shared

and specific understanding of what AI-education deployment actually involves.

I want to be specific about what the pilot's documentation should contain, because the value of the pilot depends almost entirely on the quality of the documentation, and the typical institution's documentation practice in initiatives of this kind is, in my observation, inadequate to the purpose. The documentation should include: a written baseline of the students' starting competence in the specific domain the pilot addresses, established through a diagnostic at the start rather than recalled from memory at the end; a weekly journal from each participating teacher, capturing what she did that week, what worked, what did not, what she changed the following week, and what she wants to discuss with her colleagues; a monthly snapshot of student work produced with AI support, with specific examples retained rather than summarised; a record of every technical failure, its cause, and its resolution; and an end-of-pilot reflection by each participating teacher and a sample of students, on what they believe the pilot produced and what it did not. The documentation produces a body of specific evidence that your institution can refer to in every subsequent conversation — with parents, with the board, with peer institutions, with the eventual mission coordinator for your district. The alternative, which is the default in most institutional pilots, is a pilot whose outcomes are remembered only through the impressions of the participants, and whose lessons cannot be transmitted beyond the conversation in which those impressions are shared. The discipline of written documentation is, again, unglamorous work. It is also the specific work that converts a pilot from an episode into an institutional learning asset.

The seventh task, throughout the hundred days, is to invest in your own learning. You cannot lead this change well without understanding, at your own working level, what the technology is and what the reform involves. This does not require you to become a technologist. It requires you to spend, over the hundred days,

perhaps fifty to eighty hours in your own study — reading this book and others, attending one or two serious conferences, spending time in institutions where AI-education deployment has already begun, having conversations with people whose judgement on these matters you respect. The leadership of this change requires that you speak from genuine understanding.

At the end of the first hundred days, you have: a senior lead, a listening report, an honest audit, a vision adopted by the governing body, a communication programme, a pilot underway, and the beginnings of your own substantive understanding. You have not deployed any technology at scale. You have not finalised any long-term vendor commitment. You have not made any claim about your institution's AI-education prowess that you cannot back up. What you have is the foundation on which the rest of the work will be built, and the foundation is more important than the visible construction that will follow it.

I want to name, because the omission is deliberate and instructive, what is not on the hundred-day list. There is no device procurement. There is no platform selection. There is no curriculum design committee. There is no teacher-training workshop. There is no announcement of a partnership with a named vendor, no press release, no ribbon-cutting, no visible deliverable that a visitor to the institution would recognise as an AI-education initiative. The absence of these items is not an oversight. It is the specific discipline that distinguishes a hundred days of foundation-building from a hundred days of visible-but-hollow activity, and the institution head who is under pressure from her board, her parents, or her own impatience to show results within the first quarter is the institution head most likely to compromise the foundation in favour of the visible. I would encourage you to resist. The board members, parents, and community observers who want to see visible results within the first hundred days are often the same constituents who will, two years later, be asking why the initiative has not produced the

outcomes it promised. The honest answer, in most cases, will be that the foundations were inadequate. The hundred-day discipline is the specific protection against that outcome.



II. The first year

The first year extends the hundred-day foundations into a broader set of institutional commitments and capacities. The pilot initiated at the end of the hundred days runs through the academic year, generating the learning that the year two rollout will rest on. Alongside the pilot, a set of parallel investments builds the preconditions for the rollout.

Let me specify what those parallel investments look like across the first year. The first is the infrastructure assessment — a school-by-school audit, in institutions with multiple campuses, of the electrical, network, and physical preconditions the deployment will require, with specific timelines and costs attached to remediation in each location. The second is the teacher-training enrolment — identifying the teachers who will receive certification in the first wave, securing their time commitment from the academic timetable, and coordinating with the SCERT or the DIET or the equivalent training institution on the specific programme the teachers will complete. The third is the content audit I described in Chapter Four, applied now not as a diagnostic but as the basis for a specific content-procurement plan that specifies what your institution will use in year two, at which grades, from which sources, at what cost. The fourth is the establishment of the data and privacy practices that will govern the year-two deployment — the specific policies on student data, the parental consent architecture, the vendor contractual requirements, the review mechanisms through which the policies are enforced. The fifth is the development of the parent engagement programme that will run alongside the year-two deployment, with specific

communications, specific forums, and specific feedback channels built in from the beginning rather than retrofitted after parental complaints have accumulated.

These five parallel tracks are not visible to casual observers, and they rarely generate media attention or board enthusiasm. They are the specific work that makes the year-two deployment succeed, and their absence in year one is the reliable predictor of year-two difficulty. An institution head who is asked, at the end of year one, what concrete deliverables the initiative has produced, should be able to answer with specifics on each of the five — the audit results, the teacher enrolment, the content plan, the privacy framework, the parent programme. The answer is unglamorous. It is also the answer that, three years later, will be remembered as the foundation on which the institution's success rested.

The first parallel investment, to be initiated in the first quarter of the year, is teacher development. A structured programme is established through which, over the course of the year, the first substantial cohort of your teachers — perhaps a quarter to a third of your teaching staff — completes meaningful AI-education training. The programme is multi-component. It includes formal training through the National AI Teacher Institute if that institution has been established, or through the best available alternative providers if it has not. It includes peer learning within your institution, where the pilot teachers share what they are learning with their colleagues. It includes time for individual teachers to engage with tools, content, and readings on their own, with institutional recognition that this time is part of their professional duties.

The second parallel investment is infrastructure. The pilot, and the subsequent rollout, require specific technological infrastructure — devices, connectivity, power, secure storage — that most Indian institutions are not currently equipped with at the scale required. The first year is the year in which you make the capital commitments, place the orders, install the infrastructure, and begin

the processes of maintenance and replacement. The infrastructure work is unglamorous. It is also essential — a pilot that runs well but cannot be scaled because the infrastructure is not ready for the scale is a pilot whose lessons are delayed unnecessarily.

The third parallel investment is content. I said in Chapter Four that the content problem is larger than the technology problem, and this is true at the institution level as well as at the national level. During the first year, your institution develops or curates the content it will use in year two — either by adopting content from the National AI Education Content Repository if it exists by then, or by developing the institution's own content with external support, or by evaluating and selecting from vendor offerings under contracts that preserve the institution's control. The content work involves your teachers centrally, because they are the people who will use the content and whose judgement about its suitability matters most.

The fourth parallel investment, which runs throughout the year, is governance. The initial vision adopted in the first hundred days is only a starting point. As the year progresses, specific governance arrangements are developed — the formal role of teachers in curriculum decisions, the data protection policies that govern student information, the procurement standards the institution will apply to vendor engagements, the mechanisms through which parents can raise concerns and have them addressed. Each of these governance artefacts is drafted, consulted on, and formally adopted. The governance work is the institutional analogue of the legislative work I describe in Chapter Thirteen.

The fifth investment, and the one most commonly neglected, is the institutional capacity for evaluation and learning. An AI-education initiative that deploys tools without systematically tracking their outcomes is an initiative that will, three years in, be unable to say whether it has succeeded or failed. The evaluation capacity requires specific data collection, specific analytical

methods, and specific reporting routines that make the institution's outcomes visible to its staff, its parents, and its governing body.

At the end of the first year, you have: a completed pilot, a trained initial teacher cohort, installed infrastructure ready for scaled deployment, a developed content programme, a set of governance artefacts that will guide the second year and beyond, and an evaluation capacity that will track how things actually go. The deployments are still modest in scope. The reform is visible to the institution's community but has not yet transformed the institution. What has been built is the capacity to transform it, over the subsequent year.



III. The two-year path

The second year is the year of scale within your institution. The pilot expands. The teacher-development cohort broadens to reach the majority of your teaching staff. The content is deployed across more grades and subjects. The infrastructure supports the full intended scale. The governance arrangements are applied in practice and refined on the basis of what the practice reveals. At the end of the second year, the institution is, in a specific sense, an AI-education institution — its pedagogical design, its teacher capacities, its technological infrastructure, and its governance all reflect the reform's commitments.

I want to name, however, a set of tasks for the second year that are specific to the transition from pilot-scale to institution-scale, because the transition is harder than many institutional leaders anticipate.

The first task is the scaling of teacher capacity. The first-year cohort was the self-selected volunteers — the teachers most enthusiastic about the reform. The second-year cohort includes teachers who are less enthusiastic, who have more concerns, who are less comfortable with the technology, who may be older and

closer to retirement. The training programme must be redesigned for this cohort. What worked for the volunteers will not work, unmodified, for the next wave. Specific adaptations — more hand-holding, more time for practice, more peer support, more attention to the teachers' own questions and resistances — are required.

The second task is the management of the inevitable problems. Year two is the year in which the infrastructure fails in ways the pilot did not reveal, the tools behave in ways the training did not prepare teachers for, the content proves inappropriate for specific grades or specific subjects, and the student experience reveals gaps the pilot did not surface. These problems are not evidence of failure. They are evidence of scale. The institutional response must be calm, systematic, and transparent — each problem identified, its cause analysed, its remediation specified, and its resolution tracked.

The specific pattern through which year-two problems emerge is worth anticipating, because the pattern is recognisable in retrospect and manageable if you know to watch for it. The problems do not arrive evenly. They cluster in the third month of the academic year, which is typically the point at which the initial enthusiasm has faded, the accumulated gaps in teacher preparation have become visible, and the infrastructure strain of daily use has begun to produce the failures that a three-month pilot did not have time to reveal. The cluster will feel like a crisis. It is, in most cases, not a crisis but the normal transition from the novelty phase to the sustained-operation phase of the deployment. The institutional response should be prepared for it. A specific review meeting at the end of the third month, with a structured agenda that includes each of the foreseeable problem categories — teacher preparation, tool behaviour, content adequacy, student experience, infrastructure reliability — is the mechanism through which the cluster is processed rather than avoided. The meeting is not comfortable. Its output is a specific list of remediations with timelines and responsibilities. The institutions that do this well

emerge from the cluster with a stronger deployment than they had before it. The institutions that try to avoid the cluster, or that treat each problem as an isolated incident rather than as part of the predictable transition, are the institutions that, six months later, are still managing the same problems in more entrenched forms.

The third task is the broadening of engagement. Year one was necessarily focused on the teachers directly involved in the pilot. Year two brings the parents, the students, the alumni, and the governing body into more active engagement with the reform. Specific structures — parent-teacher forums on AI education, student councils that address the experience of the reform from the student's perspective, alumni advisory panels that bring external expertise — are established. The engagement is not ornamental. It is the institutional analogue of the democratic accountability the national reform requires.

The fourth task is the beginning of outward-facing work. By the end of year two, your institution has accumulated specific learning that other institutions could benefit from. Documenting that learning, sharing it in forums, and collaborating with peer institutions on common challenges is both a contribution to the national reform and a continued discipline for your own institution's thinking.

I would add that the specific form in which the documentation is shared matters as much as the fact of sharing. A glossy report that describes the institution's successes, published once at the end of year two, is of limited value to other institutions — and is, in most cases, read more as a marketing document than as an operational resource. A more useful form is a running case-study document, updated quarterly, that specifies what the institution has done, what has worked, what has not, what the institution would do differently, and what specific questions remain open. The case study is written in the institution's own voice, with the institution's own teachers and students quoted where appropriate, and is published through the professional networks your

institution is part of — the association of principals in your state, the peer network of your school's sector, the national associations your board belongs to. The running case study is a different artefact from the year-end report. It is more specific, more immediately useful, and substantially more credible to the peer institution heads who are its primary audience. It is also, for your own institution, a discipline of continuous reflection that the year-end report is not.

The fifth task is the honest assessment of what has and has not worked. At the end of year two, your institution's governing body should receive, and should publish, a candid assessment of the reform's implementation — what has been achieved, what has not, what has been learned, what will be revised. The assessment is not a press release. It is a genuine attempt to capture, at the two-year mark, the institution's actual state. The assessment informs the third year's priorities.

I want to say a word about the specific act of publishing an honest assessment, because this is the point at which many institutional leaders, who have conducted themselves with integrity through the earlier stages of the reform, lose their nerve. The two-year assessment will, in almost every case, contain specific admissions that are uncomfortable. A teacher training programme that reached ninety percent of staff but produced certified competence in only sixty percent. A content deployment that worked well for Classes 3 through 5 but struggled in Classes 6 and 7. A vendor relationship that has not delivered what the contract promised. A parent constituency that remains sceptical. The temptation, when the time comes to publish, is to soften these admissions, to emphasise the positives, to present the difficulties as learning opportunities in language that protects the institution's reputation. The softening is almost always a mistake. The parents, the teachers, and the board members who read the assessment will, in most cases, already know the specific difficulties through their own observation, and an assessment that understates what they know is an assessment whose credibility is immediately damaged.

A candid assessment, by contrast, establishes the institution as an honest partner in the work, invites the community's collaboration in addressing the difficulties, and produces the trust on which the third and fourth years of the initiative will depend. The honesty is, paradoxically, the specific act that protects the institution's reputation in the medium term, because the institution's reputation is built on the credibility of its self-reporting more than on the specific results it reports.



IV. A word on common failure modes

Let me name, briefly, the specific failure modes I have seen most commonly in institutional attempts at this kind of reform, so that you can guard against them in your own work.

The first failure mode is what I call the vendor-led initiative. The institution, confronted with the reform's complexity, outsources the design and implementation to a vendor, which provides the tools, the content, the training, and sometimes even the teacher support. The initiative runs, for a period, reasonably well. Over time, the institution discovers that it has become dependent on the vendor for decisions that should have remained institutional. The vendor's priorities — which are commercial — shape the reform in ways that conflict with the institution's pedagogical priorities. The institution cannot easily disentangle itself, because the switching costs have accumulated. Avoid this by ensuring that every decision of pedagogical consequence remains with your institution, that your teachers retain ownership of their practice, and that the vendor's role is specified narrowly as a supplier of specific tools and services rather than as a general partner in the reform.

A related but distinct discipline, addressed to change management more broadly, is the cultivation of a specific narrative within the institution — a story the institution tells itself about

what the reform is for and why it is being undertaken. The narrative is not marketing. It is the shared set of meanings that allow the teachers, the students, the parents, and the administrative staff to locate their own daily work within the larger purpose the reform serves. Institutions that manage this narrative well treat AI as a tool that extends the teacher's existing commitments rather than as a technology that replaces them, treat the reform as a continuation of the institution's historical mission rather than as a rupture, and treat the specific difficulties of the transition as evidence of the work's importance rather than as reasons for retreat. Institutions that do not manage the narrative allow it to be shaped by vendor marketing, by media coverage, or by the ambient anxieties about artificial intelligence that any institution's community absorbs from the wider environment. The narrative that emerges by default is, in my observation, almost always one that undermines rather than supports the institution's work. The discipline of shaping the narrative deliberately — in staff meetings, parent communications, student assemblies, alumni newsletters — is the specific institutional practice that distinguishes reforms that are owned by their communities from reforms that are merely imposed on them.

The second failure mode is what I call the showcase failure. The institution establishes a visible AI-education deployment in one classroom or one grade that produces impressive results, generates favourable media coverage, and serves as a marketing asset for the institution. The showcase consumes a disproportionate share of the institution's resources and attention, while the rest of the institution proceeds as it was before. After three years, the showcase remains a showcase, and the institution's median classroom has been unchanged. Avoid this by committing, from the beginning, to institution-wide transformation rather than to flagship projects.

The third failure mode is what I call the teacher-bypass failure. The institution deploys AI tools and content without substantial

investment in teacher capacity, on the theory that the tools will speak for themselves and teachers will adapt around them. The teachers, in practice, either disengage from the tools or use them in ways that undermine their intended effects. The students' experience is poor. The institution concludes that the tools do not work, when what has actually not worked is the deployment design. Avoid this by committing, as I have urged throughout this book, to teacher capacity as the first priority of the institutional investment.

The fourth failure mode is what I call the metric collapse. The institution, unable to measure the outcomes it actually cares about, defaults to the metrics the tools themselves generate — engagement hours, platform usage, session counts. The metrics look impressive. The actual learning outcomes, which the metrics are substituting for, deteriorate. Avoid this by measuring what you actually care about — reading fluency, mathematical competence, computational thinking, the specific outcomes the curriculum specifies — and by treating tool-generated metrics as inputs to the investigation of learning rather than as its substitute.

I want to add a fifth failure mode to the four I have named, because it is the one I have seen most often in the past eighteen months and the one most specific to the current moment of the reform. I call it the premature-scale failure. The institution, encouraged by the positive signals from a well-run pilot, decides to scale its deployment rapidly across all grades and all classrooms in the second year rather than proceeding through the more gradual two-year path I have described. The scaling decision is driven by some combination of competitive pressure from peer institutions, enthusiasm from the senior leadership, and the institutional confidence that the pilot's success gives. The scaling, in practice, encounters a set of problems that the pilot did not reveal — because the pilot was conducted by the most committed teachers, with the most engaged students, with the most attentive supervision. The second-year expansion includes teachers who are

not committed, students whose baseline is weaker, and supervision that is stretched across too many classrooms. The problems accumulate. The expansion's quality falls below what the pilot produced, and the institution, which had been positioned to demonstrate a genuinely transformative outcome, finds itself managing a mediocre rollout. Avoid this by treating the pilot's success as necessary but not sufficient evidence for scale, by scaling deliberately across a second cohort that resembles the broader teaching staff rather than the pilot's self-selected enthusiasts, and by accepting that the two-year path is slow precisely because the slowness is what produces the durable result.

A sixth failure mode, less common but worth naming, is what I call the leader-dependence failure. The initiative, in its early years, becomes strongly identified with a specific senior leader — often the principal herself, sometimes the senior lead she designated at the beginning. When the leader moves on — to another school, to another position, to retirement — the initiative loses its primary champion and, within a year or two, the institutional commitments that the leader's presence sustained begin to erode. The specific discipline against this failure mode is to institutionalise the initiative from the beginning — to embed it in the governance structure, the budgetary commitments, the faculty development pathways, and the strategic plan in ways that do not depend on any single individual's continued presence. The institutional embedding is difficult work, and it is specifically the work that a leader with a strong personality and a strong vision may be tempted to neglect in favour of pushing the initiative forward through her own direct engagement. The neglect is understandable and, in the short run, effective. Over a decade, it is the specific failure mode that has undone more educational reforms than any other at the institutional level.



Three decisions for the reader

1. Within the next two weeks, designate your senior lead and begin the hundred-day plan.

Do not wait for perfect conditions. The hundred-day plan is designed to be workable in imperfect conditions, and the conditions for your initiative will not become more perfect by waiting. The specific act of designating the lead and communicating that designation to your institution is what begins the reform, and every subsequent step flows from that beginning.

2. Publish your institution's two-year path to its community.

At the end of the first hundred days, your adopted vision is, in effect, your public commitment. Publish it. Share it with your parents, your teachers, your students, your alumni, your governing body. The publication transforms the vision from an internal document into a public accountability. You will be held to what you publish, which is why the publication matters.

3. Join, or start, a peer network of institution heads engaged in this work.

The specific task of leading an institution through this reform is not a task that should be attempted in isolation. There are, across India, a few hundred institution heads who are serious about this work, and you should know them. Find them through the associations your institution is part of, through the conferences where these matters are discussed, through the professional networks your senior lead can access. Meet with them regularly. Share what you are learning. Learn from what they are learning.



END OF CHAPTER NINE

CHAPTER TEN

Building the Mission — A Statutory Proposal



*The institutional instrument the reform requires,
and why it must rest on legislation.*



Indian public administration has a specific pattern when it approaches a reform of civilisational consequence. The pattern, repeated across the decades with varying degrees of success, is to create a time-bound mission with a clear mandate, adequate resources, statutory protection, and a first-rate team, and to empower that mission to coordinate across the institutional landscape in service of a specific national objective. The pattern has produced, among others, the Unique Identification Authority of India which built Aadhaar, the National Health Authority which built the Ayushman Bharat digital mission, the GST Council which harmonised indirect taxation across a federal constitution, and, going back further, the Indian Space Research Organisation which built the national space programme from scratch in a geopolitical context that offered no shortcuts. These missions have succeeded to varying degrees. Their successes have been uneven, their cultures idiosyncratic, their politics sometimes difficult. But they are the specific organisational form that Indian public administration has developed, by trial and correction over seven decades, for undertakings of a scale and ambition that the ordinary ministerial structure cannot accommodate.

The mission form has, of course, its own characteristic failure modes, and it would be dishonest to propose it as the institutional answer without acknowledging them. Missions can ossify into self-sustaining bureaucracies that outlive their original purpose. They can accumulate political patronage that distorts their merit-based culture. They can become captured by the vendor ecosystems they regulate, particularly when those ecosystems are well-resourced relative to the mission's own staff. They can, in the hands of a weak CEO, become fragmented and directionless, with each directorate pursuing its own priorities under the nominal umbrella of the mission's mandate. They can, in the hands of a strong but narrow CEO, become instruments for a particular vision of the reform that differs from the vision the legislation authorised. Each of these failure modes is visible, on the record, in the histories of previous Indian missions, and the design of the AI Education Mission should anticipate them. The specific protections I have proposed – the fixed tenure for the CEO, the statutory sunset at fifteen years, the Governing Council's composition with independent members and state representation, the protections for the Evaluation and Equity Directorate's independence – are the specific arrangements through which these failure modes are, to the extent possible, resisted. No arrangement eliminates the failure modes entirely. The protections reduce their likelihood, and the reform should not proceed on the assumption that any of them is impossible.

The argument of this chapter is that the AI-education reform is, by its scale and ambition, precisely the kind of undertaking for which a statutory mission is the right institutional instrument; that the alternatives – an enhanced committee under an existing ministry, an autonomous society under the Ministry of Education, a project office under NITI Aayog – have specific weaknesses that the mission form is designed to overcome; and that the specific design of the mission, which I will lay out in this chapter, must draw on the lessons of previous Indian mission undertakings while responding to the specific characteristics of the education sector.

I want to note, because the reader who has worked in or near Indian public administration will recognise the pattern, that the proposal for a new statutory mission will encounter a specific form of institutional resistance that is worth naming. The resistance comes not from opposition to the reform itself — most serious officials across the central and state machinery support the reform's broad purpose — but from the accumulated institutional interests of the existing bodies whose mandates overlap with the mission's proposed remit. Each of those bodies has its own leadership, its own staff, its own constituencies, and its own view of what its role should be in any new arrangement. A proposal that transfers functions to a new body is experienced, by the existing bodies, as a diminishment, even when the functions in question are ones they have not been performing at the scale or quality the reform requires. The political task of establishing the mission therefore includes a specific negotiation with each of the affected bodies, in which the mission's mandate is clarified, the existing bodies' continuing roles are honoured, and the specific points of coordination are agreed. The negotiation cannot be avoided. It can be conducted well or badly, and the manner of its conduct shapes the mission's subsequent relationships across the institutional landscape.




Figure 10.1 | National AI Education Mission — statutory architecture.



I. Why a statutory mission, and not an alternative

Let me begin with the question that the Cabinet note must answer before any other. Why does this reform require a new statutory body, when the existing institutional landscape is already populated with bodies that could, in principle, take on the work? The Ministry of Education has departments that address every aspect of the reform's remit. NCERT has the curriculum capacity. NIEPA has the planning and research capacity. NCTE has the teacher-education regulatory capacity. CBSE has the national-scale assessment capacity. The SCERTs have the state-level adaptation capacity. NITI Aayog has the inter-ministerial convening capacity. MeitY has the technology-infrastructure capacity. A determined reformer, working through these existing bodies, could in principle coordinate the reform without creating anything new.

The answer is that she could not. The coordination problem the reform presents is not a problem that the existing bodies, even if each performed perfectly within its own mandate, can solve among themselves. The problem is that no one of them has the authority, the resources, or the mandate to direct the others, and no mechanism exists for resolving disagreements among them at the pace the reform requires. A reform that depends on twelve institutional bodies acting in coordinated service of a single national objective, without any one of them being able to direct the others, is a reform that will proceed at the pace of its slowest link, with the pace of each link determined by its own institutional priorities rather than by the reform's schedule. This is, empirically, what has happened to previous attempts at comparable scale.

The alternatives to a statutory mission each have specific weaknesses that recommend against them. An enhanced committee under the Ministry of Education would be dominated by the Ministry's own institutional culture and priorities, and would struggle to command the engagement of MeitY, NITI Aayog, and the state governments whose cooperation the reform requires. An autonomous society under the Ministry would have

greater flexibility than a ministerial committee but would lack the statutory protections against budget reductions, leadership changes, and mandate dilutions that the reform's ten-year horizon requires. A project office under NITI Aayog would have the inter-ministerial convening capacity but not the budgetary and operational authority to direct implementation.

The statutory mission has three specific advantages that together justify the additional effort its creation requires. It has durable authority — its powers are conferred by Parliament rather than by executive order, which means they survive changes of government and cannot be diluted through ordinary administrative revision. It has independent resources — its budget is a dedicated line in the Union Budget, not subject to the year-to-year discretion of the finance ministry in the way that ministerial allocations are. It has recruitment flexibility — it can hire outside the ordinary civil-service rules, which is essential for a reform that requires specialised technical and pedagogical talent that the existing cadres do not uniformly supply. Each of these advantages corresponds to a specific failure mode that has affected previous education reforms, and the statutory mission form is the specific remedy that Indian institutional experience has developed for these failures.

A fourth advantage, which the preceding list does not explicitly name, is the specific protection that a statutory instrument provides against the kind of quiet administrative revision through which substantive commitments are sometimes eroded in Indian public administration. A commitment that exists in a Cabinet note, or in a ministry resolution, or in a departmental circular can be revised by a subsequent Cabinet note, a subsequent resolution, a subsequent circular — often with no public debate and no parliamentary scrutiny. A commitment that exists in a statute requires a parliamentary amendment to revise, which produces specific public visibility, specific parliamentary scrutiny, and specific political costs that the lesser instruments do not entail.

The mission's substantive commitments — the compensation framework for teachers, the equity commitments to disadvantaged states, the independence protections for the Evaluation and Equity Directorate, the children's AI safety code — should be codified in the Act rather than left to the mission's own internal policy documents. The codification does not eliminate the possibility of subsequent revision, but it places the bar for revision substantially higher, and the higher bar is the specific protection that a reform of this duration requires. The discipline that follows for the Act's drafters is that the Act should be more detailed, and more substantive, than the administrative default for such instruments — with specific provisions written out in the legislation rather than deferred to rules that the executive can revise at will. The discipline lengthens the drafting process. It also produces the specific instrument that can sustain the reform across the political transitions of the coming fifteen years.

I want to name these failure modes specifically, because the abstraction of institutional weakness conceals the concrete mechanics of how reforms have been diminished. The first failure mode is the leadership-turnover failure. The Ministry of Education has, over the past fifteen years, seen eleven different Secretaries occupy the position of Secretary (School Education). Each brings her own priorities, her own understanding of the sector, her own political relationships with the Minister, the state governments, and the other central departments. Each inherits the previous Secretary's commitments and, within a year or two, is in the position of deciding which of those commitments to maintain and which to redirect. A reform that requires ten or twelve years of sustained institutional commitment cannot proceed through this pattern without being progressively reshaped by each change of leadership. The statutory mission's CEO, by contrast, holds a fixed tenure that spans multiple Secretaries, and her strategic framework, once adopted by the Governing Council, is not subject to reinterpretation by each new departmental head. The second failure mode is the budget-dilution failure. Ministerial allocations

are, in each annual Union Budget cycle, subject to negotiation between the ministry and the finance ministry, with the negotiation shaped by the political salience of the moment, the fiscal space available, and the other priorities competing for attention. A commitment that was adequate in year three may be diluted in year four to accommodate a different priority, without the dilution being visible until the sector begins to experience its consequences. The statutory mission's budget, protected by the Act that establishes it, is not subject to the same annual renegotiation, which is the specific protection that makes a decade-long commitment credible. The third failure mode is the coordination-collapse failure. Reforms that depend on the coordinated action of multiple ministries and multiple levels of government are, in the absence of a dedicated coordinating institution, coordinated through committee meetings whose minutes are produced but whose decisions are rarely enforced. The statutory mission's authority to issue binding directives within its mandate is the specific protection against this pattern, and the protection is not replicable through any non-statutory instrument.



II. The legal instrument

I propose that the mission be established through an Act of Parliament — the National AI Education Mission Act — that specifies its mandate, its powers, its governance, its budget, and its protections. The Act should be tabled within twelve months of the Cabinet decision that initiates the reform. In the interim, a Cabinet decision and a Ministry of Education resolution can establish the mission on a provisional basis, but the statutory foundation must follow within the twelve-month horizon, because a mission that operates without statutory foundation is a mission whose permanence is not assured and whose ability to make the long-term commitments the reform requires is compromised.

The Act's structure, which I sketch below, draws specifically on the Aadhaar Act of 2016, the National Health Authority's constituting framework, and the Right of Children to Free and Compulsory Education Act of 2009. Each of these Acts offers lessons – positive and negative – for the drafting of the National AI Education Mission Act. The drafter's task will be to synthesise those lessons into an instrument specific to the reform's requirements.

The Act's Preamble should establish the constitutional basis for the mission and its overall purpose. The constitutional basis rests on the concurrent jurisdiction of the Centre and the states over education, and on the Directive Principles regarding state provision for education. The Preamble should name, specifically, the reform's civilisational purpose – the preparation of Indian children for a world that artificial intelligence will reshape – and its commitment to equity across the dimensions I specified in Chapter Seven.

Chapter I of the Act should define the key terms – AI education, AI tutor, mission, Governing Council, Chief Executive Officer, state coordinator, and the other specific terms the Act uses. The definitional work is unglamorous but consequential, because much of the subsequent interpretive difficulty in Indian legislation arises from ambiguous definitions.

Chapter II should establish the mission as a body corporate with perpetual succession and a common seal, specify its headquarters and field offices, and create the Governing Council as its supreme decision-making body. The composition of the Governing Council should be specified in the Act itself rather than delegated to rules, because the Council's composition is the specific safeguard against capture by any single political or bureaucratic interest.

Chapter III should specify the mission's mandate and the specific authorities through which it performs them. The five

statutory authorities — to issue binding standards, to disburse direct grants, to operate the content repository, to commission evaluations, to coordinate private deployments — should be codified here, with the specific conditions and limitations on each authority.

Chapter IV should establish the office of the Chief Executive Officer, specify the qualifications and tenure, set out the recruitment process, and enumerate the CEO's specific powers.

Chapter V should specify the mission's budget — not the actual amount, which will be determined in the annual Union Budget, but the specific appropriation mechanism, the protection of the mission's budget against arbitrary reduction, and the rules governing the mission's discretion over its allocated funds.

Chapter VI should establish the Act's relationship with state governments, recognising the Concurrent nature of education while specifying the areas in which the mission's standards are binding on states and the areas in which state adaptation is permitted.

Chapter VII should establish the mission's accountability mechanisms — the reporting to Parliament, the publication of its performance against specified metrics, the audit arrangements, the mechanisms for citizens to raise grievances, and the protections for whistleblowers within the mission and its contractors.

Chapter VIII should address the mission's sunset and review. The mission I propose has an initial statutory life of fifteen years, with provision for a comprehensive Parliamentary review at year ten and possible extension through amendment. The sunset is not symbolic. It is the specific mechanism that prevents the mission from becoming a permanent feature of the institutional landscape independent of whether it continues to serve its purpose.

Chapter IX should specify the transition arrangements — the specific steps through which the mission takes over functions currently performed by other bodies, the arrangements for

personnel transfers, the protection of ongoing contracts and commitments, and the mechanisms for resolving disputes about jurisdiction.

I want to add one further chapter to the structure I have sketched, because the omission would be consequential. Chapter X of the Act should establish the provisions for the mission's engagement with the private sector — the specific rules governing its procurement, its contracting, its partnerships with private vendors and civil society organisations, and its conflict-of-interest protections. The rules are not trivial. Indian public procurement, operating under the General Financial Rules and related frameworks, has a characteristic inflexibility that was developed to protect against corruption but that has, in its accumulated effect, produced a procurement culture that favours compliance over quality and that frequently selects bidders who have learned the procurement process over bidders who can actually deliver the work. The mission requires a more flexible procurement framework — modelled on the arrangements that UIDAI and NPCI have operated under — that permits quality-based selection, multi-year framework agreements, and specific carve-outs for technology and content procurement where the standard rules produce poor outcomes. The framework must, of course, include robust conflict-of-interest protections, mandatory disclosures, and audit arrangements. But the framework must also be capable of producing, in practice, the partnerships with high-quality private providers that the mission's work requires, and the current default procurement framework is not. Chapter X of the Act is the specific legislative vehicle through which this flexibility is granted, with the specific safeguards codified. Its drafting is one of the most consequential pieces of work in the Act as a whole, and it should not be left to the routine drafters of the general procurement rules.



III. The Governing Council

The Governing Council is the mission's supreme decision-making body. Its composition, its procedures, and its authorities are the specific arrangements through which the mission is held accountable to the broader political and professional community while retaining the operational autonomy the reform requires.

I propose that the Governing Council comprise the following members. The Prime Minister, as chair, whose involvement at this level signals the reform's political priority and whose authority resolves disputes the Council cannot otherwise resolve. The Union Minister of Education, as vice-chair, whose day-to-day political stewardship of the reform falls to her. The Union Ministers of Finance, Electronics and Information Technology, and Skill Development and Entrepreneurship, whose participation ensures that the related portfolios are engaged in the reform's decisions rather than working at cross-purposes with it. Six Chief Ministers, serving rotating two-year terms drawn from across the political spectrum and the regional geography, whose participation ensures that state perspectives are incorporated. The Vice-Chairperson of NITI Aayog. The Principal Scientific Adviser to the Government of India. Six independent members, appointed by the Prime Minister from among educators, technologists, children's rights advocates, and other qualified individuals, whose independence ensures that professional and civil-society perspectives are represented at the highest level.

The Council meets twice a year. Its decisions are binding on all constituent bodies of the mission. Its quorum is ten, with the requirement that the quorum includes at least one representative from the states, at least one independent member, and the Union Minister of Education or a duly authorised representative. Its proceedings are, in part, public – the agendas and decisions are published, though the deliberations are not.

The specific inclusion of six Chief Ministers is a considered choice, drawn from the GST Council's precedent, and is the specific mechanism through which federal legitimacy is built into the mission's design. A mission that excludes Chief Ministers from its supreme decision-making is a mission that will be perceived, by the states, as a Central programme being imposed on them. A mission that includes them – and includes a range of them, including from states governed by parties in opposition to the party in power at the Centre – is a mission whose decisions carry weight across the Indian federation.

The independent members deserve specific comment. Their inclusion is essential, because the Council otherwise risks becoming an echo chamber of elected and appointed officials whose perspectives, however expert, do not capture the full range of relevant considerations. The independent members must be genuinely independent – appointed on merit, with backgrounds that contribute specific expertise the Council needs, and with the personal standing to speak freely in meetings where the Prime Minister is present.

I want to be specific about what independence in this context requires, because the phrase has become so conventionalised in Indian institutional design that it has lost much of its operational force. The six independent members should meet four specific tests. The first is the financial test – the individual should not have, in the five years preceding her appointment, held a financial interest in any company, foundation, or organisation whose business is materially affected by the reform's decisions, and should not hold such an interest during her tenure. The second is the employment test – the individual should not be in current or recent employment with any such organisation, and should not take up such employment for a period of three years after the conclusion of her tenure. The third is the political test – the individual should not hold, or have recently held, an elected office or a formal position in a political party, and should not have been

nominated as the representative of a political party to any other body. The fourth is the professional-standing test — the individual should have, at the time of her appointment, a professional standing and a body of work in the relevant field that is credible to her peers independently of any association she has had with the government. The four tests together exclude many of the candidates who, in current Indian institutional practice, are routinely appointed as independent members. They also produce, in aggregate, a pool of candidates whose independence is structurally protected rather than merely nominal, and whose participation in the Council can be expected to introduce genuinely independent perspectives rather than to validate decisions that have already been taken elsewhere. The Act should codify the four tests as specific eligibility criteria, and the appointment process should include a transparent review mechanism that tests each candidate against each of them.



IV. The Chief Executive Officer

The CEO is the operational head of the mission. Her appointment, her tenure, her qualifications, and her powers are among the Act's most consequential provisions, because the effectiveness of a mission is shaped, more than by any other single factor, by the quality of its chief executive.

I propose that the CEO should hold the rank of Secretary to the Government of India, serve a fixed five-year term that can be renewed once, and be required to be, above all else, an educator.

The rank of Secretary matters because it determines the CEO's operational authority within the Central institutional landscape. A CEO holding lower rank is a CEO who will, in practice, be unable to command the engagement of the departments whose cooperation the mission requires. A CEO holding the Secretary's rank operates within the existing hierarchy of Central authority,

which is the specific mechanism through which her decisions carry weight across the government.

The fixed five-year term, renewable once, is a specific protection against two distinct risks. The first is the risk of the CEO being moved before the reform's long-horizon work is complete. The second is the risk of the CEO becoming entrenched beyond the point where her continuation serves the mission.

The requirement that the CEO be, above all else, an educator is the specification I most want to emphasise and the one most likely to be resisted in the drafting process. The default instinct in Delhi, when appointing the head of a mission with a substantial technology component, is to appoint a technologist or a senior civil servant with a technology portfolio. For the mission's reasons this default is wrong. The mission's hardest problem is not a technology problem. The hardest problem is the delivery-to-the-classroom problem, which requires the kind of institutional understanding and pedagogical judgement that someone whose career has been about teachers, children, and learning outcomes will bring. A technologist at the head of the mission will, however good her intentions, optimise the mission for the metrics technologists understand – engagement, platform usage, deployment counts – rather than for the metrics the reform actually needs – literacy gains, teacher capacity developed, children's learning outcomes improved.

I want to anticipate the objection that the preceding paragraphs have understated the importance of technological fluency in the CEO's role, because the objection will be raised and deserves a direct answer. The AI-education mission is, undeniably, a mission with a substantial technology dimension. The platforms, the tools, the data infrastructure, and the vendor relationships all require technological judgement at the top of the mission. A CEO who is technologically illiterate cannot adequately supervise the Infrastructure and Platforms Directorate, cannot adequately evaluate the Repository's technical choices, cannot adequately hold

the private-sector partners to account on the technical dimensions of their deliverables. The objection is valid. The answer is not that technological fluency does not matter – it matters substantially – but that it is acquirable, at the level the CEO requires, through deliberate engagement over the first year of her tenure, whereas the pedagogical and institutional judgement that the role requires is the product of decades of professional formation that cannot be acquired on the job. The CEO who is, first, an educator can learn the technology she needs to supervise, through sustained engagement with the Chief Technology Officer she appoints, through the technical briefings the directorate delivers, through her own reading and conversation. The CEO who is, first, a technologist can learn the pedagogy she needs, in principle, through the same process – but the evidence from comparable institutional appointments in Indian education suggests that the technological-first leader does not, in fact, make this pedagogical acquisition the priority of her tenure, and defaults instead to the technology-centric framing of the reform that her formation has equipped her for. The asymmetry is the specific reason the pedagogical-first appointment is the right default, even granting the technological substance of the mission's work.

The CEO's deputy – the Chief Technology Officer of the mission, reporting to her – should be the technologist the mission also requires. The CTO's appointment is substantially easier to make than the CEO's, because the specifications are clearer and the pool of candidates is larger. The CTO's relationship with the CEO is the specific pairing through which the mission's dual technical-pedagogical character is balanced at the top. Both appointments matter. Neither substitutes for the other. The Act should specify both positions, with clear separation of their mandates and clear arrangements for their interaction.



V. The four directorates

Beneath the CEO, the mission is organised into four directorates, each led by a director with the rank of Additional Secretary. The directorates correspond to the four substantive dimensions of the reform – curriculum and content, teacher capacity, infrastructure and platforms, and evaluation and equity – and their design reflects the specific division of labour the reform requires.

The Curriculum and Content Directorate is responsible for the mission's work on the grade-by-grade roadmap, the National AI Education Content Repository, and the ongoing curriculum refinement that the reform's evolution will require. It works in close partnership with NCERT, the state SCERTs, and the statutory curriculum bodies whose jurisdictions overlap with the directorate's remit. Its director is, appropriately, an educator with substantial curriculum-development experience.

I want to say a word about the specific form of the partnership with NCERT, because the relationship is among the mission's most sensitive and the specific arrangements will shape the directorate's effectiveness. NCERT is an existing statutory body with a distinguished history, a large staff, and substantial institutional equities at stake in any new arrangement that touches curriculum. A mission that is seen by NCERT's leadership as usurping NCERT's mandate will find, in practice, that the cooperation it formally requires is grudgingly provided and that its own work is slowed by the need to navigate the resulting tensions. A mission that is seen as respecting NCERT's role, while operating in the specific space that NCERT's existing mandate does not adequately cover, will find cooperation forthcoming and mutual learning possible. The specific architectural choice I would propose is that the Curriculum and Content Directorate operate, for its curriculum work, as a commissioner of NCERT rather than as a competitor to it. The directorate sets the specifications, provides the budget, and holds the final quality judgement. NCERT, with its established staff

and its constitutional relationship with the state SCERTs, executes substantial portions of the curriculum-development work under the directorate's commissioning. The arrangement is less tidy than a full vertical integration within the mission, and it requires specific protocols to manage the inevitable points of friction. It also preserves NCERT's institutional standing, draws on its accumulated capacity, and creates a structural alliance that protects the reform from the kind of institutional opposition that has weakened earlier attempts. The specific protocols – delegation agreements, joint review mechanisms, agreed escalation pathways – should be developed in the mission's first year, with NCERT as a full participant in their design.

The Teacher Capacity Directorate houses the National AI Teacher Institute I proposed in Chapter Six, the ongoing training programmes the reform operates, the certification framework, and the peer-mentor networks that sustain ongoing professional learning. It works in partnership with NCTE, the state teacher-education infrastructure, and the teachers' council within the mission's governance.

The Infrastructure and Platforms Directorate addresses the technology layer of the reform – the DPI platform, the procurement frameworks, the device deployment support, the connectivity and power infrastructure that schools require. It works in partnership with MeitY, the CDAC, the NIC, and the state-level technology agencies whose support the deployment requires.

The Evaluation and Equity Directorate houses the mission's monitoring and evaluation function, its equity accountability, the public-facing dashboard, the citizen grievance mechanism, and the commissioning of independent external evaluations. Its work is the mission's internal conscience, and its independence from the other directorates is important.

I want to be specific about what the independence of this directorate requires in operational terms, because the phrase "independence from the other directorates" is easily stated and rarely implemented in Indian public administration. The director of the Evaluation and Equity Directorate should report, for purposes of her substantive evaluations, directly to the Governing Council rather than to the CEO, and her evaluations should be published in the form the directorate produces them rather than in a form that the CEO or the other directors have approved. Her budget should be protected from annual negotiation with the other directorates, and her staff should be recruited through a process that does not include the other directorates' approval. She should have the authority to commission external evaluations from independent research organisations, and the commissioned evaluations should be published within a specified period of their completion, regardless of whether the findings are uncomfortable for the mission. She should, finally, have the explicit protection of the Act against retaliation — against her, against her staff, or against her budget — in response to evaluations that the mission's leadership finds unwelcome. These protections are substantial. They are also the specific arrangements through which the directorate can function as the mission's conscience rather than as its publicist, and in the absence of the protections the directorate will, within three or four years, have become the latter rather than the former. Indian institutional experience on this specific question is not encouraging, and the drafters of the Act should consult the arrangements that, in other sectors, have worked and not worked, with specific attention to what distinguishes the two.



VI. The federal question

The mission operates in a federal constitutional context in which education is on the Concurrent List, with both the Centre and the states having jurisdiction. The specific arrangements for

the mission's interaction with state governments are among the Act's most delicate provisions, because they must reconcile the requirement for national coherence with the constitutional autonomy of state governments and the practical reality that the ground-level machinery of school education is overwhelmingly state-run.

The approach I propose has four elements. The first is the national-standard core — the curriculum framework, the content repository, the device and platform standards, the safety and privacy code — which the mission issues and which is binding on government and government-aided schools across the country, with the states' participation in the standard-setting secured through their representation on the Governing Council and through the state coordinators' voices in the mission's operational consultations.

The second element is the state-specific implementation pathway — the specific plan that each state develops, in consultation with the mission, for how it will implement the reform within its own context. The pathway specifies the state's adaptations to the national core, the pace of its rollout, the infrastructure and training investments it will make, the equity commitments it will sustain, and the milestones against which its progress will be measured.

The third element is the outcome-linked grant structure — the specific fiscal transfers through which the mission supports each state's implementation. Sixty percent of the mission's budget is disbursed to states as outcome-linked grants, with pre-published metrics on device deployment, teacher certification, and learning outcomes; the remaining forty percent is spent centrally on content, platforms, and the mission's own operations.

The fourth element is the dispute-resolution mechanism. Disputes between the mission and a state, or between states, are resolved through a specific procedure — negotiation, mediation by

a committee that includes state representation, and final resolution by the Governing Council if the earlier stages do not produce agreement.



VII. The mission's first year

The Act's passage is the beginning of the mission's work, not the end of the reform's preparation. The mission's first year is a specific undertaking, with specific tasks that must be completed if the subsequent years are to proceed well.

In the first quarter, the CEO is appointed, the four directors are recruited, the initial staff cadre is drawn from deputation and lateral hires, the interim headquarters is established, and the Governing Council holds its first meeting. The Council's first decisions — the initial budget, the five-year strategic framework, the provisional operational priorities — set the direction for the subsequent work.

In the second quarter, the mission produces the draft standards documents — the curriculum framework, the content specifications, the device and platform standards, the teacher certification framework, the children's AI safety code — for public consultation. The consultation runs for sixty days, with specific outreach to teacher organisations, parent associations, civil-society groups, state governments, and the private sector.

In the third quarter, the finalised standards are issued, the first procurement framework tenders are published, the National AI Teacher Institute is formally established, and the National AI Education Content Repository project is launched at NCERT with its initial commissioning calls.

In the fourth quarter, the five pilot districts are identified and agreed with the relevant state governments, the first framework agreements are signed with qualified vendors, the teacher-training

cascades are initiated, and the mission's public dashboard goes live with the initial input metrics. By the end of the first year, every foundational element of the reform exists in formal institutional terms, even if it has not yet reached a classroom at scale.

A specific note on the five pilot districts, because the selection of these districts is more consequential than it may at first appear. The temptation, in selecting pilot districts, is to choose districts where the conditions for success are most favourable — districts with strong SCERT support, districts with competent district education offices, districts where the political environment is favourable, districts where the infrastructure is relatively intact. The pilots in these districts are more likely to produce visible success in the first two years, and the visible success supports the political narrative the mission needs in its early years. The temptation is understandable. It is also, in my view, wrong. A pilot conducted in favourable conditions produces results that cannot be generalised to the national deployment, because most of the country does not operate under the favourable conditions the pilot benefited from. The pilots should instead be selected to represent, in aggregate, the range of conditions the national deployment will encounter — one district each from the relatively strong state systems of Kerala and Tamil Nadu, one district each from the middle-capacity systems of Karnataka and Rajasthan, one district from a weaker-capacity system such as Bihar or Uttar Pradesh, and perhaps a sixth from a north-eastern state where the specific challenges of geography, language, and political context produce conditions unlike those in the rest of the country. The pilots will be harder. Some will, in their first year, struggle more than a favourably-conditioned pilot would have. The institutional learning the struggles produce is substantially more valuable than the learning a smooth pilot produces, and the learning is what the subsequent scaling rests on. The Governing Council should, in its first meeting, adopt selection criteria for the pilots that explicitly prioritise representativeness over ease of success, and the mission's

internal culture should be shaped from the beginning to value the hard-pilot learning over the easy-pilot vindication.

The first year is entirely about architecture. No classroom is touched by the mission's direct work. The work is in Delhi, in the state capitals, in the SCERTs and DIETs. It is the work on which everything else rests. A mission whose first year is spent on visible public-facing deployments — on ribbon-cuttings and flagship pilots — is a mission whose subsequent years will be spent struggling with the foundational arrangements that were not adequately built in year one.



Three decisions for the reader

1. If you are in government, draft the Cabinet note within ninety days.

The specific act that sets the mission in motion is the Cabinet note that seeks approval for its establishment. The note should be drafted within ninety days of the political decision to proceed, tabled within a further thirty days, and approved with the associated budget provision within a further sixty days. The statutory Bill follows within twelve months. If the Cabinet note is delayed beyond ninety days, the political energy for the reform will begin to dissipate, and the subsequent steps will become more difficult than they needed to be.

2. If you are a Chief Minister or your office, signal engagement early.

The mission's architecture depends on substantive engagement from state governments across the political spectrum. A Chief Minister whose state signals early that it intends to be a full partner in the reform — through representation on the Governing Council, through rapid development of the state-specific

implementation pathway, through visible investment in the state's own capacity — is a Chief Minister whose state will benefit more from the reform than one whose engagement is delayed or grudging.

3. If you are an independent member candidate, prepare.

The six independent members of the Governing Council will be among the most consequential appointments the reform makes. If your professional background qualifies you for consideration, prepare now for the possibility that your name will come forward. Preparation involves thinking seriously about what independence means in this context, what contribution you would want to make, and what you would be willing and unwilling to do to preserve your independence while exercising influence.



END OF CHAPTER TEN

CHAPTER ELEVEN

Paying for It — The Fiscal Architecture



The numbers, the sources, and the mechanisms that transform intention into entitlement.



Every serious reform in Indian public policy succeeds or fails on a question of money. The money question has three dimensions — the quantum required, the sources from which the quantum is raised, and the mechanisms through which it flows from source to use — and a reform that addresses any one of these without the others will find itself, some years into its implementation, encountering a fiscal crisis that an earlier design could have avoided. This chapter addresses all three dimensions, with the specificity the reform requires and the honesty that a policy-influence book demands.

I want to say a word, before I proceed to the numbers, about why the fiscal discussion is so often neglected in the kind of book this is — and, indeed, in much of the Indian policy-reform literature. The neglect is not accidental. Writers on education reform, including specialists whose technical expertise is substantial, often treat fiscal questions as matters for a separate conversation conducted by finance-ministry officials and economists. The separation is convenient. It allows the educational argument to be made in its own terms, without the pressure of the fiscal trade-offs that would, if confronted, compromise some of the more ambitious recommendations. The separation is also, I believe,

a mistake. A reform proposal that has not been costed honestly is a reform proposal whose political and fiscal credibility is, from the start, compromised. The Cabinet members who will ultimately decide whether to commit to this reform will not do so on the strength of the educational argument alone. They will ask, specifically, what the commitment costs, where the money will come from, and what trade-offs it implies against other priorities. A proposal that cannot answer these questions is a proposal that will not survive its first Cabinet meeting. I therefore treat the fiscal chapter not as an accounting appendix to the substantive argument but as a load-bearing part of it, and I have tried to present the numbers with the honesty the stakes demand.

Let me begin with the total envelope. A five-year rollout of the AI-education reform, covering all government and government-aided schools and supporting the broader institutional architecture the reform requires, demands a total fiscal envelope of approximately one hundred and twenty thousand crore rupees, or approximately twenty-four thousand crore per year at full deployment. This estimate breaks down as follows.

I want to note, before the breakdown, that the envelope I am proposing is materially larger than the current annual Union expenditure on the Samagra Shiksha programme, which is the principal existing vehicle for centrally-supported school education in India. This is a substantial observation. It means that the reform cannot be accommodated within the existing school-education allocation without either displacing a substantial fraction of that allocation or obtaining fresh money at a scale that is politically demanding. I will return to this later in the chapter. For now, the reader should know that the numbers are, by the standards of existing Indian school-education expenditure, large, and that the political work of justifying them is not trivial.

Device procurement, approximately forty thousand crore over five years. This component supports the deployment of approximately fifty million learner-accessible devices across

government schools, with varying specifications across the three stages, at unit costs that reflect Indian manufacturing capacity. The estimate includes replacement cycles — devices have a useful life of approximately four years, which means that by year five of the reform, the first wave of devices begins to need replacement.

Connectivity and power infrastructure, approximately fifteen thousand crore. This component addresses the broadband, WiFi, and electrical infrastructure that many government schools currently lack, particularly in rural and semi-urban contexts. The component also supports the backup power arrangements — solar installations where grid power is unreliable, battery backup systems where outages are frequent.

Content production and platform engineering, approximately ten thousand crore. This component funds the National AI Education Content Repository I specified in Chapter Four, the DPI platform I specified in Chapter Five, and the ongoing content development and platform maintenance that the reform requires.

Teacher training and certification, approximately twenty thousand crore. This component funds the National AI Teacher Institute I specified in Chapter Six, the cascade training through the SCERT-DIET-BRC chain, the peer-mentor networks, and the ongoing professional-learning infrastructure. The component does not include the compensation increments tied to certification, which I will treat as a separate line because its political and fiscal characteristics are distinct.

Mission operations, monitoring, and evaluation, approximately five thousand crore. This component funds the mission's own institutional existence — the directorates, the field offices, the evaluation function, the public-facing dashboard, the grievance mechanism. The component is deliberately modest relative to the overall envelope.

A separate accounting question the fiscal framework must address, because it shapes which of the above components is

politically easier to fund, is the capex-opex split. Of the hundred-and-twenty-thousand-crore envelope, approximately fifty-five thousand crore is capital expenditure — devices, infrastructure, platform engineering, the initial Repository build — and approximately sixty-five thousand crore is operating expenditure — teacher training and certification delivery, content maintenance, mission operations, ongoing connectivity charges, Equity Fund disbursements. The split matters because Indian fiscal politics treats capex and opex asymmetrically. Capital outlays are politically easier to justify and tend to be released more reliably, because they produce visible physical deliverables and because they can be spread across state-level counter-part funding. Operating outlays are politically harder — they require sustained annual commitments, they produce less visible deliverables, and they are more vulnerable to dilution in subsequent budget cycles. The reform's fiscal architecture must therefore defend the opex commitments with particular care, through the statutory protections I proposed in Chapter Ten and through the specific public-commitment mechanisms I proposed earlier in this chapter. A reform that is well-funded on capex and under-funded on opex is a reform that will, within three years, have deployed substantial physical infrastructure that it cannot operate.

Equity Fund, approximately thirty thousand crore. This component, which I specified in Chapter Seven as the specific fiscal mechanism through which the reform's equity commitment becomes real, supports the states and districts with the weakest starting baselines. It funds additional device density, intensive teacher mentorship, infrastructure upgrades, and specific content and training adaptations for the areas of deepest disadvantage.

These components sum to one hundred and twenty thousand crore over five years. The components should be understood as approximate magnitudes rather than as precise budget lines — the detailed budgeting will involve substantial refinement as the mission's work develops — but the overall magnitude is, I believe, robust. It is the amount that doing the reform seriously actually

costs. Proposing to do it for less is proposing either to do less than the reform requires or to pretend that costs can be borne elsewhere when they cannot.

I want to subject this estimate to an honesty test, because the reader deserves to know how robust the number actually is. The estimate is built up from assumptions that are themselves uncertain — the per-school device cost, the cost per teacher of certification training, the cost per language of content origination, the overhead rate for the mission's own operations. If each of these underlying assumptions is off by twenty percent in the direction of higher cost, the overall envelope rises to perhaps one hundred and fifty thousand crore over five years. If the assumptions are off in the direction of lower cost — for example, if Indian manufacturing capacity brings device prices down faster than I have assumed, or if Bhashini's infrastructure reduces the content-origination burden more than I have credited — the envelope could be as low as ninety thousand crore. The range of honest uncertainty is therefore between ninety and one hundred and fifty thousand crore, and the point estimate of one hundred and twenty is the midpoint of a genuinely wide confidence interval. I mention this not to undermine the estimate's usefulness but to caution against treating it as more precise than it is. The Cabinet note that commits to the reform should commit to a budget envelope with explicit provision for revision in years three and four, when the actual unit costs will be known with much greater precision, and the mid-course revision should be treated as a normal feature of the fiscal architecture rather than as an embarrassing departure from the original plan.

I should also address, because the question will be asked, how the numbers compare to what other countries have spent on comparable reforms. The comparisons are difficult to make cleanly, because the starting conditions differ — the existing infrastructure, the baseline teacher quality, the scale of the student population, the fiscal framework of the education system. But some rough comparisons are possible. Estonia's digital-education transformation,

conducted between roughly 2000 and 2015, consumed approximately one point two billion euros in today's money, for a school population of perhaps one hundred and thirty thousand children — a per-child lifetime investment of approximately seven hundred and fifty thousand rupees. Singapore's equivalent investments, harder to disaggregate from the general education budget, are estimated at several billion US dollars across a comparable period for a school population of roughly half a million — a broadly similar per-child figure. India's per-child investment under the reform, if the envelope is sustained, would be approximately ten thousand rupees across the five-year period for a school population of two hundred and fifty million — roughly seventy-five times less than the Estonian benchmark, even accounting for purchasing-power differences. The comparison does not prove that the Indian envelope is insufficient, because the Indian reform is designed to operate at a different unit cost than the European reforms, with heavier reliance on Indian manufacturing, Indian talent, and the existing infrastructure of Indian public schools. But the comparison does suggest that the envelope is not extravagant by any international standard, and that the more likely risk is under-funding rather than over-funding.

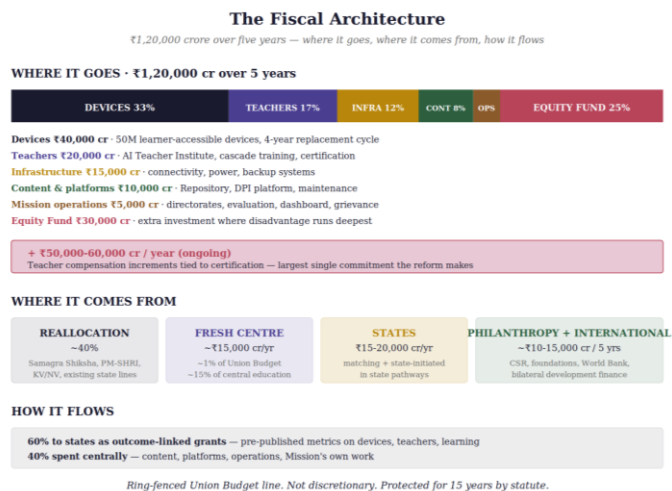


Figure 11.1 | The fiscal architecture — where the money goes, where it comes from, how it flows.

I want to address, briefly, a separate fiscal line that I acknowledged in Chapter Six. The compensation increments tied to teacher certification — the ten percent increment for certified teachers plus the additional five percent for school-level AI-education coordinators — represent an ongoing fiscal commitment of approximately fifty to sixty thousand crore per year at steady state, assuming most of the teaching workforce is eventually certified. This commitment is larger than the entire device procurement budget. It is also, in my judgement, the specific fiscal commitment that most directly determines whether the reform's teacher capacity agenda succeeds.



I. Where the money comes from

The question of sources has specific dimensions that the Cabinet note must address. I will treat them in order.

The first source is reallocation. Approximately forty percent of the reform's investment envelope — perhaps fifty thousand crore of the one hundred and twenty thousand — can be reallocated from existing education-technology line items across multiple programmes. Samagra Shiksha has substantial technology-in-education allocations that, with coordinated planning, can be directed toward the reform's requirements. PM-SHRI schools have technology components that can be integrated with the mission's deployment. The Kendriya Vidyalaya and Navodaya Vidyalaya systems have their own device and connectivity budgets. The state governments have analogous allocations across their own budget lines. The reallocation is not a matter of cutting these programmes; it is a matter of coordinating their technology spending under a common architecture so that the aggregate deployment serves a coherent national objective.

The reallocation is, in practice, substantially harder than its description suggests, and I do not want to underplay the

institutional difficulty. Each of the programmes from which reallocation is proposed has its own constituency, its own administrative structure, its own delivery arrangements, and its own reporting cycles. The programme managers are, for good reasons, attached to their existing architectures and resistant to reallocation that they experience as a reduction in their autonomy. A Samagra Shiksha allocation that is integrated into the mission's deployment is, from the programme manager's perspective, an allocation over which she has lost specific discretion, even if the aggregate benefit to her programme's beneficiaries is equivalent or greater. The coordination therefore requires specific political authority above the level of the programme managers — ultimately, the Cabinet's authority, expressed through specific Cabinet decisions that direct the reallocation — and a sustained negotiation process through which the programme managers' legitimate concerns are addressed while the reallocation is nevertheless accomplished. The negotiation is not a technical matter that can be delegated to a coordination committee. It is political work that requires the attention of the Minister of Education, with the Prime Minister's backing where state-level complications arise. A reform that assumes the reallocation will happen smoothly because it is obviously sensible is a reform that has misunderstood the politics. A reform that plans specifically for the negotiation, with dedicated political attention and a realistic timeline of two to three years for the reallocation's full effect, is a reform that has a chance of accomplishing it.

The second source is fresh central allocation. Approximately fifteen thousand crore per year of genuinely new central money — money that is raised through general taxation or borrowing and that does not displace other priorities at a one-to-one ratio — is required to fund the reform at the pace and quality it needs. This is a substantial commitment, and its political and fiscal cost is real. Fifteen thousand crore is approximately one percent of the total Union Budget at current scale, or approximately fifteen percent of the annual central education budget. Neither ratio is prohibitive.

Both ratios require that the political leadership of the reform argue its case against competing demands.

The third source is state governments. Education is on the Concurrent List, and the reform's full delivery requires substantial state investment alongside central investment. The states together already spend approximately three lakh crore per year on school education, of which a portion is the states' own contribution to Centrally Sponsored Schemes and a portion is state-initiated expenditure. The reform's marginal state-level cost, on top of existing expenditure, is of the order of fifteen to twenty thousand crore per year, distributed across twenty-eight states and eight union territories.

The fourth source, which deserves specific attention, is private and philanthropic co-investment. There is, currently, substantial private and philanthropic interest in the AI-education space in India — from corporate CSR funds, from major Indian philanthropies, from international foundations, and from the impact-investment community. The mission should establish a specific framework through which this interest can be channelled into the reform's priorities, with clear governance, transparent reporting, and protections against the capture risks I named in Chapter Four. A realistic magnitude for this contribution, over the first five years, is of the order of ten to fifteen thousand crore.

I want to say a specific word about the governance of philanthropic and CSR co-investment, because the risks are substantial and the current default arrangements are not adequate to the scale the reform envisages. The first risk is that the co-investors, individually or collectively, acquire disproportionate influence over the mission's decisions, through the specific lever of the funds they provide. An Indian philanthropy that contributes five hundred crore to the reform over three years, even if the contribution is formally unrestricted, acquires in practice a relationship with the mission that gives its views and preferences a weight that an equally valid but non-contributing voice would

not have. The second risk is that the co-investment is directed toward the components of the reform that are most attractive to the co-investors — typically the visible, measurable, media-friendly components — while the less visible but more consequential components remain under-funded. The third risk is that the co-investment introduces specific conditionalities that, individually, are reasonable but that, in aggregate, constrain the mission's strategic discretion to a degree that compromises its integrity. The governance arrangements must address each of these risks. The framework I would propose is that co-investments of above a specified threshold — perhaps one hundred crore — must be made through a blind trust structure, in which the funds are pooled with other contributions and the allocation to specific uses is made by the mission's Evaluation and Equity Directorate according to the reform's published priorities, with the co-investor receiving an annual report of how the pooled funds were used but having no operational influence over the allocation. Smaller contributions can be accepted on more flexible terms, but with transparent disclosure of all contributions above a smaller threshold — perhaps one crore — and with public reporting of the contributors and the uses to which their contributions were directed. The framework is stricter than the current Indian default for philanthropic engagement with public institutions, and it will produce some reduction in the total volume of co-investment that is attracted. The reduction is a price worth paying for the integrity of the mission's independence, and the framework should be codified in the Act rather than left to the mission's discretion.

The fifth source, which I mention for completeness without recommending it as a primary instrument, is international development finance. The World Bank, the Asian Development Bank, and several bilateral development partners have existing programmes in Indian education, some of which could be adapted to support specific components of the reform. International finance should be treated cautiously — it comes with conditions that sometimes constrain sovereign decision-making — but for specific

technical assistance, specific equity-focused interventions, and specific evaluation work, it is a defensible supplementary source.

A final observation on the sources question concerns the specific relationship between the reform and the existing Central-sector programmes it should integrate with rather than compete against. Indian policy has a characteristic pattern, sometimes called new-scheme syndrome, in which each major reform creates a new implementation vehicle rather than operating through the existing pipes. The new vehicle duplicates administrative overhead, fragments the reporting machinery the states are asked to maintain, and consumes political attention that should have been directed at the substantive work. The AI-education reform should resist this pattern deliberately. The Mission is a new coordinating institution, justified by the coordination problem the existing bodies cannot solve among themselves; but the delivery of its programmes should, wherever possible, operate through Samagra Shiksha's existing Central-sector architecture, through PM-SHRI's existing flagship-school framework, through DIKSHA's existing digital platform, and through Bhashini's existing language infrastructure. The Mission coordinates and sets standards; the existing programmes execute, under the Mission's specifications. The integration is administratively more complex than standing up a parallel delivery apparatus, but it is substantially more efficient, substantially more likely to survive political transitions, and substantially more respectful of the institutional equities of the programmes that already exist. The Cabinet note establishing the Mission should codify this integration principle explicitly, so that the default assumption — in every subsequent implementation decision — is integration rather than replication.



II. How the money flows

The specific arrangements for fiscal flow are at least as consequential as the total quantum. Indian public administration has tried, with mixed success, a variety of fiscal transfer mechanisms for education. Each has specific characteristics that make it suitable or unsuitable for the reform's requirements.

Centrally Sponsored Schemes, which have been the dominant model under Samagra Shiksha, have the advantage of leveraging state capacity and of incentivising state investment through matching requirements. They have the disadvantage of requiring political alignment between the Centre and the states for timely disbursement. They also tend, over time, to become inflexible.

Direct grants to autonomous bodies, as with UIDAI or NHA, work well for centralised delivery but do not translate automatically into state-level execution. Their strength is the autonomy they confer on the implementing body; their weakness is the separation they can create between the central implementer and the state machinery whose cooperation the implementation ultimately requires.

Finance Commission devolution gives states broad discretion over their use of centrally-transferred resources, which suits the federal structure well but makes it difficult to tie funds to specific outcomes.

A variation on this mechanism, which has been experimented with under the Fifteenth Finance Commission's arrangements, is the performance-linked grant component — a specific portion of the devolved resources whose release is conditional on the state meeting agreed performance indicators. The experiment has had mixed results. In some states the performance linkage has produced genuine improvements in reporting and in programme operation. In others it has produced optimisation for the specific indicators at the expense of broader programme quality, or outright manipulation of the reporting through which the indicators are assessed. The AI-

education reform's grant architecture should learn from both patterns. Performance linkage is valuable when the indicators are carefully chosen, difficult to game, and used as inputs to a substantive dialogue rather than as triggers for automatic financial consequences. It is harmful when the indicators are chosen for administrative convenience, easily manipulated, and tied to fiscal transfers in ways that produce perverse incentives. The specific indicator framework that I advocated earlier in the chapter — focused on device deployment, teacher certification, and disaggregated learning outcomes, rather than on aggregate usage metrics — is the framework that best fits the productive use of performance linkage, and the mission's fiscal architecture should embed this framework rather than defaulting to the weaker alternatives.

Performance-linked grants, as used in the foundational literacy and numeracy mission Nipun Bharat, tie central fiscal transfers to state performance against specified metrics. The mechanism has the advantage of aligning state incentives with the reform's outcomes, and the disadvantage of requiring substantial measurement capacity to implement credibly.

The mechanism I propose for the AI-education reform is a hybrid, drawing on each of these precedents. The mission receives its own budget line from the Centre. Sixty percent of that budget is disbursed to states as outcome-linked grants, with pre-published metrics on device deployment, teacher certification, and learning outcomes. The remaining forty percent is spent centrally on content, platforms, and the mission's own operations. Within the outcome-linked grant structure, the mission retains a specific discretion to intensify support for states that are performing below expectation and to undertake direct-to-district interventions where state machinery is underperforming systematically.

I want to be specific about what the outcome metrics should, and should not, include, because the design of the metrics is the specific mechanism through which the incentive structure of the reform is shaped. The metrics that should be included, at the state

level, are: the proportion of target government and government-aided schools that have received the infrastructure deployment, the proportion of target teachers who have received the first-level certification, the proportion of students at each target grade who are demonstrably achieving the learning outcomes the curriculum specifies, the equity-disaggregated versions of each of the above, and specific measures of the delivery-chain quality — response times on grievances, frequency of Rapid Response Unit interventions required, the compliance of vendor partnerships with the reform's standards. The metrics that should not be included, or should be included only with specific caution, are the aggregate AI-tool usage metrics that the tools themselves generate — engagement hours, session counts, platform active users. These aggregate metrics are easily produced, easily gamed, and not well-correlated with the learning outcomes the reform cares about. A state whose fiscal transfers are linked substantially to these metrics will, within a year, be optimising for them rather than for learning, which is the specific pathology I named earlier as the metric-collapse failure. The choice of metrics is therefore not a technical afterthought but a primary design question, and the mission's early work should include substantial consultation — with measurement specialists, with state implementers, with the evaluation research community — on the specific metric framework that best serves the reform's genuine purposes.

The Equity Fund I specified in Chapter Seven has a different fiscal architecture. Its thirty thousand crore envelope over five years is disbursed not against generic outcome metrics but against specific equity-focused commitments — the disadvantaged districts identified, the under-resourced schools targeted, the linguistic communities served, the accessibility investments made. The Fund's disbursement is managed by a dedicated unit within the Evaluation and Equity Directorate, with its decisions published in full and its annual report including specific district-level allocations.



III. The political economy of the fiscal commitment

I want to close this chapter with an honest observation about the political economy of making a fiscal commitment of the scale the reform requires. The commitment is politically difficult. It asks the finance ministry to ring-fence a substantial annual allocation for a fifteen-year horizon, to accept reduced discretion over that allocation through the mission's structure, and to forgo the allocation's availability for competing priorities that will, inevitably, press for attention during the mission's life.

The political sponsor of the reform — the Prime Minister, or a senior minister with the Prime Minister's backing — must be willing to defend the allocation against these competing pressures not once but continuously, across budget cycles and political transitions. The defence is winnable, but it requires specific argumentation. The generic argument — that education is important and AI is the future — is insufficient, because every claim on the fiscal envelope is supported by some generic argument for its importance. The specific argument is that the reform's outcome is civilisational, that the cost of not doing it is larger than the cost of doing it, that the window in which it can be done is limited by the pace of the technology's development and the aging of the teacher workforce, and that the specific institutional arrangements the mission embodies protect the allocation from the inefficient deployment that comparable commitments in other sectors have sometimes suffered.

I want to unpack the "cost of not doing it" argument, because it is the argument most often gestured at and least often made concretely. Suppose the reform is not undertaken, or is undertaken at a fraction of the envelope I have specified. What, specifically, is the cost? In the short term — the three to five year horizon — the cost is largely invisible, which is among the reasons the fiscal commitment is politically difficult to secure. The children who are in school today will continue through school, pass through their examinations, and enter adulthood, with or without the reform.

Their life trajectories will differ from what they would have been under the reform, but the differences will emerge only in the adult years, when many of them are competing for employment in an economy increasingly shaped by AI-capable workers from other countries. In the medium term — the ten to fifteen year horizon — the cost becomes more visible. India's workforce competitiveness in a wide range of sectors depends substantially on the cognitive preparation of the workers entering it, and a country that has under-invested in AI literacy at school level will, a decade later, be drawing on a workforce that is systematically less prepared than the workforces of comparable economies. The employment consequences, the wage consequences, the sectoral-competitiveness consequences accumulate. In the long term — the twenty to thirty year horizon — the cost becomes civilisational. A generation of Indians who have not had the preparation the reform would have provided is a generation whose capacity to participate in, and shape, the AI-driven economy of the 2040s is diminished, with the diminishment passing to their children and, eventually, to the country's aggregate position in the world. The argument is not that these costs are certain. Some of them depend on how the technology itself develops, which is uncertain. But the expected value of the costs, weighted by reasonable probabilities, exceeds the cost of the reform by a substantial margin, and this is the specific form in which the fiscal case for the commitment must be made in the Cabinet note and the budget negotiations.

The window argument deserves equally specific treatment. Why is the window limited? Two reasons. The first is the teacher cohort argument. The teachers who will deliver the reform in its early years are the teachers who are currently in service, and their average age is such that a substantial proportion of them will have retired before a delayed reform reaches them. A reform that begins in 2027 can build capacity among teachers who will still be in classrooms in 2040. A reform that begins in 2033 has lost access to a substantial share of the veteran workforce whose cumulative experience is one of the country's most valuable educational assets.

The second is the technology-trajectory argument. The pace of AI development is such that the technical parameters of the reform – the capabilities of the AI tutors, the nature of the classroom integration, the specific pedagogical opportunities the technology enables – are evolving year by year. A reform that commits in 2027 can design itself around the capabilities of that moment, with adjustments as the technology develops. A reform that commits in 2033 is designing for a technology environment that is substantially different from the one the earlier design would have addressed, and is doing so after six years in which the unprepared students have encountered the evolving technology without the curricular framework the reform would have provided. The window is not closing irreversibly, but the cost of delay accumulates in ways that make an earlier commitment substantially cheaper and more effective than a later one. The argument should be made in exactly these terms, with the specific numbers behind each component, in the fiscal negotiations that the commitment requires.

The political economy of the commitment also has a dimension that deserves specific comment. A reform whose fiscal commitment is visibly branded as the initiative of a single governing party will be vulnerable to reversal when the political balance changes. The mission's framing, from its first announcement, should be non-partisan – a national reform whose success is a matter of civilisational consequence, in which the principal opposition parties are invited to take substantive ownership roles through the Governing Council's composition and through the mission's transparency commitments. The non-partisan framing is not naivety. It is the same discipline that has protected Aadhaar across three changes of government.



Three decisions for the reader

1. If you are in government, ring-fence the allocation in the next Union Budget.

The Union Budget following the mission's announcement should include a specific, ring-fenced allocation for the mission — not a sub-allocation within the Ministry of Education's envelope, but a dedicated line with its own scheme code. The specific act of ring-fencing is the test of the government's fiscal seriousness about the reform, and it is the signal the sector reads to assess whether the reform is real or ornamental.

2. If you are in state administration, model the reform's fiscal implications for your state.

Within six months of the mission's establishment, your state's finance and education departments should produce a detailed fiscal model of the reform's implications for your state — the central transfers you expect to receive under the outcome-linked grant structure, the state matching requirements, the additional state-initiated investments the reform's success requires, and the fiscal space within which these costs must be accommodated. The model is the analytical foundation for the state's implementation pathway.

3. If you are a philanthropist or institutional funder, engage the mission's co-investment framework early.

The specific design of the co-investment framework matters substantially for whether private and philanthropic contributions accelerate the reform or fragment it. Philanthropists and institutional funders who engage with the mission's design in its early months, who help shape the framework's specific provisions,

and who align their own commitments with the framework's priorities are funders whose contributions will be most consequential.



END OF CHAPTER ELEVEN

CHAPTER TWELVE

The Delivery Chain — From Shastri Bhavan to the Classroom



Why most reforms fail at the middle of the chain, and the design disciplines that shorten it.



Between a decision taken in Delhi and a Monday morning class in Barabanki, there are seven administrative layers, each with its own logic, its own incentives, and its own specific friction. This chapter is about how to navigate them. The reform cannot succeed without a deliberate strategy for doing so, and no chapter of this book would be honest without confronting what has killed so many earlier reforms at the same middle layers.

I want to name, because the observation is uncomfortable and therefore often omitted from Indian policy writing, that the specific pattern of middle-layer failure I am about to describe is not a pattern that afflicts a few outlier states while most of the country functions smoothly. It afflicts the great majority of Indian states to varying degrees, including states whose administrative reputation is comparatively strong. A principal secretary in a well-regarded state, speaking to me candidly in preparation for this chapter, estimated that perhaps sixty percent of the substantive content of a Central directive reaches the classroom in a form recognisable to its authors, and perhaps forty percent either does not reach or reaches in distorted form. Her own state is considered, on most comparative metrics, among the stronger third of Indian states. I do not cite her by name because her candour was given in

professional confidence. I cite the estimate because it is, in my experience, within the range that other senior officers, in other states, have privately offered when pressed on the same question. A reform that designs itself on the assumption that Central directives reach classrooms with high fidelity is a reform that has misunderstood the system it operates within. A reform that designs for the sixty-percent fidelity, with specific disciplines to raise the figure and specific mechanisms to detect and respond to the shortfall, is a reform with a chance of producing the outcomes it claims.

The seven layers, which I will recall for the reader, are these. At the top sits the Centre — the Mission, the Ministry of Education, the PMO. Below that sits the state government — the Department of Basic Education or its equivalent, headed administratively by the Principal Secretary and politically by the state's education minister. Below that sits the State Council of Educational Research and Training. Below that sits the Directorate of State Educational Training, which in several states operates through regional offices. Below that sits the district-level education administration under the District Education Officer. Below that sits the Block Resource Centre, which supervises schools across a block of a district. And at the bottom, where the reform must actually succeed, sits the school — the headmaster and the teachers.

I want to add, because the school layer is often treated as a passive recipient of the chain's decisions rather than as a substantive participant in it, that the headmaster is a specific actor whose own agency within the chain deserves explicit attention. A headmaster in the Indian public school system has, in most states, limited formal authority over substantive matters — she does not hire or fire her teachers, does not control her budget beyond very small operational allocations, does not set her school's curriculum or its assessment framework. Her formal role is largely administrative. Her informal role, in the school that functions well, is substantially larger — she sets the tone of the school's

professional culture, she protects her teachers from the external pressures that would otherwise disrupt their work, she represents the school's specific circumstances upward through the chain when the chain's default assumptions do not fit her school's reality, and she negotiates with parents, community members, and visiting officials on her school's behalf. The informal role is the specific leverage point at which the reform's delivery can be strengthened or weakened at the school level. A reform that equips headmasters with the authority, the information, and the professional standing to exercise this informal role effectively is a reform that has substantially reinforced its delivery chain's final link. A reform that treats headmasters as administrative functionaries, subordinate to the chain's upward authority, is a reform that has weakened its own final link in the specific way that Indian educational administration has, for decades, tended to weaken it. The specific disciplines that strengthen the headmaster's role – delegation of specific authorities from the chain's higher layers, explicit protection of the headmaster's professional judgement against bureaucratic second-guessing, investment in headmaster training that addresses the leadership dimensions of the role rather than only its compliance dimensions – are disciplines the mission should embed in its own engagement with the school layer from the first year of its operation.

Each layer has its own information system, its own reporting schedule, its own performance incentives, and its own survival dynamics. A memo issued from Shastri Bhavan in March does not simply appear on the noticeboard of the Barabanki school in April. It is interpreted by the state department, often with local modifications. It is rephrased by the SCERT into training material. It is disseminated through the Directorate and DIET networks into district-level training programmes. It is communicated from the District Education Officer to the Block Resource Centre. It finally reaches the headmaster through a block-level meeting at which several other directives are competing for her attention. Each translation introduces small but cumulative distortions. Each delay

consumes weeks or months. By the time the memo's substantive intent has travelled from the top of the chain to the bottom, weeks have become months, the original formulation has been modified through six interpretive steps, and the teacher who finally receives the instruction may be implementing something that bears limited resemblance to the Centre's original specification.

This is not a failure of intention at any specific layer. It is the systemic friction of a chain that was not, in most cases, designed with direct-to-classroom delivery in mind. The administrative architecture of Indian education evolved to handle textbook distribution, examination administration, and the occasional policy directive; it was not designed to handle the frequent, precise, technology-supported communication that a reform of this character requires.

I want to illustrate the friction with a specific example, because abstraction continues to conceal it. In 2021, an earlier education reform under the Nipun Bharat mission issued a detailed guideline on the use of specific diagnostic assessments for foundational literacy and numeracy. The guideline was issued in May. By August, when I visited three districts in Uttar Pradesh to understand its implementation, the situation on the ground was the following. In Lucknow district, the District Education Officer had received the guideline in June, had forwarded it to the Block Education Officers in early July, and the Block Education Officers had communicated it to headmasters at the block-level meeting in late July. The headmasters I spoke to had, in most cases, received a two-page summary in Hindi that differed in several specifics from the Centre's original formulation. The teachers, who are the people the guideline was ultimately addressed to, had in most cases not seen the guideline in any form; they had received verbal instructions from their headmasters to begin a new assessment procedure, without the pedagogical context the original guideline had carefully provided. In Barabanki district, the chain was the same but slower — the teachers had not yet received even the

verbal instructions. In Balrampur district, the guideline had, at the time of my visit, not been communicated to headmasters at all; the DEO was holding it pending a clarification from the state directorate. Three districts within a single state. Three different states of implementation. A guideline issued in May was, in August, at various stages of its translation into classroom action, with the translation itself already having introduced specific distortions from the original intent. This is the friction. It is not a failure of any one officer. It is the specific systemic pattern that the reform's direct-to-classroom disciplines must address.

Three design disciplines can shorten the chain. I want to elaborate on each in this chapter, because the specific operational practice each requires is the difference between a reform that works and one that does not.

The three disciplines I will describe are not novel in the literature of public administration. They have been attempted, with varying success, in other Indian reforms and in reforms across comparable countries. What is specific to the AI-education reform is the combination — the particular way in which the three disciplines must operate together to address the particular friction that the education chain generates. An institution head reading the chapter will recognise versions of each discipline from her own experience of earlier reforms. The specificity of the argument is in how the three must be adapted to the AI-education context, and how their joint operation produces effects that the disciplines deployed individually do not produce.



Figure 12.1 | The seven-layer delivery chain — and the three disciplines that shorten it.



I. Direct-to-school digital delivery

The first discipline is direct-to-school digital delivery. The mission, or more likely the mission's extension of the existing DIKSHA platform, should build a direct channel through which its materials — content, training, monitoring instruments, operational communications — reach teachers' and headmasters' devices without passing through the intermediate administrative layers. The administrative layers still matter. They translate, they adapt, they enforce. But the content and the training and the communications reach the teacher directly and simultaneously nationwide.

The specific operational implication is substantial. When the mission updates a Class 6 data-literacy module, the update should appear on every Class 6 teacher's device across the country within days, not months. When the mission identifies a common

misconception in how teachers are handling the machine-as-fallible theme at the foundational stage, the corrective training material should be pushed to affected teachers immediately. When a teacher in Kerala discovers a classroom technique that works well, the mission should have a mechanism to surface that technique, verify its effectiveness, and make it available to every teacher who teaches the same grade across the country, within a time horizon measured in weeks rather than years.

The technical infrastructure for this direct channel is not trivial. It requires that every teacher has a device capable of receiving the communications, that the channel operates reliably in the network conditions Indian schools actually experience, that the content is offline-capable for periods when connectivity fails, and that the channel accommodates the twenty-two languages the mission's content must reach. These requirements are substantial, but they are tractable. DIKSHA already operates in most of the required conditions, and its extension to the specific communication functions the mission requires is an achievable technical project within the first year of the mission's operation.

I want to address, at the risk of moving into operational detail that the general reader may not care about, the specific question of how the direct-to-school channel should handle the two-way flow of information, because the channel is valuable only if it supports communication in both directions rather than only from the centre to the classroom. The downward flow — content, training, announcements, operational communications from the mission to the teacher — is the easier half. The upward flow — teacher feedback, classroom observations, specific problems encountered, specific techniques developed — is harder, both technically and organisationally. Technically, the upward flow requires structured formats through which the teacher's reports can be aggregated and analysed, rather than free-text messages that the mission has no capacity to process at scale. Organisationally, the upward flow requires the specific commitment, by the mission, that the reports

will be read, will inform decisions, and will produce visible responses within a specified time. A channel that collects teacher feedback without acknowledging or acting on it will, within a year, stop receiving substantive feedback — teachers will, reasonably, conclude that their time is not well spent on reports that produce no response, and the channel will devolve into a one-way broadcast medium. The specific commitments the mission should make are: every structured report submitted through the channel receives an acknowledgement within forty-eight hours; every issue raised is assigned an owner within seven days; every resolution is communicated back to the reporting teacher within thirty days, or within a longer specified period if the resolution requires more extensive work. The commitments are demanding for the mission's operations. They are also the specific discipline through which the upward channel remains substantive rather than collapsing into formality, and the investment in meeting them is an investment in the reform's operational integrity.

The political and administrative implications of direct-to-school delivery are less comfortable than the technical ones. The intermediate administrative layers, particularly the SCERT, Directorate, and DIET cadres, may experience direct-to-school delivery as a bypass of their authority — and they are, in one specific sense, correct. The direct channel does bypass their role as content intermediaries, substituting the mission's direct reach for their traditional function of adaptation and dissemination. The mission must acknowledge this frankly and must redesign the intermediate layers' functions so that they continue to add value rather than simply mediate the chain.

The redesign has a specific shape that deserves elaboration. The SCERT, under the redesigned model, moves from being an information-relay institution to being a state-specific content-adaptation and capacity-support institution. Its primary work becomes the adaptation of the national curriculum and content to the state's specific linguistic, cultural, and pedagogical context, and

the operation of the master-trainer cadre that delivers the in-person components of the teacher-training programme. The Directorate — where it exists as a separate body from the SCERT — moves from being a dissemination intermediary to being the operational coordinator of state-level implementation, responsible for the specific deployment decisions, the vendor coordination, and the accountability for state-level outcomes. The DIET moves from being a training institution that receives teachers for periodic programmes to being the anchor institution for ongoing teacher support within the district — hosting the peer-learning circles, supporting the school-level coordinators, conducting the classroom observations that feed back into the training programme's continuous improvement. The Block Resource Centre, which historically has had the weakest role in the chain, acquires new importance as the site of frequent small-group teacher engagement and as the first point of contact for school-level problems that cannot be resolved at the school itself. Each of these redesigns requires a specific capacity-building programme for the institutions in question, because the skills that the redesigned roles require are not identical to the skills the current roles have cultivated. The capacity-building programmes should be developed by the mission in consultation with the institutions themselves, should be funded from the reform's envelope, and should be sequenced so that each institution is prepared for its redesigned role before the redesign takes effect operationally. The sequencing matters. A redesign that takes effect before the institution is prepared for it is a redesign that produces, in the short term, a capability vacuum that the chain cannot function without.

The redesign is not a demotion of the intermediate layers. It is a refocusing. The layers move from mediating information flow, which the direct channel now handles, to performing the specifically local functions that the direct channel cannot handle — the state-specific adaptation, the in-person training, the local support, the community engagement. These functions are genuinely important. They are also, in most cases, not what the

intermediate layers currently do well, which means the redesign is accompanied by specific capacity development for the layers themselves.



II. Named mission coordinators

The second discipline is the naming of specific individuals, at the state and district levels, as the mission's coordinators in their geography. This is not a new bureaucratic position. It is the designation of a specific officer from the existing cadres — typically a senior SCERT officer at the state level, and a senior DIET officer at the district level — as the named owner of the reform within her territory.

The practical importance of the naming is substantial, and it is not obvious until one has watched what happens in its absence. When a reform is addressed to the state government or the district administration, the reform is addressed to nobody in particular. The specific officer receiving the communication is one of many, her responsibility for the reform is one of many, and the mission has no specific person to contact when something needs attention. The mission calls the state, is told to call the department, is told to call the joint secretary, who is told to call the section officer, who has been reassigned. The chain of responsibility is diffused, and the reform is, in effect, under nobody's specific supervision at each layer.

When a specific officer is named, the dynamic changes. She is personally responsible for the reform's performance in her state or district. Her career is, in a specific sense, tied to the reform's outcomes in her geography. She attends monthly mission convenings. She holds her state or district colleagues accountable for reform milestones. She is the person the mission knows by name and can call directly when something is going wrong. The pattern of naming specific officers with specific accountabilities is

the distinguishing feature of every successful large-scale reform in recent Indian administrative history — the UIDAI enrolment drive, the Swachh Bharat construction programme, the Nipun Bharat foundational literacy push.

The naming has specific design requirements. The officers must be genuinely senior — not junior officers given a coordination label, but people of substance whose authority within their own geographies is real. They must be given the time to do the reform work seriously — not the reform added to a portfolio of ten other responsibilities, but a dedicated substantial allocation of their time. They must be given the authority to convene their state or district colleagues, to escalate issues to the mission, and to take specific decisions within agreed frameworks without having to revert to the Centre for every adjustment. They must be given the support — the analytical staff, the access to the mission's systems, the travel and meeting budgets — that the role requires. And they must be given some specific career benefit for the investment the role represents.

I want to expand on the career-benefit point, because it is the specific incentive mechanism through which the naming of coordinators produces sustained engagement rather than symbolic compliance. The Indian civil service, at both the Central and state levels, operates with career-progression structures that do not, in most cases, specifically reward engagement with reform initiatives. An officer who invests three years in the intensive work of coordinating a reform's implementation in her state is, in terms of her subsequent career, often no better positioned than an officer who occupied a comparable posting during the same period without the reform-specific engagement. The specific asymmetry between the additional effort the role requires and the absence of additional career reward is a specific disincentive that, across the cadre, produces a pattern of officers seeking to minimise rather than maximise their engagement with reform work. The reform should disrupt this pattern. The specific mechanisms available

include: a formal recognition of the coordinator role in the officer's annual confidential report, with the recognition linked to specific outcomes the coordinator has produced rather than to generic commendation; a specific weight given to the coordinator role in the promotion decisions that follow the posting, negotiated with the state and Central services' personnel authorities; the availability of specific training and exposure opportunities – international study visits, enrolment in premier educational programmes, secondments to the mission's headquarters – that are reserved for officers who have served effectively in coordinator roles; and a named-coordinator professional network, operated by the mission, that creates a specific cohort of officers with shared experience and mutual support across their subsequent careers. The mechanisms are substantive but not extravagant, and they are the specific form in which the reform communicates to the cadre that coordinator roles are career-building rather than career-neutral.



III. The Rapid Response Unit

The third discipline is the establishment of a Rapid Response Unit within the mission – a small team, perhaps twenty officers, drawn from deputation and private-sector hires, whose function is to travel to failing districts, diagnose bottlenecks, and either fix them in partnership with the local administration or escalate to the state government for political resolution.

Every large delivery programme produces, within its first year, a tail of districts, schools, or states where the standard process has failed to produce results. The standard response in Indian public administration is to wait for these to surface through annual evaluations, debate them in committees, and address them in the next programme cycle. This response is too slow for a reform of the AI-education kind, because the children in the failing districts are

falling further behind every term that the failure persists. A rapid response capacity is the specific mechanism through which the mission can intervene in real time.

The Unit's operational model is specific. When performance metrics show that a district is falling significantly behind its expected trajectory, or when grievances from the district exceed a threshold, or when a named coordinator flags a specific problem that she cannot resolve, the Unit is dispatched. A team of two to three officers travels to the district within a week. They spend seven to ten days on the ground — visiting schools, meeting teachers, talking to parents, understanding the specific bottleneck, and working with the local administration on a specific remediation plan. The plan is agreed with the District Education Officer, and, where necessary, escalated to the state government for political support. The Unit leaves behind a specific action plan with timelines and responsibilities, and returns after six months to verify that the remediation has taken effect.

The Unit is the operational equivalent of a fire service. It should not be needed most of the time. When it is needed, it should be fast, respectful, and effective. It should leave behind a local capability rather than create a dependency. Its very existence — the knowledge, across the administrative chain, that failures will be surfaced and addressed rather than allowed to persist — is a substantial discipline, even when the Unit itself is not active in a given geography.

I want to address the delicate question of how the Rapid Response Unit's interventions are conducted politically, because the technical effectiveness of the Unit depends, in practice, on the political legitimacy with which its interventions are perceived. A Central unit dispatched to a state district to diagnose problems in state-operated schools is, whatever its technical mandate, engaging with the political authority of the state government, the professional authority of the state's own officers, and the community authority of the district's local institutions. A Unit that

operates without sensitivity to these authorities will, within a few interventions, have generated the kind of state-level resentment that reduces its subsequent effectiveness across the system. The specific disciplines that protect the Unit's legitimacy are: the intervention is always conducted in consultation with the state government, with the state's mission coordinator participating in the Unit's work from the first day; the Unit's findings are shared with the state government before being incorporated into the mission's reporting, with an opportunity for the state to provide its own response; the Unit's communications with the local administration and community are conducted through the state's officers rather than around them, with the Unit's officers supporting rather than substituting for the local chain of authority; and the Unit's reports, where they identify failures, are specific about the distinction between failures that are attributable to local or state conditions beyond the administration's control and failures that are attributable to specific administrative decisions. The disciplines are demanding. They are also the specific arrangements through which the Unit becomes a respected institutional resource rather than a politically intrusive instrument, and the mission's leadership should model and enforce them from the Unit's first operations.

The Unit's staffing deserves specific comment. Its officers must be people who can travel frequently, who are comfortable in district-level settings, who have the analytical capacity to diagnose bottlenecks in unfamiliar contexts, and who have the interpersonal skills to work respectfully with officers they are, in a sense, evaluating. The combination is rare. It is best achieved through a mix of deputation from the central and state services and lateral hires from consulting, research, and social-sector backgrounds who bring specific analytical capacities.



IV. A broader observation on delivery

I want to close this chapter with a broader observation about delivery that applies beyond the three specific disciplines I have described.

Indian public administration has, over decades of attempting reforms at scale, developed a specific skill in the production of plans, frameworks, guidelines, and directives, and a less developed skill in the actual delivery of the reforms these documents describe. The gap between planning capacity and delivery capacity is not a reflection of incompetence — the officers who work on Indian reforms are, on average, no less capable than their counterparts in comparable countries — but of institutional incentives that reward the production of documents more than they reward the outcomes those documents are intended to achieve.

The specific incentive asymmetry deserves to be made concrete, because it operates at every level of the administrative chain and is often not visible to officers who are inside it. A district education officer is evaluated, in most states, through a reporting cycle that assesses her performance on the documents she has produced — the reports filed, the meetings conducted, the circulars issued, the returns submitted — rather than on the specific learning outcomes her district's children have achieved. The asymmetry is not accidental. Documents can be counted, verified, and filed; outcomes require substantial measurement capacity that most state administrations do not operate. The officer, responding rationally to the incentives she faces, allocates her attention toward the measurable documents and away from the unmeasurable outcomes. Over a career, the allocation becomes habitual. The officer who spent twenty years in a system that rewarded documentation will, when appointed to a role that formally expects outcome-orientation, continue to produce the documents that her accumulated habits know how to produce. The reform's insistence on outcome-orientation is therefore not a matter of issuing a different instruction to the same officers. It is a matter of rebuilding,

across the chain, the specific measurement capacity that makes outcome-orientation operationally possible, and of reshaping the incentives that have produced the current habits. Both are multi-year undertakings. The reform's fiscal envelope should include the specific resources — the measurement infrastructure, the training programmes, the officer-evaluation reforms — that these undertakings require, rather than treating them as administrative matters to be handled by the states' own systems.

A reform that seeks to change this pattern must, from its earliest design, privilege delivery over planning. The mission's calendar should be dominated by delivery reviews rather than by planning exercises. Its reports should be dominated by specific district-level performance data rather than by national-level aggregates that conceal delivery gaps. Its rewards — for its own staff, for its partner institutions, for the officers in the administrative chain — should flow from outcomes achieved rather than from documents produced. The cultural shift this represents is substantial, and it will be resisted, often implicitly rather than explicitly, by the administrative culture the reform operates within.

The resistance takes specific forms that are worth anticipating. The most common form is the documentation-substitute pattern, in which officers, when asked to report on specific delivery outcomes, produce extensive documentation of their planning and process that substitutes for evidence of the outcomes themselves. A district whose schools are not in fact operating the AI-education programme as specified produces, when asked, a report that documents the training programmes conducted, the equipment delivered, the communications issued, the meetings held — a documentation that is, in its own terms, accurate, but that does not address the specific question of whether the classroom experience of the children has changed. The mission's response to this pattern must be specific: reports are evaluated not on their documentation of inputs but on their demonstration of outcomes, with the

outcomes assessed through independent verification that bypasses the chain's own self-reporting. A second form of resistance is the exception-framing pattern, in which failures to deliver are attributed to specific exceptional circumstances — a district with unusual geographic challenges, a state with a particularly difficult political environment, a set of schools affected by a specific local disturbance. Each individual exception is often valid. The aggregate pattern of exceptions, across a national deployment, is not. The mission's response must be to acknowledge specific exceptions where they are genuinely exceptional, and to disregard them where they have become the default framing through which the chain explains its non-delivery. A third form of resistance is the timeline-extension pattern, in which commitments that have not been met on their original schedule are re-committed for later dates, with the re-commitment treated as equivalent in weight to the original. The mission's response must be to treat re-commitments as specific events that require specific justification, with repeat re-commitments triggering escalation to the Governing Council and, where necessary, to the political leadership of the mission. None of these responses is easy. Each requires the mission to hold an uncomfortable conversation with an officer or an institution that is, in other respects, performing its role in good faith. The willingness to hold the uncomfortable conversation, consistently and across the chain, is the specific behavioural discipline that distinguishes a delivery-focused mission from a planning-focused one, and the discipline must be modelled from the top — by the CEO, by the directors, by the Governing Council — if it is to become the mission's culture rather than remaining its aspiration.

The delivery chain I have described in this chapter is, in the end, a chain of people. The quality of the reform's delivery depends on the quality of the officers along the chain, on the quality of the support the mission provides them, on the quality of the relationships that build up between the mission and the field over years of engagement. An investment in the people of the chain is,

in every serious sense, the most consequential investment the mission can make.



Three decisions for the reader

1. If you are in state administration, nominate your mission coordinator with care.

The officer you designate as the mission's coordinator in your state will, over the next five to eight years, be the single most important individual in your state's AI-education journey. Choose her with more care than is typical for similar designations. She should be senior, competent, respected by the cadres she will need to work with, and genuinely interested in the reform. She should have a career horizon that allows her to serve in the role for at least three years and ideally longer. She should be protected, by specific instruction from the Chief Secretary, from being pulled into other urgent matters at the expense of the reform.

2. If you are in the mission's early team, build the direct-to-school channel before the classroom-facing deployments.

The direct-to-school digital channel is foundational infrastructure. Its development is unglamorous — it produces no ribbon-cuttings, no minister visits, no media coverage — but it is the specific infrastructure on which the reform's subsequent deployments depend. The temptation, in the mission's first year, is to prioritise the classroom-facing pilots that are more visible. Resist this temptation.

3. If you are a teacher or a headmaster, use the channels the mission provides.

When the direct-to-school channels become available, use them. Report the problems you encounter. Share the techniques you develop. Respond to the mission's requests for feedback. The channels work only if the teachers and headmasters at the bottom of the chain engage with them substantively rather than treating them as additional administrative burden.



END OF CHAPTER TWELVE

PART IV



The Execution Layer

CHAPTER THIRTEEN

Outcomes and Accountability



What the reform will be measured against, and who will be named as accountable for the measurement.



A reform's accountability architecture is the specific set of arrangements through which, three years into its operation, someone can answer the question of whether it is working. The question sounds simple. The arrangements that permit an honest answer are not. This chapter specifies the arrangements I propose the Mission should adopt, in three parts. The first part specifies the seven national outcomes the reform will be held to, with their metrics, their cadence, and the institutions accountable for each. The second part specifies the three-level metric framework — classroom, district, state — through which the national outcomes are produced from specific data at specific levels of the chain. The third part specifies the dashboard architecture through which the metrics become visible to the political leadership, the policy community, and the citizens whose children the reform is meant to serve.

I place this chapter at the beginning of Part IV, the execution layer, because accountability is logically prior to risk and to phasing. A reform that cannot specify what it is trying to produce cannot credibly specify what would go wrong, and cannot sequence its work across time. The seven outcomes below are the specific things the reform is trying to produce. The chapters that follow — on failure modes and on phasing — rest on them. A reader whose time is short, and who wishes to take the argument of Part IV in its most compressed form, can read this chapter alone

and proceed directly to Chapter Sixteen. She will have the substantive skeleton of the accountability framework. She will lose the discussion of what to do when the skeleton cracks, which is what the next two chapters provide.



I. The seven national outcomes

A reform without specified outcomes is a reform without accountability, and a reform without accountability, in Indian public administration, is a reform that will, within three years, have drifted into whatever shape the path of least institutional resistance takes it. The first operational discipline is therefore the specification of outcomes the reform will be measured against, at specific horizons, with specific metrics, and with specific consequences for the institutions responsible for them. I propose seven national outcomes, each with its own metric framework, each reported publicly on the Mission's dashboard, and each tied to specific fiscal and political consequences if it is not met.

The seven outcomes I would propose, in descending order of weight, are these. First, foundational literacy and numeracy uplift attributable to AI-tutor deployment, measured as the additional proportion of children reaching the National Curriculum Framework's Class 3 and Class 5 benchmarks beyond the counterfactual trajectory. Second, the AI Readiness Index by district, a composite measure combining infrastructure coverage, teacher certification, and demonstrated student competence at each grade. Third, the teacher capacity ratio — the proportion of the grade-appropriate teaching workforce that has completed first-level certification and is actively teaching with the tools. Fourth, the content availability ratio — the proportion of the Repository's target artefact population available at required quality in each of the scheduled languages. Fifth, the equity gap closure rate, measured as the narrowing, or widening, of outcome differentials

between the disadvantaged-district decile and the advantaged-district decile. Sixth, the delivery integrity index, measured through independent classroom observation of whether the Mission's issued guidelines are actually being implemented at the school level. Seventh, the child-safety compliance rate, measured through audit of vendor deployments against the Children's AI Safety Code proposed in Chapter Sixteen.

Each of these outcomes is, by design, measurable without the manipulation that would render the measurement useless. Each is disaggregated by state, by district, by gender, by community, and by disability status, in accordance with the equity discipline I specified in Chapter Seven. Each has a specific target trajectory agreed at the beginning of each three-year cycle and published at the cycle's outset, rather than being negotiated retrospectively. And each has a specific officer — at the Mission, in the state, at the district — named as accountable for the outcome in her territory, with her name attached to the outcome's progress in the Mission's public dashboard.

I want to say a specific word about why I have chosen the seven outcomes I have, and what I have consciously excluded. I have excluded the aggregate usage metrics that AI tools themselves generate — session counts, engagement hours, platform active users — because these metrics are easily produced, easily gamed, and weakly correlated with the learning outcomes the reform actually cares about. A reform whose accountability is substantially shaped by these metrics is a reform that will, within a year, have optimised its deployment for appearances rather than for children. I have excluded, also, the various composite indices that aggregate across dimensions in ways that obscure the specific performance on each. A composite index that combines infrastructure, teacher capacity, and learning outcomes into a single number is a number whose movement tells you little about what is actually happening underneath, and a reform that is measured by such a number will find itself unable to diagnose the specific problems its deployment encounters. The seven outcomes I have chosen are, instead, specific enough to be diagnosable,

disaggregated enough to be equitable, and few enough to be politically legible.

Outcome	Metric	Cadence	Accountable
FLN uplift via AI	Δ % reaching NCF Class 3 / Class 5 benchmarks vs counterfactual, disaggregated	Annual, by district	State Principal Secretary (Education) Mission Curriculum & Content Director
AI Readiness Index (district)	Composite of infrastructure coverage, teacher certification %, student competence at grade	Quarterly	District Education Officer Named State Coordinator
Teacher capacity ratio	% of grade-appropriate workforce certified and teaching with tools	Quarterly	Mission Teacher Capacity Director State SCERT Director
Content availability ratio	% of target Repository artefacts at quality threshold per scheduled language	Quarterly	Mission Curriculum & Content Director NCERT Head
Equity gap closure	Ratio of bottom-decile to top-decile outcome performance, multi-dimensional	Semi-annual	Mission Evaluation & Equity Director
Delivery integrity index	Independent classroom-observation concordance with issued guidelines	Semi-annual	Mission CEO Rapid Response Unit Head
Child-safety compliance rate	% of vendor deployments passing audit against Children's AI Safety Code	Quarterly	Statutory Safety Authority (per Ch 16)

Figure 13.1 | The Seven National Outcomes – metric, cadence, accountable institution.

The table above is, in my proposal, a document the Mission publishes in its first ninety days and revises only through a formal process of public consultation. The officers named in the rightmost column are, specifically, the officers accountable to the political leadership and to the public for the outcome's progress. The accountability is personal rather than institutional. An outcome that is the responsibility of a named officer is an outcome that someone will attend to. An outcome that is the responsibility of an institution in the abstract is an outcome that, in the Indian administrative practice, tends to receive the attention that no one in particular is obliged to give it.



II. The three-level metric framework

The outcomes above are national. They disaggregate into specific metrics at the classroom, the school, the block, the district, the state, and the national level, and the disaggregation matters because an outcome observable only at the national level is an outcome no one along the chain can be held to in real time. The metric framework I propose is set out in Figure 13.2. At the classroom level, the metrics are pedagogical — children's actual competence against the curriculum, the teacher's use of the tools, the classroom environment's functioning. At the school level, they aggregate the classrooms and add the specific school-level variables — infrastructure uptime, teacher certification rate, headmaster-reported bottlenecks. At the block level, the metrics aggregate across the schools and add the peer-learning health that operates at that level. At the district level, the metrics are the composite AI Readiness Index, the equity disaggregation within the district, and the Rapid Response flags. At the state level, the metrics aggregate the districts, track equity movement across the state, and include the quality of the state's own administrative chain. At the national level, the Mission's dashboard aggregates upward from all of these.

Each level's metrics are visible primarily to the officers at that level, with structured visibility into the levels immediately above and below. The headmaster sees her own school's metrics in detail, and sees the block and district aggregates within which her school sits. The District Education Officer sees her district's metrics in detail, sees the block-level roll-ups, and sees the state aggregate within which her district sits. The Principal Secretary sees her state's metrics in detail, sees the district-level roll-ups, and sees the national frame. The Mission CEO sees everything, with the ability to drill down to the specific district or school when a specific issue surfaces. The visibility architecture matters because it determines which officer, at which level, is actually equipped to respond to which deviation. A deviation that is visible only to the CEO is a deviation that cannot be responded to at the speed the reform requires.

Level	What is measured	Reviewed by	Frequency
Classroom	Child competence vs grade outcomes Teacher AI-tool usage log Tutor engagement quality Classroom infrastructure status	School AI-education coordinator Headmaster	Weekly
School	Grade-level competence aggregated % teachers certified on roll Infrastructure uptime %	Headmaster BRC coordinator	Monthly
Block	School-level rollups Peer-learning circle health	BRC coordinator DEO	Monthly

	Headmaster-reported bottlenecks		
District	District AI Readiness Index Equity disaggregation within district Rapid-Response flags	DEO Named State Coordinator	Monthly
State	All districts rolled up Equity gap movement Grant-linked outcome tracking	Principal Secretary CM's office	Quarterly
National	All seven national outcomes Cross-state comparison Political dashboard for PMO/CMs	Mission CEO Governing Council	Quarterly (Council) Monthly (CEO)

Figure 13.2 | The Six-Level Metric Framework – classroom, school, block, district, state, national.

A word about the specific cadence, because the cadence is the feature of the framework most easily diluted in practice. Weekly classroom review is demanding. It is also the frequency at which the teacher's own classroom practice is naturally reviewed, and the AI-education metrics should be integrated into the teacher's existing weekly reflection rather than added as a separate burden. Monthly school, block, and district review corresponds to the cadence of existing administrative reporting, with the AI-education metrics becoming part of the monthly review the headmaster, the BRC coordinator, and the DEO already conduct. Quarterly state and national review is the cadence at which political attention can realistically be sustained — more frequent than quarterly produces review fatigue at the political level, less frequent than quarterly allows problems to accumulate. The specific cadence is therefore not a technical choice but a specific

calibration to the rhythms at which the Indian administrative system already operates. A reform that imposes cadences out of step with existing rhythms is a reform that will, within a year, have been quietly absorbed into the existing rhythms with its specific cadences forgotten.



III. The dashboard

A word about dashboards, because the dashboard is the specific artefact through which the metric framework becomes operational rather than aspirational. The Mission's public dashboard should, in my proposal, be a single web property with a specific architecture. A top-level view displays the seven national outcomes with their current values and targets. A state-level drill-down displays the same outcomes for any chosen state, with the state's named coordinator's name and photograph attached. A district-level drill-down displays the district-level metrics with the DEO's name and photograph attached. A school-level drill-down displays the school's own metrics with the headmaster's name attached. The drill-downs are navigable by any citizen with a web browser, without authentication, subject only to the specific privacy protections that protect individual children's data from aggregation into identifiable form.

The specific design choice — that the dashboard attaches names to numbers at every level — is what distinguishes a functioning accountability instrument from a decorative one. A metric without a name attached is a metric no one owns. A metric with a specific person's name attached is a metric that will, within weeks of its first adverse reading, have produced the phone call, the site visit, the corrective action that a nameless metric never does. The specification is uncomfortable for some officers who would prefer the protection of institutional anonymity. The protection they prefer is the specific protection the reform must

withhold, and the withholding is what makes the accountability framework work.

The dashboard's technical architecture deserves attention. It should be built on the same DIKSHA extension that supports the Mission's direct-to-school digital channel, leveraging the platform's existing authentication and data infrastructure rather than standing up parallel systems. Its data should flow through automated pipelines from the existing administrative reporting systems wherever possible, with manual data entry required only for the specific classroom-observation metrics that cannot be captured automatically. Its update frequency should match the metric framework's specified cadence — weekly for classroom data, monthly for block and district data, quarterly for state and national. Its public accessibility should be subject only to the privacy protections that protect individual identifiability, not to the access restrictions that the administrative system will otherwise default to imposing. A dashboard that is public is a dashboard that exercises the accountability discipline. A dashboard that is internal is a dashboard that has become another administrative reporting system, and the Indian administrative system does not need another reporting system; it needs an accountability instrument.

The dashboard's political use is the specific use for which it is most valuable. A Prime Minister who receives the dashboard, in working form, once a quarter, with her attention directed to the specific outcomes that have fallen behind their trajectories, is a Prime Minister equipped to hold her Cabinet to the reform's commitments. A Chief Minister who receives the state-level drill-down, showing her state's position relative to the national trajectory and to peer states, is a Chief Minister equipped to make the specific decisions that will improve her state's position. A Principal Secretary who receives the district-level drill-down is equipped to convene the specific DEO whose district is falling behind and to make the specific intervention the situation requires. The dashboard is, in each of these uses, a substitute for the

narrative reports that the Indian administrative system has traditionally produced — reports that are long, that are filled with process details, and that rarely permit the reader to extract, in the time the political desk has available, the specific issues that require the political desk's attention. The dashboard is what gives the political desk its time back. It is also what makes the political desk's use of its time visible to the specific officers whose work the political desk is reviewing, and the visibility is what produces the behavioural discipline throughout the chain.



Three decisions for the reader

1. If you are in the Mission's early leadership, publish Figure 13.1 in the first ninety days.

The specific act of publishing the outcomes table, with the accountable institutions named, is the specific commitment that converts the reform from rhetoric into accountability. Publish it on the Mission's website. Attach the names of the specific officers, with their photographs and contact information. Accept the criticism that follows. Revise, if the criticism is well-founded. Republish. The public availability of the outcomes, with names attached, is itself a disciplining instrument, and the earlier the instrument is engaged, the stronger its effect through the life of the Mission.

2. If you are a Chief Secretary, commission your state's implementation of the three-level metric framework within six months.

Your state will need to build the specific data pipelines that produce the classroom, school, block, and district metrics at the cadences the framework specifies. The build is substantial — it is not a spreadsheet exercise — and it will require the specific investment of your state's IT capacity, your SCERT's methodological expertise, and your district administrations'

cooperation. Begin within six months. Complete within eighteen. Publish your state's dashboard, accessible to citizens, as soon as the pipelines are operational rather than waiting for perfect data. The value of the publicly-visible dashboard exceeds the cost of its residual imperfections, and the imperfections will be addressed faster under public scrutiny than under internal review.

3. If you are a citizen, use the dashboard.

When the dashboard is public, look at your own district. Look at your child's school, if the school-level drill-down is available. Compare what the dashboard reports to what you observe in your child's actual educational experience. Where the two diverge, raise the divergence — with the school, with the block officer, with the DEO, with your elected representative. The dashboard's accountability function depends on citizens using it. Citizens who use it hold the officers at each level to the commitments the dashboard records. Citizens who ignore it allow the officers to report whatever the reporting systems permit. Your use of the dashboard is, in a specific and under-acknowledged sense, the final layer of the accountability architecture the reform is building.



END OF CHAPTER THIRTEEN

CHAPTER FOURTEEN

Risks and Failure Modes



The ten specific ways this reform will fail, and the early warnings and pre-emptive mitigations for each.



I want to do something uncomfortable now, which is to name, in advance, the specific ways in which this reform will fail — not to predict its failure, but to surface the failure modes the reform must guard against. Every reform of comparable ambition in Indian history has failed in recognisable ways, and the reforms that have succeeded have done so in substantial part because their designers and implementers anticipated the failure modes and built specific mitigations into the design. The reader who has completed Chapter Thirteen's outcomes framework will recognise that an accountability architecture without a failure-mode framework is an incomplete accountability architecture — a system that specifies what it wants but not what it fears, and that is therefore unprepared when what it fears begins to occur. This chapter specifies what the Mission should fear.

I have organised the chapter around ten specific failure modes, ordered by the product of their probability and their severity. Each is paired with an early-warning indicator — the specific signal that would suggest the failure mode is beginning to manifest — and a pre-emptive mitigation that can be activated before the failure becomes irreversible. The ten are not exhaustive. A reform of this scale will encounter specific difficulties this chapter cannot anticipate, and the list should be revised annually by the Mission's Evaluation and Equity Directorate in light of the specific issues that

emerge. What the list does is establish the institutional posture — the specific expectation that the reform will, probably, fail in some recognisable way, and the specific commitment to identify the failure early and to respond to it before it becomes irreversible. A reform that operates from this posture is substantially more robust, in my observation, than a reform that operates from the posture of confident success. The confident posture is easier to maintain at ribbon-cutting ceremonies. The self-critical posture is the one that produces reforms that survive.



I. Tech-first deployment

The first failure mode is the pattern in which the reform's visible early outputs are devices, platforms, and vendor announcements rather than teacher capacity and content. It is the failure mode most likely to occur because it is the failure mode that produces the most politically attractive early milestones. Devices can be counted. Platforms can be demonstrated. Vendor announcements can be photographed. Teacher capacity, by contrast, is invisible until it is operational; content is invisible until it is used; and the failure mode occurs when the political pressure for visible early milestones displaces the substantive early work on the less visible components. The reform that fails through this mode is a reform that has, by the end of year one, deployed a substantial number of devices into classrooms without the teachers prepared to use them or the content to animate them, and that is then unable to demonstrate learning outcomes in year three. The devices sit unused or are used poorly; the children receive the appearance of the reform rather than its substance; and the political sponsor who was rewarded for the year-one visible milestones is the political sponsor who is held responsible for the year-three outcome failure.

The early-warning indicator is specific: a Q1 device-procurement allocation that exceeds the Q1 teacher-training allocation, or vendor contracts signed before the corresponding teacher-training commitments are secured. The pre-emptive mitigation is the budget discipline I specified in Chapter Eleven — the rule that no device procurement is approved until the corresponding teacher-training commitment is secured in the same fiscal year, codified as a specific condition in the Mission's fiscal framework. The rule sounds technical. It is, in fact, the single most consequential protection against this failure mode, because it forces the Mission's political sponsor to sequence her visible milestones in the order the reform's substance requires rather than in the order political pressure would otherwise produce.



II. Teacher resistance

The second failure mode is produced by the combination of inadequate consultation, inadequate compensation, and the professional fear that the tools will be used to diminish rather than augment the teacher's role. It is the failure mode Chapter Six has been at the greatest length trying to prevent, and the one most likely to manifest if Chapter Six's recommendations are taken selectively rather than comprehensively. A teacher who has not been consulted on the reform, who has not been compensated for the additional work the reform demands, and who suspects — often correctly — that the tools are being deployed to justify administrative reductions in her autonomy, is a teacher who will, rationally, withhold the engagement the reform requires. The withholding is not sabotage; it is the specific professional response to specific professional threats. It will, in aggregate, produce a reform that has the tools but not the pedagogical engagement that makes the tools work.

The early-warning indicators are the rate at which teachers opt out of the first training cohorts (a rate above fifteen percent is a red flag), the qualitative feedback from early trainings (when the dominant register shifts from curious to wary), and the specific pattern of teacher-union communications about the reform (when unions shift from cautious engagement to explicit opposition). The mitigations are the four categories of protection I specified in Chapter Six — protection against deskilling, against surveillance, against algorithmic management, against rushed implementation — each codified in the Mission's statutory framework rather than left to the Mission's discretion. The protections are the specific promises the reform must make to the profession at its outset, kept in the specific form they were made, for the duration of the reform. A reform that breaks these promises is a reform that loses the profession, and losing the profession is the specific condition under which every other component of the reform fails.



III. Political discontinuity

The third failure mode is the change of government that results in the reform being de-emphasised, rebranded, or structurally revised beyond the point of its continuity. The Indian electoral cycle produces, on average, a change of government at the Centre every ten to fifteen years and in most states every five. A reform whose horizon is fifteen years will, with near certainty, traverse at least one change of government at the Centre and several at the states through which it operates. A reform designed to be protected against these transitions is a reform that survives them; a reform designed on the premise of continuous political sponsorship is a reform that will, at the first transition, discover that its continuity depends on decisions the incoming political leadership has no incentive to make.

The early-warning indicators are a decline in political speeches referencing the reform, a reduction in Cabinet-level attention in subsequent budget cycles, and, more specifically, any proposal to rebrand the Mission or to revise the outcomes framework. The mitigations are the statutory mission form specified in Chapter Ten, the non-partisan Governing Council composition that includes Chief Ministers across the political spectrum, and the specific public-commitment instruments that create political costs for revision. None of these mitigations makes the reform immune to political change — no design can — but together they raise the political cost of dilution to a level at which most incoming governments will find it easier to continue the reform than to reverse it. The precedent, within Indian institutional practice, is Aadhaar: a reform whose original design was controversial, whose continuation across three changes of government was secured through the specific statutory and institutional arrangements that made its dilution politically costly, and whose specific operational character has survived political transitions that its designers could not have predicted.



IV. Budget dilution

The fourth failure mode is the slow-motion reduction of the reform's fiscal envelope across successive budget cycles. It is the mode most likely to produce the reform's failure in years four and five, after the political excitement of the launch has faded and the ordinary pressures of competing priorities begin to claim a share of the ring-fenced allocation. A reform whose annual commitment is nominally protected but is, in practice, reduced by five percent a year through various administrative mechanisms is a reform whose delivery capacity, by year five, has been reduced by approximately a quarter. The reduction is invisible in any single budget cycle. It is decisive across the sequence.

The early-warning indicator is any proposed revision to the Mission's dedicated line in the Union Budget, including the specific tactics that are commonly used to reduce allocations without formally revising them: delaying releases to later quarters, imposing administrative conditions that slow disbursement, or shifting specific components into other budget lines that are separately subject to dilution. The mitigation is the statutory ring-fence in the Act, the public commitments made at the reform's announcement, and the specific insulation instruments that have protected other Indian commitments — the specific arrangements that protected the Fifteenth Finance Commission's devolutions, for instance, or that protected the MGNREGA allocations during their most politically contested years. The fiscal framework I proposed in Chapter Eleven codifies these insulations. Their codification is the specific protection. A Mission whose budget is visibly protected against dilution is a Mission whose officers can plan across the reform's horizon; a Mission whose budget is subject to annual renegotiation is a Mission that cannot plan beyond the current fiscal year, and cannot therefore deliver the multi-year work the reform requires.



V. Vendor capture

The fifth failure mode is the scenario in which a small number of private-sector providers, through the specific mechanisms described in Chapters Four and Five, come to shape the reform's decisions in ways that serve their commercial interests at the expense of the children's learning. The mechanics were specified in those chapters: specification-shaping capture, editorial-board capture, data-standard capture. The aggregate effect, when the mechanics operate together over several years, is a reform whose public character has been compromised without any single decision being visibly improper. The children continue to receive the reform; the reform continues to produce outputs; and the

specific shape those outputs take has been determined, quietly, by the commercial imperatives of the providers who have become embedded in the reform's decision-making.

The early-warning indicators are concentration of the commissioning process among a small number of providers (fewer than three providers winning more than seventy percent of tenders), specification choices that specifically advantage incumbent products, and the progressive convergence of the reform's standards on the technical architecture of a particular provider's offering. The mitigations are the procedural architecture specified in Chapters Four and Five — the DPI platform that prevents data-format capture, the open Repository standards that prevent content-format capture, the rotating review committees that prevent governance capture, the conflict-of-interest disclosure requirements that make the capture mechanisms visible when they begin to operate. None of these mitigations is sufficient on its own. Their aggregate application, consistently and across all of the Mission's commissioning and governance processes, is the specific discipline that protects the reform's public character over the duration of its operation.



VI. Delivery-chain collapse

The sixth failure mode is the pattern in which the administrative chain between the Centre and the classroom fails to transmit the reform's specific content with fidelity, producing an implementation that diverges substantially from what was designed. The mechanics were specified in Chapter Twelve — the seven-layer chain, the systemic friction at each layer, the specific distortions introduced at each translation. A Mission that issues excellent guidelines is, in the Indian administrative practice, a Mission whose guidelines are implemented at perhaps sixty percent of their intended fidelity, with the residual forty percent

either not implemented or implemented in distorted form. The specific gap is the specific failure mode, and the gap widens as the chain lengthens.

The early-warning indicator is the growing gap between the Mission's issued guidelines and the independent classroom-observation data on what is actually happening in schools. When the gap exceeds thirty percent, on the specific dimensions the Mission is tracking, the failure mode is manifesting and requires intervention. The mitigations are the delivery-chain disciplines I specified in Chapter Twelve — the direct-to-school digital channel that bypasses administrative intermediation for substantive content, the named coordinator model that creates specific personal accountability at the state and district levels, and the Rapid Response Unit that intervenes in specific districts where the standard process has failed. Each of these disciplines is separately necessary; together they are a set of mutually reinforcing protections against the specific pathology of long administrative chains. A Mission that operates all three, from year one, is a Mission whose guidelines reach the classroom at fidelities substantially higher than the Indian administrative default. A Mission that operates only one or two is a Mission whose delivery will, predictably, resemble the Indian administrative default more closely than it will resemble its own intentions.



VII. Equity erosion

The seventh failure mode is the pattern in which the reform, despite its equity commitments, produces outcomes that widen rather than narrow the existing inequities. The mechanics were specified in Chapter Seven: the same technology that could equalise can also, if deployed carelessly, further entrench the inequities it could have closed. The failure mode manifests when the reform's early deployments reach the better-positioned schools

first, when the content arrives in English and Hindi before the other scheduled languages, when the teacher training reaches the better-trained teachers faster than the weaker-prepared ones, and when the net effect of all these sequencing decisions is a widening rather than a narrowing of the educational divide the reform claimed to address.

The early-warning indicator is specific: divergence in district-level outcome metrics between the disadvantaged-district decile and the advantaged-district decile across any two consecutive reporting cycles. The mitigations are the Equity Fund I specified in Chapter Seven, the disaggregated measurement framework codified in Chapter Thirteen, the external evaluation cycles that surface divergence the Mission's own reporting may obscure, and the specific sequencing discipline — beginning pilots in states that represent a range of capacities rather than only the strongest — that I argue for in Chapter Fifteen. The equity protections are mutually reinforcing. A Mission that applies them all is a Mission whose reform will, over a decade, narrow the divide. A Mission that applies them selectively will, despite its equity rhetoric, widen it.



VIII. Data-governance failure

The eighth failure mode is an incident in which children's data is misused, leaked, or otherwise mishandled in a way that compromises the reform's public legitimacy. The mechanics are familiar from comparable technology deployments in other sectors — a vendor's data-handling practice is inadequate, a security vulnerability is exploited, a regulator's approval is found to have been granted on incomplete information, or, most commonly, a specific use of the data is discovered that the parents and teachers had not consented to and that the public regards as improper. The incident need not be large to be consequential. A single well-publicised case of a child's AI-tutor conversation being used for a

commercial purpose, or of a student's learning record being disclosed to a third party, is enough to produce the kind of public backlash that can set the reform back by years. Recovery from such an incident is substantially harder than its prevention.

The early-warning indicator is any material deviation from the data-protection protocols of the Children's AI Safety Code, discovered through the Mission's own audit infrastructure or surfaced by external scrutiny. The mitigations are the audit infrastructure specified in Chapter Sixteen, the vendor contractual framework specified in Chapter Five, the statutory safety authority that enforces the Code, and the specific governance arrangements that make any deviation visible before it becomes a public incident. The most consequential specific protection is the requirement that every vendor deployment in government schools is subject to pre-deployment audit against the Code, with the audit's results public. A pre-deployment audit catches the specific vulnerabilities that post-deployment incidents would otherwise expose. It is slower. It is, over the reform's horizon, substantially less costly than the incidents it prevents.



IX. Content inadequacy

The ninth failure mode is the pattern in which the Repository fails to produce content at the scale, language coverage, or quality required for the classroom deployments, producing classrooms equipped with tools but no material to animate them. The mechanics were specified in Chapter Four: the scale of the content requirement is substantial, the specific editorial discipline required is scarce, the origination work cannot be accelerated through simple translation, and the institutional capacity to produce the Repository at the scale the reform requires does not yet exist and must be built alongside the reform's other components. A Repository that falls behind its commissioning schedule produces

a reform in which the tools are deployed but the children's specific experience with them is degraded by inadequate or inappropriate content.

The early-warning indicator is the Repository's commissioning pipeline falling behind its published schedule by more than two quarters in any major language. The mitigations are the acceleration protocols that must be built into the Repository's operational framework — specific contingency commissions that can be activated when a language falls behind, a reserve editorial capacity that can be deployed to acceleration, the specific quality trade-offs the Repository is prepared to accept when the alternative is a deployment without content at all. The trade-offs are uncomfortable to specify in advance, but a Repository that has specified them is a Repository that can respond when the failure mode begins to manifest. A Repository that has not specified them is a Repository that will, when the failure mode occurs, either accept the deployment gap or make the trade-offs under pressure in ways that produce poor outcomes.



X. Implementation fatigue

The tenth failure mode is the slow erosion of institutional attention after the first two or three years, as the novelty of the reform fades and the sustained operational discipline it requires becomes harder to maintain. It is, in some ways, the most difficult failure mode to guard against, because it does not manifest through a specific event that can be responded to; it manifests through the gradual decline of the specific institutional behaviours that produced the reform's early success. The weekly classroom reviews become less regular. The monthly block reviews become perfunctory. The quarterly state reviews begin to be delegated to junior officers. The Mission's own internal reviews become routine rather than substantive. The behaviours that were critical in years

one and two lose their critical character in years four and five, and the reform's delivery erodes accordingly.

The early-warning indicators are specific: a decline in the quality and frequency of the Mission's internal reviews, rising staff attrition at the Mission and its state-level counterparts, and the specific pattern of officer-level communications becoming less substantive over time. The mitigations are structural rather than procedural. External evaluation cycles, commissioned from independent research organisations and published publicly, create the specific external pressure that sustains internal discipline. Parliamentary review at year ten creates a specific political milestone against which the Mission must demonstrate its continued performance. The distributed editorial structure of the Repository, with nodes across multiple Indian cities rather than a single Delhi concentration, makes the Mission's work more resilient to the attention decline that tends to affect Delhi-centralised institutions. None of these mitigations eliminates the failure mode. Together they slow it, and slowing it across the reform's horizon is the specific discipline that distinguishes reforms that survive into their second decade from reforms that do not.



XI. The failure-modes table – for the risk register

The ten failure modes above are the specific failure modes I consider most likely. Figure 14.1 collects them into a single table, suitable for the Mission's risk register and for the Cabinet note that establishes the Mission. The table should be revised annually. The annual revision is not a formality; it is the specific moment at which the Mission's leadership reviews what has changed in the reform's risk profile and adjusts the mitigations accordingly. A risk register that is not revised is a risk register that has become decorative.

Failure mode	Early-warning indicator	Pre-emptive mitigation
1. Tech-first deployment	Q1 device-procurement allocation exceeds teacher-training allocation	Budget rule: no device tender approved until matched teacher-training commitment is secured (Ch 11)
2. Teacher resistance	Opt-out rate in first training cohort > 15%; negative qualitative feedback dominates	Teacher protections codified in Act; consultation-first design; compensation increment (Ch 6)
3. Political discontinuity	Drop in ministerial speeches / budget-speech mentions; Cabinet-level attention declines	Statutory mission form with 15-year sunset; non-partisan Governing Council framing (Ch 10)
4. Budget dilution	Any proposed revision to dedicated Union Budget line for Mission	Statutory ring-fence in Act; public commitments; 15th Finance Commission-style insulation (Ch 11)
5. Vendor capture	Concentration of commissioning among ≤3 providers; incumbents winning > 70% of tenders	DPI architecture; open Repository standards; rotating review committees; conflict-of-interest disclosure (Chs 4, 5)
6. Delivery-chain collapse	Gap between issued guidelines and independent classroom observation > 30%	Direct-to-school channel; named coordinators; Rapid Response Unit (Ch 12)
7. Equity erosion	Bottom-decile / top-decile outcome ratio widens over any two consecutive cycles	Equity Fund formula-based allocation; disaggregated metrics; external evaluation (Ch 7)
8. Data-governance failure	Any material deviation from Children's AI Safety Code; any reportable data incident	Audit infrastructure; vendor contractual framework; statutory safety authority (Chs 5, 16)

9. Content inadequacy	Repository commissioning > 2 quarters behind published schedule in any language	Parallel commissioning across languages; adaptation-not-translation discipline; 200-person editorial core (Ch 4)
10. Implementation fatigue	Decline in internal-review frequency/quality after year 2; staff attrition rising	External evaluation cycles; Parliamentary review at year 10; distributed editorial structure (Chs 4, 10)

Figure 14.1 | Ten Failure Modes – early-warning indicator and pre-emptive mitigation.



Three decisions for the reader

1. If you are drafting the Mission's Cabinet note, include Figure 14.1 as an annexure.

Cabinet notes that identify specific risks, with specific mitigations, are Cabinet notes that survive the approval process more reliably than Cabinet notes that do not. The Cabinet members reviewing the note are looking, among other things, for evidence that the proposing ministry has thought seriously about what could go wrong. The failure-modes table is that evidence. Include it. Revise it based on the specific concerns the Cabinet's review surfaces. The revision strengthens the note rather than weakening it.

2. If you are on the Governing Council, make risk-register review a standing agenda item.

The Council meets twice a year. At least one of the two meetings should include a substantive review of the risk register, with the specific early-warning indicators reviewed against the current data. The Council that treats risk review as an annual

exercise is a Council that will, when a failure mode begins to manifest, find itself notified late. The Council that treats risk review as a semi-annual discipline is a Council positioned to intervene before specific failures become entrenched.

3. If you are a journalist, policy analyst, or civil-society researcher, test the mitigations.

The mitigations specified above are, in effect, commitments the reform has made to protect itself against each failure mode. The commitments are testable — they correspond to specific institutional behaviours that can be observed and reported on. A sustained external scrutiny of whether the mitigations are actually operating, as distinct from whether they have been announced, is among the specific external pressures that sustains the internal discipline of the reform. The pressure is not adversarial; it is a specific contribution to the reform's own success. Journalists, researchers, and civil-society organisations who hold the Mission to its mitigations are doing the work the internal review cycles, left to themselves, would not fully do.



END OF CHAPTER FOURTEEN

CHAPTER FIFTEEN

The Phased Rollout



Five years of work, sequenced against the country's institutional capacity and the reform's fiscal architecture.



A reform of this scale cannot be done everywhere at once. The question is not whether to phase it but how: in what sequence, with what state clustering, at what pace across the fiscal cycle, and with what specific milestones against which the phasing's progress is assessed. This chapter specifies a five-year phased rollout, organised around four specific design disciplines: the temporal sequencing of substantive work across the five years, the geographic sequencing across states of different capacities, the capex-opex composition of the fiscal commitment at each phase, and the specific milestones that serve as the phase-gates between one year's work and the next.

The phasing I propose is not the only possible phasing. It reflects my best judgement about what sequences most effectively given the specific institutional capacities and constraints the reform operates within, and it is offered as a starting point for the Mission's own planning rather than as a fixed prescription. The reader who finds specific elements of the phasing wrong will, I hope, engage with the specific elements rather than dismissing the framework as a whole. The phasing is revisable; the discipline of phasing itself is what the chapter argues for.



I. The temporal sequencing – years one through five

The first design discipline is the temporal sequencing. Figure 15.1 sets out the five-year trajectory I propose. The phasing has specific design features worth naming. Year one is entirely about architecture – no classroom in the country is directly touched by the Mission's deployment, because the Mission in year one does not yet exist in the form required to touch classrooms. The temptation, in year one, to produce visible classroom-level outputs is the temptation I flagged in Chapter Nine as the showcase failure, and it must be resisted. Years two and three are the pilot-and-learn phase, during which the reform's operational assumptions are tested against reality in a controlled number of districts, and the lessons are absorbed into the scaling plan. Year four is when the national ambition becomes visible to the ordinary parent and teacher – the point at which the reform either delivers on its promise or is seen, publicly, to have failed. Year five is consolidation and the first external review, which determines whether the second five-year phase proceeds on its current basis or on a revised one.

Year	Primary focus	Geography	Key milestones
Year 1	Architecture, not classrooms	Centre + 5 pilot districts agreed	Act passed; Mission constituted; Governing Council seated; 5 pilot districts agreed; standards published for consultation; Repository commissioning begins; DIKSHA extended with direct-to-school channel; dashboard specification published
Year 2	Pilot deployment + content wave 1	5 pilot districts + expansion to 50 districts across 15 states	First 500 schools deployed; first 50,000 teachers in training; Repository v1 available in 8 major languages; first Rapid Response Unit

			interventions; public dashboard live with baseline data
Year 3	Scale wave 1 + lessons absorbed	150 districts across all states	First independent external evaluation published; mid-course corrections applied; teacher certification reaches 500,000; equity disaggregation reporting mandatory; first cohort of B.Ed. graduates enters classrooms under reformed curriculum
Year 4	Scale wave 2 + secondary-stage rollout	All districts in 20 high/medium-readiness states; pilot in remaining states	Secondary-stage electives live in CBSE and 10 state boards; content available in all 22 scheduled languages at quality threshold; 2 million teachers certified; Equity Fund second tranche released
Year 5	Consolidation + review preparation	National; all states; all grade levels	Universal availability claim testable; first 5-year review for Parliament; Act amendment consultation begins; second 5-year phase planning; capex investment winds down; opex steady-state commitment begins

Figure 15.1 | The Phased Rollout – Years One through Five. The pace is deliberately uneven: year one is architecture; years two and three are pilot and scale-learning; years four and five institutionalise.

A reader who has worked in Indian public administration will notice that the phasing is, compared to the norms of the sector, substantially more gradual than is typical. A more aggressive phasing – nationwide rollout within two years, as some political pressures would favour – is technically possible but operationally

inadvisable. The specific institutional capacities the reform requires cannot be built, in most states, faster than the phasing above permits. A reform that attempts the faster pace is a reform that will encounter the failure modes of Chapter Fourteen at substantially higher intensities than a reform that phases itself deliberately. The gradual phasing is, in this specific sense, a protection against the reform's own ambitions.



II. The geographic sequencing – state clustering

The second design discipline is the geographic sequencing. The selection of which states engage at which phase is politically sensitive, and the temptation is to begin with the strong-capacity states – Kerala, Tamil Nadu, Maharashtra, Karnataka – because the early results will be better and the political optics more favourable. I would advise against the pure-strong-capacity approach, and I want to say specifically why.

A pilot conducted in favourable conditions produces results that cannot be generalised to the national deployment, because most of the country does not operate under the favourable conditions the pilot benefited from. A Kerala pilot tells us what is possible when the teacher cadre is well-trained, when the SCERT is capable, when the state political leadership is supportive, and when the infrastructure baseline is relatively strong. A Bihar pilot tells us what is possible at the conditions under which the reform will actually have to succeed nationally. Both are useful; the Bihar learning is more transferable to the national scale. The pilot framework should therefore include, in the first five pilot districts of year one, a mix of state capacities. My specific recommendation is: one district from a strong-capacity state (say Kerala or Tamil Nadu), one from a medium-capacity state with strong political sponsorship (say Karnataka or Maharashtra), one from a medium-capacity state with more contested political conditions (say Odisha

or Rajasthan), one from a weaker-capacity state with sufficient local administrative commitment (say Uttar Pradesh or Madhya Pradesh), and one from a state whose specific challenges of geography, language, or political context are distinctive (say Assam, Jharkhand, or one of the smaller north-eastern states). The five together produce a representative sample of the conditions the reform will encounter nationally.

The state clustering for years two and three expands this logic. The fifty districts of year two should represent, in aggregate, the range of Indian conditions rather than concentrating on the most favourable. The hundred-and-fifty districts of year three should extend the representation further, with specific attention to districts where earlier pilots have surfaced specific difficulties — these difficult districts are the specific learning opportunity that a rapid national scale-up would forego. Only by year four, when the lessons of the pilots have been absorbed and the Mission's delivery capacity has been tested, should the rollout expand to all districts in high-and-medium-readiness states. The weaker-readiness states continue their pilots through year four and scale in year five, under specific arrangements that respect their capacity constraints rather than overwhelming them. The phasing is, in this specific sense, a form of equity — the reform's faster phases are calibrated to the states that can absorb them, and the slower phases protect the states that would otherwise be set up to fail.

Urban-rural differentiation within each state is a separate dimension of the phasing. Urban schools typically have stronger infrastructure, better-trained teachers, and greater household engagement than rural schools. A state-level phasing that deploys uniformly within the state will, in practice, produce earlier and better outcomes in urban schools and delayed or weaker outcomes in rural schools. The reform must either accept this pattern (producing, over the decade, a widening urban-rural gap) or deliberately counteract it (deploying additional infrastructure investment in rural schools in each phase, so the effective timing of

deployment matches rather than diverges from the urban timing). The second path is more expensive and more administratively demanding. It is the path the reform's equity commitments require.



III. The capex-opex composition

The third design discipline is the capex-opex composition of the fiscal commitment at each phase. I discussed in Chapter Eleven the specific split – approximately fifty-five thousand crore in capital expenditure across devices, infrastructure, platform engineering, and initial Repository build; approximately sixty-five thousand crore in operating expenditure across teacher training, content maintenance, mission operations, connectivity, and Equity Fund disbursements. Figure 15.2 shows how the split unfolds across the five years, and the unfolding has specific implications for the reform's fiscal politics that are worth naming.

Year	Capex (₹ crore)	Opex (₹ crore)	Total (₹ crore)	Notes on composition
Year 1	8,000	3,000	11,000	Front-loaded platform and Repository build; minimal classroom deployment; Mission operations start-up
Year 2	14,000	9,000	23,000	First wave of device procurement in pilot districts; teacher training ramps; Repository wave 1; dashboard build
Year 3	15,000	13,000	28,000	Scale device procurement into 150 districts; training at scale; content maintenance becomes substantial

Year 4	12,000	16,000	28,000	Capex begins to moderate as deployments complete; opex dominated by teacher certification and compensation increment ramp-up
Year 5	6,000	24,000	30,000	Capex largely replacement/expansion only; opex at steady-state including full compensation commitment
5-year total	55,000	65,000	120,000	Overall split: 46% capex / 54% opex; capex front-weighted in Years 2-3, opex back-weighted in Years 4-5

Figure 15.2 | Capex and Opex by Year – the reform's fiscal envelope, by composition and year. Note the shift from capex-heavy in the early years to opex-heavy by year five, which is the specific transition Chapter Eleven flagged as politically vulnerable.

The composition matters because Indian fiscal politics treats capex and opex asymmetrically. Capital outlays are politically easier to justify and tend to be released more reliably, because they produce visible physical deliverables and can be spread across state-level counterpart funding. Operating outlays are politically harder — they require sustained annual commitments, they produce less visible deliverables, and they are more vulnerable to dilution in subsequent budget cycles. The reform's fiscal shape is such that the politically easier capex is concentrated in the early years, when the Mission's political capital is strongest, and the politically harder opex becomes dominant in the later years, when the political attention the reform receives has begun to fade. This is the specific pattern Chapter Eleven flagged as the budget-dilution failure mode's structural setup. The mitigation is the statutory protection of the opex commitment, built into the Act and into the Mission's governance arrangements, so that the later-year fiscal

politics do not overwhelm the substance of the commitment the earlier years made.

A separate observation: the year-five opex envelope of thirty thousand crore does not end in year five. It continues, at approximately this magnitude or somewhat higher, indefinitely — because the teacher compensation increment, the ongoing content maintenance, the platform operations, and the connectivity commitments are all continuing commitments rather than one-time expenditures. The hundred-and-twenty thousand crore five-year envelope is the reform's initial investment; the thirty-thousand-crore-per-year steady-state that follows is the reform's permanent operational cost. A reform that treats the five-year envelope as the whole commitment is a reform that will find, in year six, that it has no budget for the operational work the reform continues to require. The Cabinet note that commits to the five-year envelope must also, explicitly, commit to the steady-state funding that follows it. The explicit commitment is what makes the long-horizon sustainability of the reform credible.



IV. Phase-gates and the governance of transitions

The fourth design discipline is the specific phase-gates that govern the transitions between years. Without phase-gates, the phasing becomes a plan rather than a discipline — each year's work proceeds on the assumption that the previous year's work was completed, whether or not it actually was. With phase-gates, each transition requires a specific assessment that the conditions for the next phase have been met, and the Mission's leadership has the explicit authority to slow, pause, or redirect the phasing if the conditions have not been met.

The specific phase-gates I would propose are these. The year-one-to-year-two gate requires: the Act passed and assented to; the Mission's Governing Council seated; the five pilot districts agreed

with their respective state governments; the standards documents published for consultation; the Repository commissioning operational with at least twenty commissions awarded across at least six languages; the direct-to-school channel technically operational; and the initial budget released. The year-two-to-year-three gate requires: the pilot deployments operational in all five pilot districts; at least eighty percent of the year-two teacher-training cohort certified; the first external evaluation commissioned; and the dashboard live with baseline data. The year-three-to-year-four gate requires: the external evaluation published and responded to; teacher certification reached at 500,000; mid-course corrections applied; and, critically, the equity disaggregation reporting demonstrating that the bottom-decile and top-decile districts are not diverging. The year-four-to-year-five gate requires: secondary-stage electives operational in at least ten state boards; content available at quality threshold in all twenty-two scheduled languages; two million teachers certified; and the Equity Fund's second tranche disbursed on the formula-based allocation.

The phase-gates are not merely procedural. Each is a specific institutional commitment. The Mission that misses a phase-gate is a Mission that must publicly acknowledge the miss and specify the response — whether the phasing is slowed, whether specific components are redirected, whether the Mission's leadership requires specific intervention from the Governing Council. The acknowledgement is uncomfortable; it is also, in my observation, the specific discipline without which phased rollouts progressively slip their timelines without anyone at the Mission being forced to confront the slippage. A Mission that misses three consecutive phase-gates without institutional response is a Mission whose phasing has become fictional. A Mission that acknowledges each miss and revises the phasing in response is a Mission whose phasing continues to discipline its work.



Three decisions for the reader

1. If you are drafting the Mission's five-year strategic framework, codify Figure 15.1 as the phasing baseline.

Your framework will require revision as the specific conditions of year one become clear, but a documented baseline is the reference against which the revisions are themselves disciplined. Publish the baseline in the first six months of the Mission's operation. Revise it annually. Document each revision with the specific reasons. The revisions are the record of the Mission's learning; the baseline is what makes the record coherent.

2. If you are a state-level implementer, resist the pressure to accelerate ahead of the national phasing.

Some states — typically the strongest-capacity ones — will be tempted to move faster than the national phasing permits, either because they have the capacity to do so or because their political leadership wishes to demonstrate leadership on the reform. The temptation is understandable. It is also, in most cases, counterproductive. A state that moves substantially faster than the national architecture supports will find itself building components that will need to be rebuilt when the national standards catch up, investing in training that diverges from the national certification framework, or deploying content that will be superseded by the Repository's eventual offerings. The faster state's initial lead becomes a structural disadvantage. Patience with the national phasing is, for the stronger-capacity states, the specific discipline their own subsequent effectiveness depends on.

3. If you are a fiscal officer — Finance Commission, budget secretariat, or state finance department — attend specifically to the year-five transition.

The year-five transition, from capex-heavy investment to opex-dominated steady-state operation, is the specific moment at which the reform's fiscal character changes and at which the

specific protections that have supported the investment phase must shift to the specific protections that will support the operating phase. The Cabinet note that establishes the Mission should specify, explicitly, the fiscal architecture of the post-year-five period. The architecture should be revisited annually from year three onward, so that when year five arrives the transition is anticipated rather than improvised. A transition that is improvised is a transition that, in Indian fiscal practice, tends to produce the specific cuts that the budget-dilution failure mode would otherwise produce anyway. Anticipation is the specific discipline that protects the reform at the specific fiscal vulnerability that year five represents.



END OF CHAPTER FIFTEEN

PART V



*The Legal,
Comparative, and
Human Frame*

CHAPTER SIXTEEN

The Legislative and Regulatory Layer



The specific legal commitments the reform requires, and why they matter to people who are not lawyers.



This is the shortest chapter in the book, and the one that the primary reader may be most tempted to skim. I invite her to read it anyway, because the legal and regulatory framework the reform operates within is not an abstract matter for specialists. It shapes, directly, whether children's data is protected, whether vendors can be held to account, whether teachers' professional autonomy is respected, and whether the reform's protections against harm actually prevent the harms they are designed against. The specific legal provisions matter because, in the absence of specific law, the defaults of market power and institutional inertia produce outcomes that the reform's designers did not intend and that the reform's users cannot easily correct.

I will address four specific legal questions that the reform raises — the application of the Digital Personal Data Protection Act of 2023, the proposed Children's AI Safety Code, the intellectual property framework, and the regulatory architecture. I will treat each with as much brevity as responsible treatment allows, because the detailed legal work that each topic invites is the specialised work of lawyers and legislative drafters, of which this book's author is neither.



I. The DPDP Act and its application to classroom AI

The Digital Personal Data Protection Act of 2023 provides India's foundational framework for personal data protection, and its provisions apply directly to the data generated by AI tools deployed in classrooms. The Act defines a child as any person under the age of eighteen. It prohibits the processing of a child's personal data in a manner likely to cause any detrimental effect on the child's well-being. It prohibits the tracking, behavioural monitoring, or targeted advertising directed at children. It requires verifiable parental consent for processing a child's data. These provisions are the legislative floor on which a classroom AI system must build.

The operational translation of these provisions into specific standards for AI-education vendors is, however, non-trivial, and the reform requires that specific translation to be done explicitly rather than left to the discretion of individual vendors, schools, or state governments. Specifically: what constitutes verifiable parental consent in the context of a government school where the AI tool is part of the prescribed curriculum and parents may not have meaningful alternatives to consent? What specific categories of data are considered likely to cause detrimental effect, and who makes that determination? What retention periods apply to learning records, and what deletion obligations arise when a child moves schools or completes her schooling? What rights of access does the child herself have to her own data, and how are those rights exercised given that many children are too young to exercise them alone?

Each of these questions has a legal answer, and the answer matters operationally for the reform. The mission, in collaboration with the Data Protection Board once it is fully constituted, should produce within its first year a detailed operational guidance document that specifies how the DPDP Act applies to AI-education deployments in schools. The guidance should cover the full lifecycle of the data — collection, storage, processing, sharing,

retention, deletion — and should specify the controls that vendors must implement and the audit mechanisms through which those controls are verified. The guidance is not a substitute for law. It is the operational bridge between the Act's general provisions and the specific reality of the classroom, and without it the Act's protections are weakened by the practical difficulty of applying them in specific contexts.

One specific provision of the DPDP Act deserves particular attention. The Act permits the government to notify certain data fiduciaries as Significant Data Fiduciaries, with enhanced obligations including the appointment of a Data Protection Officer, periodic audits, and data protection impact assessments. I would argue that AI-education vendors operating at scale in government schools should be notified as Significant Data Fiduciaries, given the scale of the data they handle and the specific sensitivity of data about children. The notification requires action by the Central Government on the advice of the Data Protection Board, and the reform's political leadership should ensure that the notification occurs early rather than being delayed through the ordinary procedural pace that such notifications sometimes suffer.



II. The Children's AI Safety Code

Beyond the DPDP Act, I propose the development of a dedicated Children's AI Safety Code — a subordinate regulation, issued under the Mission's statutory authority, that addresses the specific risks of AI systems in education that general data-protection law does not squarely address.

The Code should cover, at minimum, six areas. First, content-safety guardrails with explicit lists of prohibited content categories — sexual content involving minors, content promoting self-harm, content that exploits children's developmental vulnerabilities — and required human-review protocols for content moderation

decisions. The specific list of prohibited categories should be developed with input from child protection specialists, and should be updated as the understanding of harms evolves.

Second, a prohibition on advertising of any form inside any AI system deployed in a school. The prohibition should be categorical. No exceptions for educational advertising, sponsored content, or native advertising. A child in a classroom is in a specific context of limited meaningful consent, and commercial messaging directed at her in that context is, in my judgement, incompatible with the purpose of public education. This provision will be resisted by vendors whose commercial models depend on advertising revenue, and the resistance should be declined. Vendors that cannot operate within the prohibition are vendors that should not be operating in Indian classrooms.

Third, transparency requirements that oblige any AI tutor or tool to explain its outputs in plain language on request from a child, parent, or teacher. The requirement is substantive, not cosmetic. A tutor that produces an explanation so technical that the child cannot understand it has not complied with the transparency requirement. The Code should specify that explanations must be in the child's language of instruction, at a reading level appropriate to her grade, and must address the specific question she asks rather than producing a generic statement. The vendors that find this requirement technically challenging will adapt, and the adaptation itself is valuable, because it pushes the industry toward systems that are genuinely understandable rather than opaque.

Fourth, audit requirements including annual third-party audits whose summary results are published. The audit should cover the vendor's compliance with all of the Code's provisions, the specific design choices the vendor has made that affect children's safety, and the specific incidents the vendor has identified during the audit period. The audit is not a marketing exercise; it is the specific mechanism through which the public can verify that the vendor's claims about compliance are substantiated. The auditors

must be accredited by the mission, must rotate to prevent capture, and must be subject to specific professional standards analogous to those that apply to financial audits.

Fifth, incident-reporting obligations that require vendors to disclose harmful outputs, data breaches, and significant accuracy failures within a defined window — forty-eight hours for serious incidents, two weeks for lesser incidents. The disclosure should be to the Mission's Evaluation and Equity Directorate, to the affected schools, and, in aggregated form, to the public. The disclosure obligation is the specific discipline that prevents vendors from concealing problems for commercial reasons, and its enforcement requires specific penalties for non-disclosure that are visible enough to serve as a genuine deterrent.

Sixth, a right to human appeal against any AI-generated evaluation that affects a child's academic standing. If a student receives an AI-graded assessment, or an AI-generated recommendation about her placement, or an AI-derived judgement about her progress, she must have the right to have the decision reviewed by a human — specifically, a qualified teacher who has the authority to revise the decision. The right is not cosmetic. It is the specific protection against the algorithmic management I warned against in Chapter Six, translated into a right that individual students can exercise in their own cases.

The Code's enforcement architecture requires its own institutional design. The Mission's Evaluation and Equity Directorate can serve as the Code's day-to-day regulator, but binding enforcement — the power to de-list a vendor, to fine a non-compliant school system, to order remedial action — should rest with a statutory body at sufficient distance from the mission's own delivery responsibilities to be credible. The Data Protection Board is one candidate. A dedicated Children's Online Safety Authority, proposed in several legal reform conversations over the past three years, is another. Either works; what is not acceptable is an

architecture in which the mission both delivers the programme and judges its own compliance.



III. The intellectual property question

The intellectual property framework for AI-education content deserves specific legislative attention, because the default frameworks produce outcomes that are, in the Indian context, substantially suboptimal.

Three specific questions arise. First, when a child uses an AI tutor to help draft an essay, who owns the resulting work? The default answer in most jurisdictions is the child, but the default assumes a creator-tool relationship that does not cleanly apply to substantial AI-generated content. The Indian legal position should be clarified explicitly: the child owns her work, the AI's contributions do not generate a separate copyright claim by the AI's developer, and the vendor may not assert rights over student-generated content as a condition of service.

Second, when a teacher uses the Mission's platform to generate a lesson plan, is that lesson plan copyrighted to the teacher, to the state, or to no one? I would argue the teacher, with specific provisions that allow the Mission and the state to use the lesson plan in aggregate pedagogical research and, with the teacher's consent, as a model for other teachers. The teacher's ownership is the specific recognition of her professional authorship, which the reform's broader commitment to teacher respect requires.

Third, when a vendor's model has been trained on Indian schoolchildren's interaction data — under strict de-identification but still traceable in aggregate to Indian classrooms — does the Indian state retain any claim on the improved model that results? This is the most difficult of the three questions, and the one whose resolution will most affect the long-term economics of the reform. My view, offered with the humility the question deserves, is that

the state should retain a specific claim — the right to require that models substantially improved by Indian classroom data be licensed for use in Indian public education at concessional terms, or that the data's contribution to the improvement be compensated in a way that flows back into the reform's priorities. The specific legal form of this claim is a matter for expert drafting. Its existence as a claim is the substantive question, and the reform should assert the claim rather than allow the default of unlimited vendor appropriation to prevail.

The mission should, within its first twelve months, publish a clear policy framework addressing each of these intellectual property questions, developed in consultation with the Intellectual Property Appellate Board, the Ministry of Commerce, and specialised legal counsel. The framework's publication is what provides the predictability that vendors, teachers, and students require for their own decisions about the reform.



IV. What the reform cannot solve legislatively

I want to close this chapter with an honest acknowledgement that not all of the reform's risks can be addressed legislatively. Some risks — the risk that teachers lose morale despite the protections I specified in Chapter Six, the risk that parents become passive rather than engaged despite the provisions I specified in Chapter Eight, the risk that the fourth theme of the curriculum is taught superficially despite the specifications I proposed in Chapter Three — are risks that law cannot prevent. They are risks that require cultural, professional, and institutional work that legislation supports but does not substitute for.

The temptation, in any reform, is to address every risk through specific legal provisions, on the theory that what is not legally required will not occur. The theory is wrong. Many of the most important components of the reform will occur, or not, for reasons

that have nothing to do with legal requirements — because the teachers believe in the work, or do not; because the parents engage, or do not; because the administrators take the reform seriously, or treat it as another passing priority. The legal framework I have described in this chapter is necessary for the reform's success. It is not sufficient. The sufficient conditions include the cultural and professional commitments that this book, in all its chapters, has been arguing for, and that no Act of Parliament can legislate.

The legal framework's proper role is to establish the baseline of protections below which the reform may not fall, and to provide the enforcement mechanisms for the commitments the reform itself makes. Beyond that baseline, the reform's success depends on the people who execute it — on their judgement, their commitment, their willingness to do good work in conditions that are, as in any reform of this scale, less than ideal. The legal framework makes the work possible. The people who do the work are what make it succeed.



Three decisions for the reader

1. If you are in government, establish the Children's AI Safety Code within eighteen months.

The Code's development requires specific consultation with child protection specialists, civil society organisations, industry representatives, and state governments. The consultation takes time, and the Code's development should not be rushed at the cost of its quality. But it also should not be allowed to drift indefinitely, because the period during which classroom AI deployments are taking place without specific safety regulation is a period in which children may be harmed by practices that a Code would have prevented. Eighteen months is the outer limit for a responsible development and adoption

timeline. Anything beyond that is evidence of insufficient political commitment to the Code's substantive purpose.

2. If you are an institution head, implement the Code's protections in advance of its legal adoption.

You do not need to wait for the Code to be formally adopted to implement its protections in your own institution. The prohibitions on advertising, the transparency requirements, the audit obligations, the human-appeal rights — each of these can be specified in your institution's own policies and incorporated into your contracts with vendors. By implementing the protections early, you contribute to the development of industry practice, and you protect your own students in the period before formal regulation arrives. The early implementation is an act of institutional leadership that the reform needs, and that your students deserve.

3. If you are a lawyer or a legal scholar, contribute to the regulatory development.

The legal work the reform requires is substantial, and it requires legal expertise that the Mission's own staff may not fully possess. If your professional practice is in a relevant area — data protection, child rights, intellectual property, administrative law — consider what contribution you can make, whether through formal consultation with the Mission, through pro bono advisory work, through academic publication that shapes the discourse, or through direct participation in the drafting committees that will do the detailed work. The legal quality of the reform's framework depends on the quality of the legal input it receives, and specific individual contributions by qualified practitioners can substantially affect the outcome.



END OF CHAPTER SIXTEEN

CHAPTER SEVENTEEN

What India Should and Should Not Copy



*Lessons from Singapore, China, the UAE, Estonia, and South Korea,
with the specific caveats that each comparison requires.*



India is not the first country to attempt a national AI-education rollout. Several countries have undertaken, over the past three to seven years, variants of what this book proposes, and their experiences offer lessons both cautionary and encouraging. I want to walk through five of these experiences specifically, because each has produced specific lessons that India's reform designers should consider, and each has specific limits on its applicability that India's reform designers should acknowledge. The chapter is shorter than its predecessors because its substantive content is comparative rather than prescriptive, and because the reader will benefit from brevity after the denser chapters on mission design, fiscal architecture, and legal framework.

I write this chapter with a specific concern in mind. Indian policy discourse has a tendency, in moments of reform ambition, to import models from other countries wholesale, without the careful adaptation that the Indian context requires. The discourse tends to divide, unhelpfully, between those who treat foreign experience as definitive guidance and those who reject it as irrelevant to Indian conditions. Both positions are wrong. Foreign experience is informative, when read carefully, and it is misleading when applied mechanically. The discipline I want to model in this

chapter is the specific reading of each country's experience — what it tells us, what it does not, and what adaptations would be required for the insights it offers to be useful to India.



I. Singapore

Singapore's approach to AI education, coordinated by the Ministry of Education and the Infocomm Media Development Authority since approximately 2017, has been characterised by heavy curriculum integration from the primary years, substantial investment in teacher training, and a tight coupling between AI education in schools and the country's workforce development strategy. Specific features of the Singapore approach include the Code for Fun programme at the primary level, the integration of AI topics across multiple subjects rather than as a standalone subject, and the Skills Future initiative that connects school AI education with adult continuing education in a coherent national skills framework.

The transferable lesson from Singapore is the integration discipline — the way Singapore has avoided the “AI as an add-on subject” trap by weaving AI literacy across existing subjects. The Singapore approach treats AI as a capability to be developed across the curriculum, which is consistent with the framing I proposed in Chapter Three. The specific implementation practices — how teachers are prepared for cross-curricular integration, how assessment accommodates the dispersed presence of AI across subjects, how the curriculum coordinators in individual schools manage the integration — are worth careful study by Indian counterparts.

The non-transferable elements are the familiar ones. Singapore's entire school system is smaller than one of India's districts. Its teaching workforce is well-paid, highly selected, and intensively trained. Its administrative capacity is among the

world's strongest. Its curricular coordination is concentrated in a single Ministry of Education with direct operational authority over schools. Approaches that rely on the centralised intensity only a city-state can manage do not generalise to a federation of twenty-eight states with nine and a half million teachers.

Beyond these structural differences, there is a specific cultural element of Singapore's approach that Indian reformers should note without attempting to import. Singapore's educational culture places a particular emphasis on performance metrics, standardised assessment, and rapid iteration based on data. This emphasis has produced, in the Singapore context, an education system that achieves remarkable outcomes but that has, in recent years, been criticised even within Singapore for the pressure it places on children. An Indian reform that imported the Singapore approach without modification would likely produce both the benefits and the costs, and the costs may be inconsistent with the broader educational values that Indian parents and educators hold. The discipline is to learn from Singapore's integration approach while being cautious about its performance-measurement culture.



II. China

China's approach, initiated around 2017 and accelerated significantly since 2020, has been the most ambitious in scale terms. It has deployed AI classrooms to tens of thousands of schools, built specialised AI-education curricula from primary through university, and invested heavily in domestic AI capability that feeds into the education system. Specific features include the New Generation Artificial Intelligence Development Plan of 2017, which named AI education as a national priority, and the subsequent AI curriculum development that has produced textbooks for primary, middle, and high school levels, deployed in multiple provinces.

The transferable lesson from China is the scale-of-ambition lesson. The Chinese government has approached AI education as a civilisational project, with sustained political attention at the highest level and budgets to match. India's reform needs the same scale of ambition. The specific mechanisms through which China has mobilised its educational infrastructure — the coordinated role of the Ministry of Education, the provincial governments, and the AI industry — offer specific insights into how such a mobilisation can be organised, even though the specific political context differs substantially.

The non-transferable elements are the political ones. Much of what China has done is possible only in a unitary political system with minimal state autonomy and with an educational culture that accommodates strong centralised direction. India's federalism is a feature to be worked with, not a bug to be overcome, and a Chinese-style command approach would collapse in the Indian political reality within weeks of its attempted implementation. More broadly, the Chinese approach has produced, alongside its achievements, specific forms of algorithmic surveillance and control in educational settings that are incompatible with the values this book's reform proposes. The emotion-recognition systems that have been deployed in some Chinese classrooms to monitor student attention and engagement are, in my judgement, specifically the kind of system the Children's AI Safety Code I proposed in Chapter Sixteen should prohibit.

Beyond the question of what India should and should not copy, there is a distinct question of how India should position what it builds. A national reform executed at the scale and diversity the Indian case requires, if it is executed well, will produce institutional knowledge that no other country in the world will have produced. Twenty-two languages, fourteen hundred district-level operational units, nine-and-a-half million teachers, two hundred and fifty million children, a federal constitution, a functioning digital public infrastructure — no other national reform has attempted AI

education at this combination of scale and complexity. The Indian reform, if it succeeds, becomes the specific reference case for the Global South and, in important respects, for the world. Indian institutional arrangements, Indian content standards, Indian teacher-capacity architectures, Indian approaches to equity-disaggregated measurement — all of these become exportable, through the specific channels of South-South cooperation and through the multilateral institutions in which Indian voice has been growing. The strategic significance of this positioning is, in my view, substantial. A Centre that treats the reform as an internal undertaking alone is a Centre that has underpriced its external implications. A Centre that treats the reform as a specific international positioning, building into its design the documentation and the knowledge-sharing arrangements that make the lessons accessible to other countries, is a Centre that is exercising the soft power the reform's scale makes available. I am not proposing that the reform's design should be distorted to serve external positioning; the reform's primary purpose is the Indian child. I am proposing that the external positioning is a legitimate secondary consideration, and that it should be planned for rather than neglected.

There is a separate question about China's approach that Indian reformers should consider carefully. The Chinese AI-education programme has been substantially integrated with the Chinese AI industry's own development, with Chinese companies — Baidu, Alibaba, Tencent, and others — providing the platforms and tools that the educational deployment runs on. The integration has produced specific benefits for the Chinese AI industry, which has acquired substantial educational-data assets through the deployment. Whether this integration serves educational purposes well, or whether it subordinates educational purposes to industrial policy, is contested even within China. India should think carefully about the appropriate relationship between its educational reform and its domestic AI industry, recognising that some degree of integration is inevitable and that too much integration risks

repeating the pattern the DPI architecture is specifically designed to prevent.



III. The United Arab Emirates

The United Arab Emirates has, over the past five years, made AI education a signature initiative of its national education strategy. Specific features include the 2017 National Strategy for Artificial Intelligence, which was among the world's earliest national AI strategies and explicitly named education as a priority area, and subsequent initiatives that have introduced AI as a subject across UAE schools, partnered with international vendors for platform provision, and invested in attracting global AI talent to the UAE ecosystem.

The transferable lesson from the UAE is the public-positioning lesson. The UAE has made AI education part of its national brand, and has used it to attract investment, international recognition, and the kind of human capital that national development requires. India could similarly use the reform as part of its global positioning as an AI-capable civilisation. The specific communication strategies, the international partnerships, the branding of AI education as a national priority — each of these offers insights for how India might position its own reform on the international stage.

The non-transferable elements are again the structural ones. The UAE is educating approximately a million school children, not two hundred and fifty million. It is a wealthy country with substantial fiscal capacity per child, and its schools are concentrated in urban contexts where infrastructure is strong. Approaches designed for small, wealthy populations do not translate automatically to enormous, economically-varied ones. More specifically, the UAE's approach has been heavily reliant on imported technology platforms and imported expertise, which is a reasonable approach for a small country without the population

base to develop domestic capability at scale, but which would be an inappropriate model for India, where the population base both enables and requires domestic capability development.

There is also a question about the UAE approach that I want to name honestly. The UAE's educational system operates within a political context that differs substantially from India's democratic context. The specific dynamics of curriculum development, of teacher autonomy, of parent engagement, and of critical discussion of AI's social dimensions, are shaped by this political context in ways that make direct comparison difficult. Indian reformers should take the UAE's ambition and its positioning seriously, while recognising that specific practices around curriculum content and classroom discussion may not adapt cleanly to the more pluralistic context India's reform operates in.



IV. Estonia

Estonia is, in my view, the most interesting of the four comparators and the one India has the most to learn from substantively. Estonia has built, over two decades, a genuinely digital education system, underpinned by a national digital-identity infrastructure, open content, and a teacher community that is genuinely comfortable with technology. Specific features include the early digitisation of education in the 2000s, the integration of programming into primary-school curricula beginning around 2012, and the continued development of AI-education components in recent years.

What Estonia has proved is that technology-enabled education can improve outcomes at national scale in a reasonably federal context, and that the infrastructural approach — identity, content, interoperability — produces better durable results than the top-down content-delivery approach. The Estonian experience demonstrates that digital public infrastructure for education is

feasible, that teachers can be successfully brought along in technology-enabled reform over a sustained period, and that the long-term benefits of infrastructural investment compound over decades in ways that single-deployment approaches do not.

India's DPI instincts naturally align with Estonia's, and the specific architectural choices I have proposed throughout this book — the open platform layer, the content repository, the teacher capacity development as a sustained multi-decade investment — reflect lessons that Estonia has, in a sense, already validated. A serious mission would send officers to Tallinn for extended study tours rather than merely commissioning reports about Estonia from Delhi-based consultants. The specific operational insights that emerge from sustained engagement with the Estonian education ministry are worth substantially more than any summary this book can provide.

The non-transferable elements of the Estonian approach are more modest than with the previous three comparators, but they are real. Estonia is a small country, with a population of approximately 1.3 million, and its educational system operates at a scale that permits management intensity that India's cannot match. Estonia has a specific cultural orientation toward digital tools that has been developed over twenty years and that India's education community will need its own time to develop. Estonia's teaching workforce is, in terms of preparation and professional status, different from India's in ways that affect how reform is received and implemented. The architecture Estonia has built is transferable; the specific cultural conditions under which it was built are not, and India's version of the same architecture will have to accommodate the cultural reality that India starts from.



V. South Korea

South Korea's recent experience offers a cautionary tale that India should pay close attention to. South Korea's announcement, in 2023, of AI-based digital textbooks across its school system was initially received as a bold move toward a genuinely AI-enabled education system. In practice, the deployment has run into substantial friction — with teachers' unions, with parents concerned about screen time, and with students whose outcomes in early deployment have been mixed. By 2024, the programme was being revised in response to concerns from multiple stakeholders, and its full implementation timeline has been extended.

The South Korean experience is not a reason for India to hesitate. It is a reason for India to sequence more carefully. The specific lessons are these: invest heavily in teacher and parent buy-in before mass deployment; pilot rigorously and publish honestly; be willing to slow down the rollout in regions where the preconditions are not yet met; do not announce full implementation before the supporting infrastructure and cultural acceptance are in place. South Korea's decision to move quickly, with less preparatory work than the programme's scale required, produced a visible political setback that a more patient approach would have avoided.

India's reform should take this lesson seriously. The temptation to announce dramatic national deployment, particularly in the period following the mission's establishment, will be strong. Resist it. A deployment that is announced with less ambition and delivered with greater care produces better long-term outcomes than a deployment that is announced with maximum ambition and delivered unevenly. The specific discipline — of measured announcements, careful sequencing, honest reporting of pilot results, and willingness to adjust in response to evidence — is what distinguishes the Estonian model

from the South Korean one, and the reform should explicitly choose the former over the latter.

Five International Comparators — What to Copy, What Not		
Each has a lesson. Each has a limit. India learns by choosing consciously.		
COUNTRY	WHAT TO COPY	WHAT NOT TO COPY
SINGAPORE	Integration discipline: AI woven across subjects rather than bolted on as standalone	City-state scale. Performance-metric culture. Centralised authority over all schools.
CHINA	Scale of ambition: civilisational project with sustained political attention & budget	Unitary command structure. Algorithmic surveillance in classrooms. Tight industry-education integration.
UAE	Public positioning: AI education as part of the national brand, attracting talent & capital	Small wealthy context. Imported platforms and talent. Political context limits open classroom discussion.
ESTONIA	DPI pattern: identity, content, interoperability. Twenty-year sustained teacher investment	Small country scale (1.3M people). India must start from a less digitally mature base.
SOUTH KOREA	Cautionary tale: what happens when you announce before you are ready	Rushed rollout without teacher & parent buy-in. Extended timeline under public pressure.
The failure mode is borrowing without consciousness. The discipline is conscious choice.		

Figure 14.1 | Five international comparators — what to copy and what not to.



VI. What this means for India's approach

The comparative review suggests specific lessons for the design of India's own reform. From Singapore, the integration discipline — AI as a capability across the curriculum, not a standalone subject. From China, the scale of ambition — a civilisational project with sustained political commitment. From the UAE, the public positioning — the reform as part of India's national brand. From Estonia, the infrastructural approach — the DPI pattern applied to education, with teacher capacity as a sustained multi-decade investment. From South Korea, the sequencing caution — careful preparation before dramatic announcement.

Each of these lessons reinforces specific elements of the reform I have proposed throughout this book. The integration discipline is

reflected in Chapter Three's grade-by-grade roadmap, which treats AI as a cognitive capability rather than as a subject to be added. The scale of ambition is reflected in Chapter Ten's mission design and Chapter Eleven's fiscal architecture, both of which propose commitments commensurate with the reform's civilisational consequence. The public positioning is a communication strategy the mission must develop, with specific attention to how India's reform is presented internationally as well as domestically. The infrastructural approach is the architectural pattern I have argued for consistently, particularly in the platform, content, and delivery-chain chapters. The sequencing caution is the specific discipline that Chapter Nine's hundred-day plan and the mission's first-year architecture embody.

India's reform will be its own reform. It will borrow from these international precedents where borrowing serves its purposes, and it will reject them where they do not fit the Indian context. The specific discipline the reform should hold is the discipline of conscious choice — choosing from each comparator what is genuinely transferable and being explicit about what is not. The failure mode is not borrowing too much or borrowing too little; the failure mode is borrowing without consciousness, which is what has produced, in past Indian reforms, the specific phenomenon of policy documents that cite international best practice without having thought carefully about whether the practice cited applies to Indian conditions.



Three decisions for the reader

1. If you are in the mission's team, commission serious comparative study early.

The mission's first year should include, as a specific deliverable, a detailed comparative study of the five international

experiences I have discussed, and of other experiences — Finland, the United Kingdom's scale-up, the African digital-education initiatives, Latin American experiments — that space permitted me to exclude from this chapter. The study should be commissioned from researchers who will spend substantial time in the comparator countries rather than produced as a desk exercise. Its findings should be published openly, both to inform the reform's own decisions and to contribute to the global discourse on AI education. The study's budget is modest relative to the reform's overall envelope; its return on investment is substantial.

2. If you are an educator, engage with international peers.

The learning that comes from direct engagement with peers in other countries is substantially more valuable than the learning that comes from reading comparative studies. Indian educators — at the institutional level and at the policy level — should build sustained engagements with their counterparts in the comparator countries. The engagements are facilitated by organisations like UNESCO, the OECD, the World Bank's education network, and specific bilateral arrangements that the Ministry of Education can establish. Your specific professional engagement with these peers will inform your own thinking in ways that will, over years, substantially improve your contribution to the reform.

3. If you are in government, resist the pressure to announce before you are ready.

The South Korean lesson is the specific lesson Indian government should take to heart. The political pressure to announce national-scale deployment, particularly in the period before a budget cycle or an election, will be substantial. Resist the pressure. The deployment that is announced when the preconditions — teacher capacity, content, infrastructure, vendor preparation — are actually in place is the deployment that delivers real outcomes. The deployment that is announced before those

preconditions are in place is the deployment that produces visible failure. The specific political courage the reform requires is the courage to delay announcements until the substance is ready to support them, even when the delay appears, in the short term, to reduce political credit.



END OF CHAPTER SEVENTEEN

CHAPTER EIGHTEEN

The Manifesto — Returning to Sunita



What the country looks like if we do this well, what it looks like if we do not, and what is required of each of us in the years that matter.



Sunita is eleven years old when this book opens. If the reform I have described in these pages is done well, she is nineteen when it reaches its first full maturity. Let me describe what she is doing at nineteen, in two versions — the one in which the reform succeeds, and the one in which it does not. I will then say what each of us — each reader of this book, in whatever role she holds — must do for the first version rather than the second to come true.

In the first version, Sunita is in her second year of an undergraduate programme at a state university about a hundred and fifty kilometres from Semra. She is studying environmental sciences because, somewhere around Class 7, she became interested in the water systems of her district and, with her Class 8 AI-literacy project, she built a simple classifier that could identify, from smartphone photographs, whether a specific kind of algal bloom was present in the village ponds. The classifier was not sophisticated. It was, however, the first time she had produced something that functioned, and the production of it changed her relationship with her own capabilities. By Class 10 she was the student whom her teachers sent to district-level science competitions. By Class 12 she was clear that she wanted to study environmental sciences, and she performed well enough on her

examinations to enter a state university, on a government scholarship that her district administration helped her apply for.

At her university, Sunita is working on her undergraduate thesis using AI tools that have become, by this point, substantially more capable than the tools she grew up with. She uses them competently, with judgement, with awareness of their limitations. She knows when to trust their outputs and when to verify them. She uses them to translate international research papers from English into Hindi so that she can read them carefully. She uses them to analyse data from the district water-quality monitoring programme she helped design in Class 11. She uses them to practise her English, because her university's instruction is in English and her foundational literacy is in Hindi. She uses them, occasionally, to help her think through difficult ethical questions about the research she is doing — not as a source of answers, but as a foil against which to test her own thinking.

Kusum Devi is fifty-five. She still teaches at the school in Semra, though she has been a senior teacher for four years now and mentors three younger teachers who joined after her. She earns approximately eighty thousand rupees a month, which includes the certification increments she accrued through completing the four levels of the National AI Teacher Institute's programme. She takes her work seriously, as she always did, but she does it now with the confidence of a professional whose expertise is respected. She has two AI tutors in her classroom for the Class 3 children she teaches, and she has become, over the years, a skilled practitioner of their use. She is occasionally invited to other districts to speak to teacher training sessions, and the invitations — and the modest honoraria that accompany them — are among the visible recognitions of the professional standing she has built.

In the second version, Sunita is twenty-three. She left school after Class 10, because the AI-education reform reached her school only partially, her Class 9 elective was taught by a teacher whose own preparation was inadequate, and the reform's failure to

produce visible benefits for her academic trajectory made continuing to Class 12 feel less worthwhile than contributing her labour to her family. She married at twenty and has a child. She works as an anganwadi helper in her village. She has heard of AI, of course — she uses a voice-based assistant on her phone sometimes, and the advertisements she encounters on social media reflect algorithmic targeting she is vaguely aware of — but she has no professional or conceptual relationship with the technology. She is not angry about this. She is, in many ways, a content young woman, and the life she is living is the life that her mother lived, with modest improvements. The reform that might have offered her a different trajectory did not reach her with enough substance to make the difference.

Kusum Devi is, in this version, fifty-nine, and her life has been harder. Her school received the reform's devices in Year Four, but the teacher training that should have accompanied them was cut short by a state-level fiscal crisis. She struggled with the technology, was reprimanded for poor implementation of the tools, and eventually stopped using them rather than continue to feel incompetent. Her salary has not risen in real terms, because the compensation increments the reform had envisaged were quietly shelved. She will retire in five years, with the sense that the final decade of her career was one in which she was asked to do more than she had been equipped to do, and in which her professional judgement was less respected than it had been at the start of her career. She does not blame the reform specifically. She blames a general decline in the conditions of her profession, which the reform was part of rather than a remedy to.

Both of these versions are plausible from where we stand today. I do not know which one will actually come to pass. I know that the specific design choices the reform makes in the next two to five years will determine, more than any other factor, which one we move toward. I know also that each of us — the reader of this book in whatever role she holds — has a specific contribution to

that determination, and that the aggregate of our contributions is what the outcome will be. There is no hidden force that will produce the good version without our work. There is no sufficient reason to expect the bad version either, if the work is done. The outcome is genuinely contingent, and the contingency is, at its deepest level, the sum of individual choices by individual people acting in their specific contexts over the coming decade.



I. What is required of the government

The government's specific contribution is the establishment of the architecture the preceding chapters have specified — the mission, the fiscal architecture, the legal framework, the delivery chain. None of these is optional. A government that establishes the mission without the fiscal architecture has established a body that cannot do its work. A government that provides the fiscal architecture without the legal framework has provided resources that can be captured by interests the reform was designed to protect against. A government that establishes all three without the delivery chain has established the apparatus of a reform that will produce documents without producing outcomes.

The government's political leadership — the Prime Minister, the senior ministers, the state chief ministers whose cooperation the reform requires — must hold the reform as a priority across budget cycles and political transitions. The specific commitment is substantial: the reform requires sustained political attention for at least a decade, which exceeds the usual attention span of political leadership on any given issue. Sustaining this attention requires specific mechanisms — the high-level Governing Council, the periodic Parliamentary reporting, the public dashboard that makes progress visible — that make it politically costly to neglect the reform even when other priorities press for attention.

The government's administrative leadership — the senior civil servants, the state administrators, the heads of the institutional landscape I described in Chapter Two — must deliver the reform through the administrative machinery they operate. The specific competence this requires is substantial: the reform is more complex operationally than most reforms the machinery has previously delivered, and the specific skills of mission delivery, as distinct from ministerial administration, will need to be developed through the work itself. The administrative leadership's specific contribution is the professional integrity with which the work is done — the refusal to let short-term political pressures produce shortcuts, the commitment to honest measurement of outcomes, the willingness to acknowledge failures and adjust the approach in response.



II. What is required of the institutional community

The institutional community — the school principals, the college deans, the university VCs, the training-institute heads — constitutes, as I argued in Chapter Nine, the shortest path from policy to classroom that India's administrative geography offers. Their specific contribution is the adoption of the reform within their own institutions, with the thoughtfulness and care the reform requires.

The specific acts this involves have been laid out in the hundred-day plan and the two-year path. The designation of senior leadership. The structured listening to teachers. The honest audit. The institutional vision developed in consultation. The careful pilot. The sustained teacher development. The infrastructure investment. The content curation. The governance arrangements. The evaluation capacity. Each of these is specific and achievable. Together, they constitute the institutional transformation that the aggregate of many institutions' reforms produces.

The institutional community's specific contribution also includes, beyond the reform's direct implementation, the development of the peer networks and professional communities that sustain the work over time. The institutional heads who engage with each other across institutions, who share their learning openly, who contribute to the national discourse on the reform's practice, are the specific individuals whose work becomes, in aggregate, the professional foundation of the reform's success. The absence of these networks is one of the specific weaknesses that has affected earlier education reforms, and the presence of them is one of the specific strengths the AI-education reform must deliberately build.



III. What is required of teachers

The teachers on whom the reform depends, and to whom Chapter Six was largely addressed, have a specific contribution that no other actor can substitute for. Their contribution is the willingness to engage seriously with a reform that, as I acknowledged, asks a great deal of them under conditions that are less than ideal.

The engagement is, at its core, the decision to invest in their own professional development — the time in training, the practice with new tools, the reading about the technology and its pedagogy, the reflection on their own teaching practice. The investment is substantial. It is, for teachers who are already overworked, a significant additional demand on their time and energy. The reform must, as I have specified, compensate this investment appropriately through the certification increments and the career-progression reforms. But the investment itself, at the individual teacher's level, is a choice the teacher must make, and no compensation structure can force it.

Teachers' contribution also includes the professional integrity to name the reform's problems when they occur — the instruction that is inadequate, the tools that do not work, the content that does not fit the context, the protections that are not being honoured. The naming is difficult in Indian administrative culture, where critique of official initiatives is often treated as disloyalty. The reform's success depends on teachers being willing, individually and through their professional organisations, to name the specific problems they encounter, so that the reform can adjust in response. The specific cultural work of building the space for this naming is itself part of the reform, and it requires teachers' active participation, not their mere compliance.



IV. What is required of parents

Parents' contribution, as I described in Chapter Eight, is the specific work of modelling, supporting, and guiding their children's relationship with artificial intelligence — work that cannot be outsourced to any institution. The specific acts are small and cumulative. The thirty-second conversations over years. The interest taken in what the child is learning. The questions asked of the school. The attention paid to the child's well-being as the technology enters her life. The modelling, through the parent's own practice, of what thoughtful adult engagement with AI looks like.

The contribution is also, collectively, the civic pressure that parents can bring to the reform's quality. Parents who engage actively with their schools, who attend parent-teacher meetings informed rather than passive, who raise concerns through the proper channels, who participate in the governance of their institutions where that participation is possible, produce a school system that is more accountable to its constituents than a system whose parents are absent or disengaged. The specific work of being an engaged parent is not a matter of heroic effort. It is the sustained

attention of people who love their children and who have taken the time to understand what their children need.



V. What is required of me, and of you

I want to close this book with a specific personal note, because the book has, in its final chapter, become more personal than policy handbooks typically are, and because the reader is entitled to know where the author stands.

I have written this book because I believe the reform is necessary, that it is achievable, and that the consequences of getting it wrong are severe enough that the time for writing about it is now. I have written it as one practitioner's contribution to what I hope will be a much larger conversation, in which many voices, some more expert than mine, will refine and correct what I have offered. I have tried to be honest about what I know and about what I do not, and I have tried to offer specific recommendations while being clear that the recommendations are mine rather than definitive answers. Whether I have succeeded in any of this is a judgement the reader will make.

What I will do, for my part, in the years that follow this book's publication, is continue the work this book represents. I will continue to engage with educators, with policymakers, with parents, with teachers, with my own daily work in a training-and-technology business. I will listen more than I will speak, because the central work of this reform is the work of listening to the people on whom its success depends. I will correct what I get wrong when I discover that I have got it wrong, because the discipline of correction is the specific discipline that distinguishes serious practitioners from those whose positions are too invested in their own correctness. I will, in whatever modest way my circumstances allow, contribute to the specific architecture the reform requires.

What I ask of you, the reader, is something I have been asking throughout the book in the three decisions at the close of each chapter. I will not repeat those decisions here. I will say that, in aggregate, the decisions are the specific mechanism by which this book might make a difference, and that a reader who has read the book without acting on any of the decisions is a reader the book has not adequately reached.

The reform this book describes is a civilisational project. It asks for commitments of time, money, political attention, and professional energy that are, together, larger than any educational undertaking India has attempted in its history. It also offers, if it is done well, the specific possibility of closing educational gaps that have persisted for generations, of preparing two hundred and fifty million children for a world their parents did not have the tools to prepare them for, and of establishing India as a country that has met the single most consequential educational challenge of the twenty-first century with the seriousness it deserves. The offer is not guaranteed. It is available to us if we do the work.

Sunita is eleven years old as I write this. She will be twelve next year, and thirteen the year after that. The decisions that will determine which version of her nineteen-year-old life actually comes to pass are being taken, in some form, today, in offices in Delhi and state capitals and school principals' rooms and parents' kitchens across the country. You, the reader of this book, are one of the people taking those decisions. Take them with the weight they deserve. Take them for her.



Three decisions for the reader

1. Act on one specific commitment from this book before the end of this month.

You have, by now, encountered perhaps forty-five specific decisions across the book's fifteen chapters. Do not wait until you have internalised them all. Pick one — any one — that corresponds to your role and your capacity, and act on it before the end of the month in which you are reading this. The specific act, however modest, is what translates the book from something you have read into something that has changed what you do. The second act will follow more easily than the first. The first is the one that matters.

2. Identify the person in your network best placed to act on this book's arguments, and give her a copy.

You may be a reader whose own role does not position you at the specific levers this book's recommendations require. You may know, however, a person who is — a principal at a school you care about, an officer in a department whose decisions affect your community, a philanthropist whose investment choices could accelerate the reform's priorities. Give her the book. The specific act of placing it in her hands, with your own note explaining why you think she should read it, is the network effect through which books of this kind actually influence the outcomes they are written to influence.

3. Name the child for whom you are doing this work, and keep her in view.

The first decision at the close of Chapter One was to name the child this reform is for, in your own mind. I ask you now to renew that naming, at the close of the book, with the specific weight of everything you have read since Chapter One. The child's name, whatever it is, is the discipline that keeps the work honest. Every institution you lead, every policy you shape, every investment you

make, every conversation you have about this reform — let it be tested, quietly, against the question of whether it improves her education. If it does, do it. If it does not, reconsider. The discipline is severe. It is also, in the end, the only discipline that produces reforms that actually matter.



END OF CHAPTER EIGHTEEN

APPENDICES

APPENDIX A

Glossary of Terms



For the reader encountering some of the reform's terminology for the first time.



AI literacy — The set of capabilities that enables a person to understand artificial intelligence as a technology, to use it competently, to evaluate its outputs critically, and to reason about its social and ethical implications. In the curriculum proposed in Chapter Three, AI literacy is developed across three stages from Class 3 to Class 12.

AI tutor — An educational software system, typically conversational, that interacts with a student to support her learning. AI tutors vary substantially in capability; the specific tutors most relevant to this reform are those operating in Indian languages, at the grade-appropriate level, on the specific subjects for which the reform deploys them.

Bhashini — The India Language Technology Mission under the Ministry of Electronics and Information Technology, which develops AI infrastructure for Indian languages including translation, speech recognition, and speech synthesis.

Block Resource Centre (BRC) — The administrative unit at the block level within a district, responsible for pedagogical supervision and teacher support across the schools in a block. The BRC is the final administrative layer above the school itself, and is one of the key delivery points for the reform's teacher training.

Central Board of Secondary Education (CBSE) — The largest national school board in India, with approximately thirty thousand schools across the country and internationally. CBSE introduced AI as a subject at the secondary level in 2019.

Children's AI Safety Code — The proposed subordinate regulation, described in Chapter Thirteen, that establishes specific safety standards for AI systems deployed in schools.

Concurrent List — The list of subjects in the Seventh Schedule of the Constitution of India on which both the Centre and the states have legislative competence. Education is on the Concurrent List.

DIET (District Institute of Education and Training) — The district-level institution for teacher education and training, of which there are approximately six hundred and fifty in the country. DIETs are the primary sites of in-service teacher training under the reform.

Digital Personal Data Protection Act (DPDP Act) — The 2023 Indian legislation establishing the framework for personal data protection, with specific provisions applicable to children's data.

Digital Public Infrastructure (DPI) — The Indian architectural approach to technology platforms, exemplified by Aadhaar, UPI, and DigiLocker, characterised by open standards, interoperability, thin public layers with private innovation, and explicit protection of user sovereignty.

DIKSHA — The Digital Infrastructure for Knowledge Sharing, the national platform for school education content and teacher resources, operated under the Ministry of Education.

Equity Fund — The dedicated fiscal line within the mission's budget, proposed in Chapter Seven, supporting states, districts, and institutions with the weakest starting baselines.

Foundational-to-Preparatory Stage — In the terminology of the curriculum proposed in Chapter Three, Classes 3 through 5, covering ages approximately 8 to 11.

Governing Council — The supreme decision-making body of the proposed National AI Education Mission, with composition specified in Chapter Ten.

KITE (Kerala Infrastructure and Technology for Education) — Kerala's state-level agency for technology in education, which has deployed significant infrastructure in state schools since 2002.

Middle Stage — Classes 6 through 8, covering ages approximately 11 to 14.

National AI Education Content Repository — The proposed national repository of curriculum materials, teacher guides, and related content for AI education, operated by NCERT with support from the Mission, described in Chapter Four.

National AI Education Mission (NAIEM) — The proposed statutory body described in Chapter Ten, with responsibility for coordinating the national AI-education reform.

National AI Teacher Institute — The proposed statutory body described in Chapter Six, with responsibility for the development of AI-education teaching capacity across the national teaching workforce.

National Council of Educational Research and Training (NCERT) — The national body responsible for curriculum development, textbook production, and educational research, established in 1961.

National Council for Teacher Education (NCTE) — The national regulatory body for teacher education programmes, including B.Ed. programmes.

National Curriculum Framework (NCF) — The framework document that specifies the overall structure, philosophy, and content of school education, most recently revised in 2023.

National Education Policy 2020 (NEP 2020) — The policy document providing the overall framework for education in India, adopted in 2020.

Nipun Bharat – The national mission for foundational literacy and numeracy, which provides a specific precedent for outcome-linked state grants.

PM-SHRI Schools – The Centrally Sponsored Scheme, launched in 2022, designating approximately fourteen thousand schools as showcase institutions for NEP implementation.

Rapid Response Unit – The proposed capacity within the mission, described in Chapter Twelve, for rapid intervention in districts where the reform's delivery is failing.

Samagra Shiksha – The Centre's overarching Centrally Sponsored Scheme for school education, covering a wide range of interventions from elementary through secondary levels.

Secondary Stage – Classes 9 through 12, covering ages approximately 14 to 18.

SCERT (State Council of Educational Research and Training) – The state-level counterpart to NCERT, with responsibility for state curriculum adaptation, teacher training, and educational research.

SIE (State Institute of Education) – In some states, the agency responsible for teacher education and training at the state level, sometimes operating alongside or within the SCERT.

END OF APPENDIX A



APPENDIX B

Index of Decisions



The three decisions from each chapter of the book, gathered for reference.



Chapter One – The Class 3 in Barabanki

1. Name the child this reform is for, in your own mind.
2. Identify the decision you are currently avoiding.
3. Decide what you would need to read, in this book, to act.

Chapter Two – What the NEP Promises, and What It Does Not

1. Read the NEP yourself, or reread it.
2. Name the single NEP commitment you believe your institution is most under-delivering against.
3. Decide whether your institution's AI-education provision will be centre-led or institution-led.

Chapter Three – The Grade-by-Grade Roadmap

1. If you are an institution head, commission your own institution's AI-education roadmap within six months.
2. If you are a parent, ask your child's school for its roadmap.
3. If you are an officer of the Centre or the state, consider whether a national roadmap should be issued.

Chapter Four – The Content Question

1. If you are an institution head, audit your current AI-education content within three months.
2. If you are a content producer, consider commissioning your work into the Repository.
3. If you are a policymaker, commission the Repository within the next twelve months.

Chapter Five – The Technology Question

1. If you are an institution head, write your technology specification independently of any vendor's offering.
2. If you are a policymaker, commission the DPI platform within the next eighteen months.
3. If you are a parent, ask your child's school for its technology procurement policy.

Chapter Six – The Teacher Question

1. If you are an institution head, ask your teachers what they need, and listen.
2. If you are a policymaker, make the teacher the explicit centre of your reform's design.
3. If you are a teacher reading this, organise.

Chapter Seven – The Equity Question

1. Disaggregate every metric.
2. Name the child the reform must most urgently serve, and test every decision against her.
3. Allocate the Equity Fund, or its equivalent at your scale, transparently.

Chapter Eight – A Letter to Parents

1. Have the first conversation with your child this week.
2. Ask your child's school the questions in Section V.
3. Notice, honestly, the AI use in your own life, and what it is modelling.

Chapter Nine – The Institution Head's Playbook

1. Within the next two weeks, designate your senior lead and begin the hundred-day plan.
2. Publish your institution's two-year path to its community.
3. Join, or start, a peer network of institution heads engaged in this work.

Chapter Ten – Building the Mission

1. If you are in government, draft the Cabinet note within ninety days.
2. If you are a Chief Minister or your office, signal engagement early.
3. If you are an independent member candidate, prepare.

Chapter Eleven – Paying for It

1. If you are in government, ring-fence the allocation in the next Union Budget.
2. If you are in state administration, model the reform's fiscal implications for your state.
3. If you are a philanthropist or institutional funder, engage the mission's co-investment framework early.

Chapter Twelve – The Delivery Chain

1. If you are in state administration, nominate your mission coordinator with care.
2. If you are in the mission's early team, build the direct-to-school channel before the classroom-facing deployments.
3. If you are a teacher or a headmaster, use the channels the mission provides.

Chapter Thirteen – The Legislative and Regulatory Layer

1. If you are in government, establish the Children's AI Safety Code within eighteen months.
2. If you are an institution head, implement the Code's protections in advance of its legal adoption.
3. If you are a lawyer or a legal scholar, contribute to the regulatory development.

Chapter Fourteen – What India Should and Should Not Copy

1. If you are in the mission's team, commission serious comparative study early.
2. If you are an educator, engage with international peers.
3. If you are in government, resist the pressure to announce before you are ready.

Chapter Fifteen – The Manifesto

1. Act on one specific commitment from this book before the end of this month.
2. Identify the person in your network best placed to act on this book's arguments, and give her a copy.

3. Name the child for whom you are doing this work, and keep her in view.



END OF APPENDIX B

APPENDIX C

A Model AI-Literacy Curriculum Specification



A concrete curricular artefact for policymakers and curriculum designers.



This appendix provides, in outline form, a specific curriculum specification for the AI-literacy provision described in Chapter Three. The specification is intended as a starting point for the more detailed curricular work that NCERT, the state SCERTs, and the proposed National AI Education Mission will undertake. It is offered in the spirit of a practitioner's contribution to specialised work that will be done by others with more expertise than the present author.

Classes 3 to 5 – Foundational-to-Preparatory Stage

Overall aim: To use AI tools to support foundational literacy and numeracy while developing the child's first grounded understandings of artificial intelligence.

Curricular footprint: Integrated across language, mathematics, and environmental studies periods, approximately two to three hours per week of AI-tutor-supported practice, plus approximately fifteen minutes per week of explicit conversation about the technology.

Core learnings by end of Class 5: the child can read fluently in her mother tongue and at a functional level in English; can perform the four operations of arithmetic confidently; can explain, in her own

words, that a machine learns from data, that it knows only what it has been shown, and that it can be wrong about things it has not been shown; has completed at least four human-in-the-loop projects in which she collaborated with an AI tutor on a specific task.

Classes 6 to 8 – Middle Stage

Overall aim: To introduce AI as a subject of study, through three parallel strands – computational thinking, data literacy, and AI literacy proper – and to deepen the integration of AI tools across other subjects.

Curricular footprint: One dedicated period per week (40 minutes) for AI literacy, plus continued integration of AI tools across other subjects.

Three strands, running in parallel across three years: Computational thinking strand (problem decomposition, algorithmic reasoning, Scratch progressing to simple Python); Data literacy strand (reading charts, identifying misleading presentations, correlation versus causation, designing simple surveys); AI literacy proper strand (what machine learning is, how classifiers are trained, hands-on training of at least three simple ML models, introduction to social and ethical questions).

Core learnings by end of Class 8: can decompose a moderately complex problem; can read a chart and identify misleading presentation; has trained at least three ML models and can explain why each succeeded or failed; can use an AI tutor productively in at least one other subject; has had at least a dozen substantive classroom conversations about AI's ethical and social dimensions.

Classes 9 to 12 – Secondary Stage

Overall aim: To offer three optional two-year electives that allow different kinds of students to pursue AI education at the depth their intended paths require, while embedding AI tools as genuine research partners across all subjects.

Three electives: AI and Society (history, economics, political science, philosophy, literature engaging with AI's social and ethical dimensions; final project: substantive essay on an AI-related policy question). Applied AI (programming in Python, introductory mathematics of machine learning, training and deployment of at least two working AI applications; final project: deployed application with documentation). Data and Design (statistics, data visualisation, design thinking; working with real Indian datasets; final project: substantive data analysis and visualisation).

Core learnings by end of Class 12: every student is capable of using AI tools as a professional adult or university student would require; can argue a defensible position on at least one substantive question about AI in society; students completing Applied AI have deployed at least two working AI applications; students completing Data and Design have produced at least two substantive data projects; students completing AI and Society have produced written work at a depth respected in university admissions.

Cross-cutting: the fourth theme – AI and Society

Running through all stages from Class 3 to Class 12, not as a separate strand but as a theme entering through other subjects and deliberate classroom conversations. Questions addressed include: What does it mean for a machine to be biased? Who decides what a machine is allowed to do? What happens to workers whose jobs machines take over? What does privacy mean when machines know everything we tell them? What does honesty mean when a machine can produce plausible false statements at scale? Pedagogical goal: familiarity with these questions as matters the child has the right to think about, the capacity to form views on, and the responsibility to revisit as she grows.



END OF APPENDIX C

APPENDIX D

A Principal's Hundred-Day Checklist



The concrete artefact for the operator-reader.



Weeks 1-2

Designate senior lead for AI-education initiative (VP, Dean, Academic Head, or equivalent). Announce designation to faculty and senior team. Allocate senior lead's time commitment (part-time if institution under 1,000 students; full-time if larger).

Weeks 3-4

Hold structured listening exercise with teaching staff. Ask: What do you know about AI? What do you want to know? What do you need from the institution to engage with this? Take detailed notes. Report back at next faculty meeting. Acknowledge specific points raised.

Weeks 4-6

Conduct honest institutional audit across six dimensions: content, teacher capacity, technology, student experience, data practices, parent engagement. Document gaps candidly. Prepare summary for internal discussion.

Weeks 6-10

Form working group (senior lead, 3-4 teachers, 2-3 parents, 1-2 senior students if appropriate). Draft institutional vision document

(5–10 pages). Circulate for consultation within the institution. Revise based on feedback. Present to governing body for formal adoption.

Weeks 10–12

Communicate adopted vision to all constituencies: teachers (full briefing), parents (written and meeting), students at secondary stage (age-appropriate presentation), governing body and alumni (summary).

Weeks 12–14

Identify one specific pilot for next academic term: one grade, one subject, one tool, volunteer teachers. Establish pilot documentation protocols: baselines, methods, weekly reflections. Brief pilot teachers thoroughly. Agree on expectations, protections, evaluation approach.

Throughout the hundred days

Invest in own learning: 50–80 hours across the period. Read this book and others. Attend one or two serious conferences. Visit institutions where deployment has begun. Have conversations with trusted peers.

End of hundred days – deliverables

Senior lead in place with time committed. Teacher listening exercise completed, with report. Institutional audit documented. Vision adopted by governing body. Communication programme complete. Pilot initiated with documentation protocols. Principal has substantive working understanding of the field.



END OF APPENDIX D

APPENDIX E

Further Reading



A selective bibliography for the reader who wishes to go further.



Indian policy documents

National Education Policy 2020. Ministry of Education, Government of India. Available in English and Hindi at education.gov.in.

National Curriculum Framework for School Education 2023. National Council of Educational Research and Training. Available at ncert.nic.in.

National Strategy for Artificial Intelligence. NITI Aayog, 2018.

Annual Status of Education Report (ASER). Pratham, published annually since 2005. Available at asercentre.org.

International references

UNESCO, AI and Education: Guidance for Policy-makers (2021).

OECD, Artificial Intelligence and the Future of Skills (multiple volumes, 2021–2024).

World Bank, Reimagining Human Connections: Technology and Innovation in Education at the World Bank (2020).

On Indian education more broadly

The ASER reports provide, over nearly two decades, the single most important evidence base on actual Indian learning outcomes.

Publications from the Centre for Policy Research, the National Institute of Educational Planning and Administration, and the Indian Institute of Management Ahmedabad's education group offer specific analyses of Indian education reform.

On AI and education internationally

The work of the Stanford Institute for Human-Centered AI's education group, the MIT Media Lab's Personal Robots Group, and the UCL Knowledge Lab offers serious research on AI in education from international contexts.

On the convergence of AI, metaverse, and quantum

For readers interested in the broader technological context in which this reform is occurring, the author's earlier book *Reimagining Reality* provides a wider treatment of the three converging technologies and their implications across multiple sectors beyond education.



END OF APPENDIX E

Bibliography



Principal works the book is in conversation with. Not a citation of every reference; a map of the intellectual neighbourhood.



Indian education policy and frameworks

Government of India, Ministry of Education. National Education Policy 2020. New Delhi: Ministry of Education, 2020.

Government of India, Ministry of Education. National Curriculum Framework for School Education 2023. New Delhi: NCERT, 2023.

Government of India, Ministry of Education. National Curriculum Framework for Foundational Stage 2022. New Delhi: NCERT, 2022.

Kothari, D.S. et al. Report of the Education Commission 1964–66 (the Kothari Commission Report). New Delhi: Government of India, 1966.

Government of India. The Right of Children to Free and Compulsory Education Act, 2009. New Delhi: Gazette of India, 2009.

Government of India, Ministry of Education. NIPUN Bharat Mission – Implementation Guidelines. New Delhi: Department of School Education and Literacy, 2021.

Government of India, Ministry of Education. Samagra Shiksha – Framework for Implementation. New Delhi: Department of School Education and Literacy, successive editions 2018–2024.

Government of India. Annual Status of Education Report (Rural) – ASER, successive editions. New Delhi: ASER Centre / Pratham, 2005 onwards.

Teachers and teacher education in India

Justice Verma Commission. Report on Teacher Education. New Delhi: Ministry of Human Resource Development, 2012.

National Council for Teacher Education. Various regulations and frameworks on pre-service teacher education, 2014–2024.

Kumar, Krishna. Political Agenda of Education: A Study of Colonialist and Nationalist Ideas. New Delhi: Sage, 2005 (2nd edition).

Batra, Poonam. 'Voice and Agency of Teachers: Missing Link in National Curriculum Framework 2005.' Economic and Political Weekly, 2005.

Sarangapani, Padma. Constructing School Knowledge: An Ethnography of Learning in an Indian Village. New Delhi: Sage, 2003.

Artificial intelligence, education, and pedagogy

Russell, Stuart and Norvig, Peter. Artificial Intelligence: A Modern Approach. Pearson, successive editions.

Holmes, Wayne; Bialik, Maya; Fadel, Charles. Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Center for Curriculum Redesign, 2019.

UNESCO. AI and Education: Guidance for Policy-makers. Paris: UNESCO, 2021.

OECD. Opportunities, Guidelines and Guardrails for Effective and Equitable Use of AI in Education. Paris: OECD, 2023.

Luckin, Rose. Machine Learning and Human Intelligence: The Future of Education for the 21st Century. London: UCL IOE Press, 2018.

Mollick, Ethan. Co-Intelligence: Living and Working with AI. Portfolio / Penguin, 2024.

Khanmigo / Khan Academy. Public materials on tutor design and evaluation, 2023–2025.

Digital public infrastructure and India's technology stack

Nilekani, Nandan and Shah, Viral. *Rebooting India: Realizing a Billion Aspirations*. Penguin Allen Lane, 2015.

Sharma, Siddharth. *Aadhaar — A Biometric History of India's 12-Digit Revolution*. Rupa, 2020.

Ministry of Electronics and Information Technology. *India's Digital Public Infrastructure — Reference materials on Aadhaar, UPI, DigiLocker, DIKSHA, Bhashini*, 2020–2025.

Bhashini — Digital India Bhashini Division. *Technical reports and platform documentation*, 2022–2025.

Equity in Indian education

Drèze, Jean and Sen, Amartya. *An Uncertain Glory: India and Its Contradictions*. Princeton University Press, 2013.

Ramachandran, Vimala (ed.). *Gender and Social Equity in Primary Education: Hierarchies of Access*. New Delhi: Sage, 2004.

Mehrotra, Santosh (ed.). *The Economics of Elementary Education in India*. New Delhi: Sage, 2006.

Shavit, Yossi and Blossfeld, Hans-Peter (eds.). *Persistent Inequality: Changing Educational Attainment in Thirteen Countries*. Westview Press, 1993.

Public administration, delivery, and reform at scale

Pritchett, Lant. 'Is India a Flailing State? Detours on the Four Lane Highway to Modernization.' HKS Working Paper, 2009.

Khera, Reetika. *Dissent on Aadhaar: Big Data Meets Big Brother*. Orient BlackSwan, 2019.

Kapur, Devesh. 'Why Does the Indian State Both Fail and Succeed?' Journal of Economic Perspectives, 2020.

Mukhopadhyay, Partha; Rangarajan, C.; and others. Various policy papers on Centre-State fiscal federalism, 2015–2024.

15th Finance Commission of India. Report for 2021–2026. New Delhi: Government of India, 2020.

Pritchett, Lant; Woolcock, Michael; Andrews, Matt. Building State Capability: Evidence, Analysis, Action. Oxford University Press, 2017.

International comparators

Ministry of Education, Republic of Singapore. Various publications on Smart Nation, EdTech Masterplan, and teacher professional development.

Kallas, Kristina and colleagues. Materials on Estonia's digital education transformation, various dates.

OECD. PISA — Programme for International Student Assessment, successive assessment cycles and volumes.

Hargreaves, Andy and Shirley, Dennis. The Fourth Way: The Inspiring Future for Educational Change. Corwin, 2009.

Zhao, Yong. Who's Afraid of the Big Bad Dragon? Why China Has the Best (and Worst) Education System in the World. Jossey-Bass, 2014.

Child safety, data, and AI governance

Government of India. Digital Personal Data Protection Act, 2023. Gazette of India, 2023.

UNICEF. Policy Guidance on AI for Children. Various editions, 2021–2024.

Ada Lovelace Institute. Reports on AI ethics, education technology, and child protection, 2020–2025.

European Commission. AI Act — text and supporting materials, 2024–2025.

Works on learning, development, and pedagogy (general)

Bruner, Jerome. *The Process of Education*. Harvard University Press, 1960.

Dewey, John. *Democracy and Education*. Macmillan, 1916.

Tagore, Rabindranath. *Personality: Lectures Delivered in America*. Macmillan, 1917.

Freire, Paulo. *Pedagogy of the Oppressed*. Herder & Herder, 1970 (English edition).

Kumar, Krishna. *What Is Worth Teaching?* Orient BlackSwan, successive editions.



A note on this bibliography. The works listed here are those I have relied on most directly in forming the book's argument, together with those I would recommend to a reader who wishes to pursue any specific theme further. The list is not exhaustive, and I have deliberately kept it readable rather than attempting the completeness that a scholarly bibliography would require. Readers who wish a more detailed set of references on any specific subject are directed to the relevant sections of Appendix E (Further Reading) and to the specific institutional publications cited in the chapter text. I have, throughout, tried to acknowledge the specific authors whose work has shaped my thinking, even where my specific conclusions diverge from theirs. Where I have failed to acknowledge a specific debt adequately, the failure is an oversight I apologise for, and I commit to correcting in any subsequent edition.

Thematic Index



*Organised by concept rather than by page.
References are to chapter numbers.*



This index is organised thematically rather than alphabetically by specific term, because the book's argument moves through a set of intersecting concepts that a strict alphabetical index would fragment. References are to chapter numbers rather than page numbers, since pagination in digital editions varies with the reader's settings. Within each thematic cluster, the specific sub-concepts are listed alphabetically, with chapter references in the order they appear in the book.

The Child at the Centre

Barabanki, Class 3 child in (Sunita) 1, 15
Composite naming of children Disclaimer, 1
Equity as the test of every decision 1, 7, throughout
Foundational literacy and numeracy (FLN) 3, 7, 13
Individual child, reform's measurement against 1, 7, 15
Lakshmi in Umaria district 7
Learning outcomes at each stage 3
Mother tongue, teaching and tutoring in 3, 4, 7
Rural child, specific needs and conditions 7, 11
Sunita, as reason for the book 1, 15
Tribal languages and children 4, 7

Curriculum and Pedagogy

Classes 3–5, foundational-to-preparatory stage 3
Classes 6–8, middle stage 3
Classes 9–12, secondary stage and electives 3
Cognitive capabilities vs new subject 3
Computational thinking strand 3
Curricular restraint at early stages 3
Data literacy strand 3
Discussion-based pedagogy, fourth theme 3, 6
Electives at secondary stage (AI and Society, Applied AI, Data and Design) 3
Human-in-the-loop projects for children 3
Machine as tool, as data-shaped, as fallible (three understandings) 3
Pedagogy, content-neutral specification 3
Portfolio-based assessment at middle stage 3
Roadmap, grade-by-grade 3

Content and Languages

Adaptation versus translation 4
Assessment content, reform of 4
Bhashini, infrastructure vs content 2, 4, 5
Commissioning process design 4
Content Repository, National AI Education 4, 10
Cultural-fit review 4
Editorial cadre, distributed across cities 4
Ideological capture of content 4
Invisibility of content work (political problem) 4
Languages, scheduled twenty-two 4, 7
Maintenance of content as living institution 4

Parallel origination across languages 4
Tribal and non-scheduled languages 4, 7
Vendor capture, commercial 4, 5

Technology and Platforms

AI tutor, current capabilities and limits 5, 3
Children's AI Safety Code 5, 13, 13
Connectivity and power infrastructure 5, 11
Data residency (Indian jurisdiction) 5, 13
Device specifications and unit costs 5, 11
DPI platform architecture 5, 10
Engagement-optimisation harms 5, 8
Foreign vs domestic foundation models 5
Framework agreements (national procurement) 5
Open public layer / thin platform 5
Procurement discipline and templates 5, 13
Section 8 operating entity (NPCI-style) 5
Switching costs and vendor lock-in 5
Vendor capture mechanics 5

Teachers – The Book's Centre

Ageing of teaching workforce 6
Algorithmic management, protection against 6
B.Ed. programme reform 6
Brojeswor, Moirangthem (Manipur) 6
Classroom coordinator role at school level 6, 12
Compensation increment (10% certified) 6, 11
Contract teachers 6
Data surveillance, protection against 6
Deskilling, protection against 6

Devi, Kusum (Barabanki) 1, 6
Kerala teacher (Balakrishnan, Renu) 6
Nine-and-a-half million, implications of scale 6
National AI Teacher Institute 6, 10
Peer learning as core layer 6, 12
Professional respect as primary register 6
Protections, four categories 6
Pre-service vs in-service 6
SCERT, DIET, BRC cascade 6, 12
Six components of preparation 6
Trust as the reform's foundation 6

Equity

Caste and community dimensions 7
Cognitive and learning disabilities 7
Disability dimension 7
Disaggregated metrics 7, 13
Equity Fund (thirty thousand crore) 7, 11
Gender gap in electives 7
Government-private gap 7
Intra-household allocation, gendered 7
Linguistic inequity 4, 7
Rural-urban gap 7, 11
Section 12(1)(c) of RTE Act 7
Tribal children 7

Parents and Families

Academic honesty conversations 8
AI companions and adolescent mental health 8
Children under 8, no AI tools 8

Conversations with children, how to have 8
Less technologically fluent parent 8
Letter to parents 8
School, questions to ask 8
Screen time at foundational ages 8
Supervision vs agency across ages 8
Voice assistants in the home 8

Institutions and Governance

Cabinet note, drafting of 10, 11
CEO of the Mission (educator-first) 10
Chief Ministers on Governing Council 10
Directorates, four (Curriculum, Teacher Capacity, Infrastructure, Equity) 10
Evaluation and Equity Directorate independence 10
Federal question 10, 11
Governing Council composition 10
Independent members (four tests of independence) 10
Legal instrument (NAIEM Act) structure 10
Mission, statutory case for 10
National AI Education Mission 10
NCERT, relationship with Mission 10
Parliamentary review at year 10 10
Procurement framework, Chapter X of Act 10
Statutory protection against revision 10
Sunset clause (15 years) 10

Fiscal Architecture

Capex / opex split considerations 11
Centrally Sponsored Scheme model 11

Compensation commitment (60k cr/year steady) 6, 11
Devolution via Finance Commission 11
Envelope (120,000 crore over 5 years) 11
Finance Commission, 15th and subsequent 11
International comparison (Estonia, Singapore) 11, 14
International development finance, cautious role 11
Outcome-linked grant mechanism 11, 13
Performance-linked grants, cautions 11
Philanthropic co-investment, governance 11
Political economy of commitment 11
Reallocation from existing programmes 11
Ring-fencing the allocation 11
Samagra Shiksha, reallocation from 11
Sources, five 11
State fiscal implications 11
Window argument (teacher cohort, tech trajectory) 11

Delivery and Execution

Decision rights (RACI) 13
Direct-to-school digital delivery 12
Failure modes, ten 13
Headmaster's role and agency 12
KPIs at three levels 13
Named coordinators at state and district 12
Outcomes, seven national 13
Phased rollout, five years 13
Planning vs delivery capacity 12
Rapid Response Unit 12
Seven administrative layers 12
State sequencing (strong vs weak capacity) 13

Legislation, Regulation, and International Comparison

Children's AI Safety Code 13, 5

Comparisons: what to copy, what not to 14

Estonian digital transformation 14, 11

Legislative layer 13

Singapore centralised model 14

South Korean teacher competitiveness 14

China, transferable and non-transferable elements 14

The Closing Return

Manifesto 15

Political leadership, specific demands of 15, 13

Returning to Sunita 15

What the book asks of its reader 15, Introduction



This index is a reader's guide rather than a scholarly apparatus. I have tried to include the concepts a reader is most likely to want to return to, at the level of specificity at which the book treats them. If a concept the reader expected to find is absent, or if a concept is listed under a theme other than the one she would have chosen, the omission is a limitation of the book's thematic organisation rather than of the concept's importance. Suggestions for improvement in subsequent editions are welcomed.

About the Author



Anil Pandey is a practitioner of education policy and implementation with experience across the Indian school system, the national policy apparatus, and the technology sector. He has worked with state governments, central ministries, civil-society organisations, and private-sector partners on questions of educational reform, teacher capacity development, digital public infrastructure, and the specific institutional arrangements through which reforms at national scale either succeed or fail. His earlier writing has addressed the intersection of technology and Indian public administration; this book is his most sustained engagement with education as a subject in its own right.

He has spent substantial time in classrooms across Uttar Pradesh, Kerala, Manipur, Tamil Nadu, Karnataka, Madhya Pradesh, Odisha, Bihar, and West Bengal, among other states, over the years preceding this book's composition. The book's argument draws on those classroom visits as much as on any policy document he has read.

He lives in New Delhi with his family. He can be reached through the publisher for correspondence related to this book, and he welcomes particularly the responses of teachers, institution heads, and state officers whose work the book is ultimately about.



A note. The biographical details above are offered in the conventional register of an author's note. Readers who wish more specific professional information about the author are directed to his publisher or to the public record. The author's principal stake in the book is the argument it contains, not the biographical credentials that surround it.