GROWTH UK | THINK TANK

The Case for Human Education

Why the right to a fully human life demands a wholly different education system

Growth UK is an independent think tank with no external funding, no commercial affiliations, and no political alignment. It works from first principles rather than from any pre-existing ideology or vested interest. This document was produced independently. It has no commercial sponsors.

PART A — THE FOUNDATION

Section 1 — The Moral Foundation

Before considering what education should look like, we must be clear about why it matters at all — and for whom. This is not a question with an obvious answer. Most education policy skips it entirely, moving immediately to curriculum, assessment, and structure as though the purpose were self-evident. It is not. And the failure to ask it is itself a large part of why reform has failed so consistently.

So let us ask it plainly. Why educate?

1.1 The Right to Be the Author of Your Own Life

Every human being has a biological and moral right to determine their own life.

Not the state. Not the economy. Not the technology of the age. Not their parents, their teachers, or the circumstances of their birth. Every person — regardless of where they start, what they inherit, or what the labour market requires of them at any given moment — has the right to be the author of their own existence.

That right cannot be exercised without capacity. A person who cannot think clearly, reason carefully, understand themselves and the world around them, or make genuinely informed choices is not free in any meaningful sense. They are simply subject to whatever forces act upon them. The forces of the market. The forces of whoever controls the information they receive. The forces of whoever designed the systems they inhabit without ever having chosen to.

Poverty of thought is as real a constraint on human freedom as poverty of income — a philosophical claim rather than a measured one, but one with a serious tradition behind it.*[2] It is simply less visible, less measured, and far less politically convenient to acknowledge.

> ***Education is therefore not preparation for work. It is not even preparation for citizenship, though it enables both. It is the deliberate process by which a human being develops the capacity to be genuinely free — to know themselves well enough, and the world well enough, to live deliberately rather than accidentally.***

This is a moral and human duty — on society, on the state, on every institution that touches a child's development — that exists entirely independently of economic conditions, technological change, or what the labour market requires at any given moment.

This moral position has always been important. It is now urgent. For the first time in human history, human thought is developing alongside a non-human intellect of rapidly increasing capability. A person without genuine self-direction is not merely unfulfilled — they are unprepared. The case for human self-determination is no longer only philosophical. It is existential.*

* *The philosophical basis for this position, the challenges to it, and why the emergence of non-human intellect makes it existential rather than merely philosophical, are addressed in Addendum 1.0 — The Moral Foundation: A Discussion.*

1.2 What the Current System's Ambition Reveals

The present education system has revealed ambitions. Not its stated ones — which are invariably lofty — but its operational ones, visible in what it measures, what it funds, and what it ignores.

As a society, we seem content that the majority of our children reach a reasonable standard in formal subjects. Mathematics to GCSE. English to GCSE. A clutch of other subjects that function primarily as gatekeeping mechanisms for the next stage of a system designed to sort rather than develop.

But for the real world — for the actual life that a human being must navigate across eighty or more years of existence — it is about achieving genuinely useful life skills. Not just for some. For all.

The gap between those two ambitions — a reasonable standard in formal subjects for most, versus genuine life skills for every single child — is the gap this document exists to close.

The scale of that gap is not abstract. A government-commissioned review published on 28 May 2026 estimated that young people not in education, employment or training — nearly one million of them in the UK right now, with projections rising toward 1.25 million — cost the economy approximately £125 billion every year. This figure, from the Milburn review commissioned by the government, is a model-based projection covering foregone productivity, welfare costs, foregone tax revenue, and long-term scarring effects. It aggregates multiple causal factors — mental health, the pandemic, employer behaviour, housing costs, and economic conditions — of which the education system is one among several. The education system is not solely responsible for this cost. But it is a significant contributing factor, and the investment case for reform is overwhelming even if education accounts for a fraction of it: the entire programme proposed in Section 11 costs less than £100 million.

1.3 Education as Enabling, Not Directing

Most people's lives are governed not by planned decisions — carefully considered paths chosen in full knowledge of the options — but by extraneous circumstances, near spur-of-the-moment choices, and the accumulated weight of what they were or were not given in their formative years.

The wider scope of education, including that given by parents and family and community, must therefore be enabling. It must provide the basic tools of life. Tools that allow each

individual — child and adult — to consider for themselves the path they wish to follow, rather than the one foisted upon them by society, by predicament, or by a system that made decisions about their potential before they had a chance to discover it themselves.

This is a fundamentally different concept of education from the one currently institutionalised. It is not about filling a vessel with the content judged appropriate by those in authority. It is about building the capacity to navigate. To think. To choose. To adapt. To contribute. To flourish.

Those capacities do not emerge from a curriculum designed around measurable outputs and institutional convenience. They emerge from a developmental environment designed around the actual needs of a human being.

1.4 An Uncomfortable Admission

There is an uncomfortable observation that belongs here, and this document will not shy away from it.

Those of us who have thought seriously about education reform — who have built careers, businesses, and lives that by conventional measures represent success — tend, when we turn introspective, to consider our own formal education as almost irrelevant to who we became. General experience, input from family, on-the-job learning, and self-directed inquiry were what actually mattered. Formal qualifications were, at best, door-openers. At worst, they were years spent memorising content that served no purpose beyond the examination that tested it.

This is not an argument against education. It is the sharpest possible argument for a different kind — one that deliberately produces what most people only ever find by accident.

1.5 Why This Moment Is Different

This argument is not new. Philosophers have made it for centuries. Reformers have made it for generations. What is new is the urgency.

We are at a specific and unrepeatable moment of convergence. The oldest failure in human society — the failure to provide every person with the conditions to develop their full humanity — is meeting the most powerful technological transformation in history. Tools are being built, right now, that could genuinely deliver on the promise of education as a human right. Tools that could provide every child, regardless of background, with access to world-class developmental support. Tools that could make the kind of education currently available only to the privileged few accessible to everyone.

Those same tools, if left to market forces and commercial incentives, will instead become the most sophisticated instruments of dependency, inequality, and control ever constructed. They will convert the process by which a human being becomes who they are into a revenue stream. They will deliver the appearance of personalised education while extracting data, manufacturing engagement, and deepening the grip of whoever owns the system on the people it purports to serve.

This document is written at the last moment at which democratic societies can meaningfully choose which of those futures they want. That window is open now. It will not remain open indefinitely.

> ***The question is not whether to educate. The question is whether education will be designed around the human being or around the systems that profit from shaping them. That choice, made now, will determine the character of human life for generations.***

1.6 The Conspiracy or Inertia Question

There is one more observation that belongs in this opening section, because it shadows every subsequent argument.

The failures described in this document are not new. The evidence has been available for decades.

The neuroscience has been accumulating for a generation. The alternative models — the schools that do things differently and get better results, the countries that start formal instruction later and outperform the UK, the programmes that develop critical thinking and produce measurably more capable people — have been visible and documented.

And yet the system has not changed. Reforms have been attempted, announced, partly implemented, and quietly abandoned. The essential architecture — the examination-driven, content-heavy, compliance-measuring machinery built in the nineteenth century for a world that no longer exists — has persisted.

Why?

One answer is structural inertia — the ordinary, grinding resistance of any large institution to its own reform. The people who run the system are overwhelmingly its products. They were rewarded by it. The assessment framework values the skills that the assessors themselves were valued for. The system selects for its own continuation. This is not malice. It is inertia.

But there is a harder question underneath the structural one. Systems that persist tend to serve the interests of those with the power to change them. An education system that produces genuinely critical, genuinely self-determining, genuinely capable thinkers in every child — regardless of background — is an education system that produces citizens who can think independently, evaluate evidence, and make genuinely informed choices about their own lives and their society. A system that produces such citizens is a better system — for those citizens and, in the long run, for the society they inhabit. Whether the failure to produce them is designed or simply allowed, the effect is the same.

This document does not claim conspiracy. It claims that the incentive structure produces outcomes identical to what a conspiracy would produce, and that this distinction, while philosophically important, is practically irrelevant to the children it affects.

What matters is that it changes.

KEY POINTS — Section 1

- Every human being has a moral right to be the author of their own life. That right requires capacity — the capacity to think, reason, understand, and choose. Education is the process by which that capacity is built.
- The current system's revealed ambition — a reasonable standard in formal subjects for most children — falls catastrophically short of this moral requirement.
- Education must be enabling, not directing. It must provide tools for life, not destinations prescribed by the state or the market.
- The admission worth making: most of those who have thought seriously about this considered their own formal education almost irrelevant to who they became. The argument is for a different education — not less of it.
- This moment is specific and urgent. The tools now emerging could deliver on the promise of education as a human right — or become the most sophisticated instruments of dependency and control ever built. The choice between those futures is available now and will not remain available indefinitely.
- The system's failure to change despite decades of evidence is not accidental. Whether structural inertia or the preservation of power, the effect is identical. What matters is that it changes.

Proceeds to Section 2 — A Voyage of Discovery

Section 2 — A Voyage of Discovery

This document does not have all the answers. That is not a disclaimer. It is the most important thing it has to say.

Every serious attempt at education reform in the United Kingdom has failed — not because the thinking was wrong or the money absent — but because every previous attempt pretended to know more than it did. It arrived with a blueprint, announced a transformation, measured the wrong things, and left the essential architecture unchanged.

This document takes a different approach. It is the opening statement of a structured, open, 25-year inquiry — a voyage of discovery that commits to asking the right questions, testing answers rigorously, and following the evidence wherever it leads, including to conclusions we have not yet imagined.

> ***The most important skill we can give a child is the ability to navigate uncertainty. It would be a remarkable contradiction to write a document arguing for that skill while pretending we already know the destination.***

2.1 The Discipline of Acknowledged Uncertainty

Serious inquiry requires distinguishing clearly between what is established, what is probable, what is unknown, and what is unknowable until you actually build the thing and find out. Education reform has almost never applied this discipline to itself.

Some things we do not know but could find out through research and properly designed pilots — whether specific approaches that work at small scale transfer to national systems, or how long it takes for early developmental improvements to show measurable outcomes in later learning.

Some things we cannot know until we build and test — whether a redesigned system coherently delivers what it promises, what unintended consequences emerge. There will be unintended consequences. There always are.

And some things we may never know — because the counterfactual cannot be run. Once a generation has passed through a new system, we cannot know with certainty what would have happened otherwise. That is not a reason *not* to build it. It is a reason to build with great care and genuine humility, and to keep asking whether it is working in the ways that actually matter.

This application of uncertainty runs through every section of this document. It is not apology. It is the standard we are arguing the education system should itself embody.

2.2 The 80-Year Human as Design Target

Every previous education reform has designed backwards from the wrong endpoint. GCSEs. A-levels. University entrance. Employment. These are not endpoints — they are early waypoints on a journey lasting eighty years or more.

A child starting school today will be alive, if current trends continue, well into the next century. They will navigate challenges, opportunities, and choices no one today can describe. The only meaningful goal is not a qualified school-leaver but a person capable of navigating whatever they encounter across a full human life. A person who:

- Can manage their body and mind throughout their life — maintaining the physical and psychological foundations that underpin everything else
- Has the capacity and disposition to keep learning — curiosity intact, intellectual confidence unbroken, the habit of inquiry maintained
- Can think critically — evaluating evidence, questioning received wisdom, and reasoning carefully to their own conclusions. This is not a luxury skill. It is the foundation of every advance in human knowledge, democratic governance, and social progress. Its absence is why so many are so easily misled.

A note on what critical thinking requires: critical thinking cannot float in empty space. You cannot reason carefully about nothing. You cannot evaluate evidence without knowing enough about a domain to recognise what counts as evidence. Expertise depends on deep stored mental models. Memory still matters enormously. Fluency frees cognitive bandwidth. The best-performing education systems in the world — Finland, Singapore, Estonia — did not achieve their results by abandoning rigorous knowledge acquisition. They achieved them by combining knowledge with reasoning, wellbeing, autonomy, and developmental support. The argument in this document is not against knowledge. It is against knowledge divorced from narrative, context, and the capacity to use it. Bad balance is the problem. Not knowledge itself.

The practical expression of this principle is what might be called the bedrock facts concept. A small number of genuinely foundational ideas — the electron and its interactions, the structure of DNA, the origin of the elements in stars, the nature of energy, the common ancestry of all life — are the roots from which everything else grows. These are not less rigorous than a chemistry syllabus full of equations. They are more rigorous, because they are the foundations from which everything else can be derived. A child who understands what an electron is and does can infer most of chemistry. A child who understands what DNA is and how it works can derive most of biology. A child who understands where energy comes from can understand every major challenge facing human civilisation.

The difference between a child who has memorised the periodic table and a child who understands that all elements heavier than iron were forged in dying stars — and can feel the extraordinary fact that the iron in their blood was made in a supernova — is not a difference of rigour. It is a difference of whether the knowledge is alive or dead. Alive knowledge connects to everything. Dead knowledge connects to nothing except the exam that tests it. The question is therefore not knowledge or skills. It is what knowledge, in what order, told as what story. A coherent narrative of how the universe works — from the Big Bang to hydrogen to stars to elements to chemistry to DNA to life to humans to society — gives a child the foundation for everything else. The civilisation model is one way to deliver such a narrative. There are others. Defining and testing the most effective approaches — for different ages, different contexts, different subjects — is precisely the kind of inquiry the governance body should be commissioning from year one.

- Knows themselves — their mind, their emotions, their motivations, and their physical needs well enough — to regulate and develop them

- Can interact socially with genuine skill and empathy — forming real relationships, reading others accurately, and navigating the human world in an age when social media has degraded these capacities in a generation
- Can collaborate — working with others toward shared goals across difference and disagreement
- Can adapt — entering new domains, acquiring new skills, functioning effectively under uncertainty, repeatedly
- Has a sense of values and purpose — enough ethical grounding and self-knowledge to make deliberate choices about how to live

This is not a utopian specification. It is a description of what thoughtful people, reflecting on their own lives, identify as the qualities that have actually served them. The tragedy is that it is almost entirely absent from any examination syllabus ever written.

> ***Ask not what a child needs to know at sixteen — but what foundations they need to build the life they choose to live, all the way to a fulfilled life at eighty. Work backwards from there. Everything else follows.***

2.3 The Structure of This Inquiry

At the end of this voyage we will have better questions, a clearer map of the territory, and a governed process for finding out what we do not yet know. That is what this document provides.

KEY POINTS — Section 2

- This document does not have all the answers. That is not a weakness — it models the thinking we are arguing education should produce.
- Every previous education reform failed partly because it pretended to know more than it did. This document takes a different approach: a structured 25-year inquiry committed to following evidence wherever it leads.
- Known unknowns fall into four categories: what we could find out, what we can reasonably infer, what we cannot know until we build and test, and what we may never know. Each demands a different response.
- The design target is not a qualified school-leaver. It is a person capable of building the life they choose to live, all the way to a fulfilled life at eighty. Work backwards from there.

Proceeds to Section 3 — The Future We Cannot Predict

Section 3 — The Future We Cannot Predict

What is new is the speed. And what is alarming is not the change itself but the specific nature of what is now changing, and the specific moment at which it is arriving.

The labour market has always changed. New industries have replaced old ones, new skills have superseded others, and each generation has found that the world of work it inherited looked different from the one that followed. Previous technology replaced physical labour, then routine cognitive labour — the tasks that followed clear rules and could be expressed as procedures.

Artificial intelligence replaces something categorically different: judgement. Analysis. The synthesis of complex information. The generation of language, argument, and code. The capacities that formal education has always treated as its highest products and its most reliable gateway to economic security are precisely the capacities now being automated — faster than any education system has yet acknowledged.

But there is something more fundamental than the speed. Every previous technology was built by humans who understood what they had built. The steam engine, the printing press, the transistor, the internet — each was complex, each had unintended consequences, but each was ultimately explicable. We knew what it did and why.

The AI systems now being deployed at scale — including the system that contributed to writing this document — are not fully understood by the people who built them. That contribution is itself an illustration of the argument: the ideas, judgements, values, and direction in this document are human. The AI accelerated the work. That is precisely the distinction between augmentation and replacement that the document argues must be preserved. The outputs can be observed. And yet the internal processes that produce them cannot be fully explained, even by their creators. That fact alone should give us pause.

Humanity has, for the first time, built a machine it does not understand. The consequences of that fact are genuinely unknown. We are not managing a known risk. We are facing an open question about the nature of human civilisation itself — in real time, without a map, and with the decisions already being made by those who have not stopped to think.

3.1 This Is Not the First Time

It is worth pausing here, because the alarm is real but the perspective is important. Cognitive outsourcing through technology is not new. Writing atrophied the extraordinary oral memory that pre-literate cultures maintained — entire cosmologies, genealogies, and histories held in human minds, without marks on surfaces. Socrates saw it coming. In the Phaedrus, he argued that writing would produce the appearance of wisdom without the reality — people would read many things without proper instruction and seem to know much while actually knowing little. He was laughed at. He was also partially right.

Then the printing press changed what it meant to be educated. Then the calculator made mental arithmetic seem unnecessary. Then the internet made memorisation look redundant. Each time, a cognitive capacity that humans had previously exercised internally was outsourced to a tool. Each time, something was genuinely lost and something was genuinely

gained. Each time, the people warning about the loss were called luddites. Each time, they were at least partially right.

The difference now is one of scope and depth. Previous tools outsourced specific cognitive functions. The AI systems being built today are present across the whole of cognitive life — writing, reasoning, researching, coding, analysing, deciding. The potential scope of outsourcing is total in a way it has never been before. And we already have early evidence of the atrophying. Students who use AI to produce work they do not understand. Professionals who can no longer write clearly without AI assistance. A generation learning to prompt rather than to think.

The Socrates warning applies with more force now than it ever did. We do not yet know whether AI augmentation leads to human flourishing or a gradual hollowing out of what makes us fully human. The difference between those two futures is not the technology. It is what we build the education system to accomplish.

> ***Every previous wave of cognitive outsourcing freed human minds for things the tools could not do. The question is whether AI will do the same — or whether, for the first time, the scope is wide enough to leave nothing distinctively human behind. We do not know. But we must design the best response we can, while we still have human foresight to guide us.***

3.2 Two Futures, One Choice

There are two futures now being simultaneously built, and the choice between them is being made right now — largely by default, largely without democratic deliberation, and largely without the people most affected having any say.

The first future starts not with the technology but with the human being. We know, with reasonable confidence, what a fully developed person looks like at the end of a twenty-five year educational journey. Curious. Critical. Self-knowing. Capable of relationship, of adaptation, of purpose. That destination is defined. The question this future asks is precise and disciplined: can this technology help us get there, and if so, how, and under what conditions? Not what can AI deliver and how might humans adapt to it — but what do humans need, and can AI serve that need without distorting it?

If it helps us get there, we use it. If it does not, we find another way. What it must never do is alter the destination itself — and that destination is already defined: a human being who can think critically, know themselves, relate to others with genuine skill, adapt across a lifetime, and live deliberately rather than accidentally. Not a more efficient worker. Not a more compliant citizen. A more fully human person. That is what this journey is for. Everything else is a means to that end, including the technology.

The second future is the one we sleepwalk into. In it, the machine becomes the destination — not by design, but by default. Human development is no longer the goal — it is the raw material. The same technology is deployed by commercial entities whose incentive is dependency, not development. The human being is no longer the author of their own life. They are the product. And the process by which a child becomes who they are becomes, simply, a revenue stream.

This is not speculation. It is the logical endpoint of the same commercial incentive structure that gave us social media — applied not to attention, but to identity. Not to what you watch, but to who you become.

The people building it are not, in the main, malevolent. The incentive architecture is. A system that knows your child better than you do, that your child cannot imagine functioning without, that accumulates irreplaceable knowledge of who they are and how they think — built by a company whose revenue depends on continued subscription and engagement — is not a mentor. It is the most powerful dependency mechanism ever conceived, wearing the clothes of education.

Steinbeck wrote of the machine that must be fed. Nobody controls it. It simply demands. The commercial deployment of AI in human development is that machine, and it is already in motion.

3.3 The Irony This Document Cannot Avoid

This document was written with AI. The AI that helped articulate this warning is itself a commercial product, built by a company with investors, a valuation, and a revenue model.

That irony is not a disclaimer. It is placed here deliberately, because it is the sharpest possible illustration of this argument. The distinction that matters is not between AI and no AI. It is between AI that serves human development and AI that feeds on it.

The first is defensible. The second requires a legislative framework that prevents it — one that places developmental data ownership with the learner, requires full transparency from any commercial provider operating in educational settings, and prohibits business models whose revenue depends on the dependency of children. That framework is described in Part B.

3.4 The Window

Democratic societies can still choose which of these futures they want. That remains true, as this document is written. But the window is closing fast. Not in decades, but today.

Once the commercial infrastructure is built and the dependencies are established — once a generation of children has grown up with a corporate AI companion they cannot imagine living without — the question of who controls it becomes academic. The answer will already be written in the terms of service nobody read, in the data contracts nobody scrutinised, and in the cognitive habits of an entire generation shaped by systems designed for retention rather than development.

The US platform companies are building now. China is building an AI-integrated education system as an instrument of national formation — well-resourced, and designed to produce citizens aligned with the state's priorities rather than their own. Western democratic governments are arguing about regulation, moving slowly, and almost entirely unprepared for the speed of what is arriving.

The vacuum is being filled. The only question is by whom, and with what values at the centre.

The disengagement pipeline is not a future risk. It is the present reality. In early 2026, close to one million young people in the UK were not in education, employment or training — a figure rising, not falling, and placing the UK in the bottom six of all OECD countries. These

are the children the current system has already lost. The question is how many more we are prepared to lose before we act.

Humans must control this technology. That is not an aspiration. It is the condition on which everything else depends. We use the tool. The tool does not use us. And the reality, right now, is that the tool is using us — and that this has been allowed to happen without decision, without debate, and without the people most affected having any say.

The children who pass through the next generation of schools will either be equipped — curious, critical, self-knowing, psychologically grounded — to navigate whatever the AI-shaped world becomes. Or they will be shaped by forces they never chose and cannot see. That choice is what the rest of this document is about.

KEY POINTS — Section 3

- AI is not a faster version of previous technological change. It is automating judgement, analysis, and synthesis — the very capacities education has always treated as its highest products. A curriculum built around knowledge transmission has no adequate answer to this.
- Cognitive outsourcing is not new. Writing, printing, calculators, the internet — each time a capacity was outsourced, something was lost and something gained. The Socrates warning applies now with more force than ever. What is different this time is scope: the potential outsourcing is total, not partial.
- Two futures are simultaneously possible. In one, AI serves human development. In the other, commercial deployment creates the most powerful dependency mechanism ever built — applied not to attention but to identity.
- The evil is structural, not personal. The incentive architecture of commercial AI makes exploitation the inevitable outcome regardless of the intentions of those building it.
- This document was written with AI. That irony is deliberate. The distinction that matters is between AI as a tool in service of human development and AI as a system designed to embed itself in development and extract value from dependency. The first is defensible. The second must not be permitted.
- The window for democratic choice is open and closing. The US is building. China is building state-formation education AI. Democratic governments are not moving at the required speed. The vacuum is being filled. The only question is by whom.

Proceeds to Section 4 — What Human Development Actually Requires

Section 4 — What Human Development Actually Requires

Human development is a biological process. It has specific physical and neurological requirements that no amount of curriculum design can substitute for. This section is about those requirements — and about why the current system ignores almost all of them.

4.1 Learning Is Biological

The desire to learn is not something education instils. It is something every child is born with. The survival of higher animal species depends on it — the young that are curious, that observe, that imitate, that experiment and adapt, are the ones that survive. Human children are the most learning-oriented creatures on earth. They arrive with insatiable curiosity, extraordinary openness, and a biological drive to make sense of the world around them.

Watch a child of three. They ask *why* incessantly. They try things, fail, try differently. They absorb language, social cues, physical skills, and cause-and-effect relationships at a rate no adult can match. They do not need to be motivated to learn. They need to be encouraged to do it.

But this is only half the picture, and the document would be incomplete without the other half. Children are also naturally distractible, impulsive, novelty-seeking, and inconsistent. Learning difficult things — mathematics, writing, sustained argument, complex skills — often requires repetition, delayed gratification, structure, and sometimes externally imposed effort. Intrinsic motivation is the most powerful driver of long-term learning, but it cannot do everything alone. The developmental argument is not a case for pure freedom. It is a case for the right balance of freedom and structure — enough structure to support the acquisition of genuinely difficult capabilities, enough freedom to ensure that acquisition is connected to genuine understanding rather than compliance. The current system has too much of the wrong kind of structure — compliance-based, anxiety-driven, externally ranked — and too little of the right kind: the disciplined, supported, purposeful effort that produces real mastery.

The current system demotivates them with remarkable efficiency. By the time most children are in secondary school, the natural curiosity and creativity that arrived at age four has been substantially replaced by compliance, anxiety about failure, and a definition of learning that means sitting still and reproducing what the teacher said. The system did not set out to do this. But it does it, consistently, at scale, and has done so for generations.

We do not need to teach children to want to learn. We need to stop teaching them not to.

4.2 The Three Foundations

Modern neuroscience has confirmed and significantly extended what developmental psychology long suspected. Human development is not simply a matter of filling a mind with content. It is a biological, social, and psychological process with specific preconditions. When those preconditions are met, development is remarkable. When they are not, the consequences compound across a lifetime.

Three things are now supported by strong and consistent evidence across multiple disciplines. All three are routinely ignored by the current system.

Foundation One — Body and Brain Are One System

A child who is hungry, cold, sleep-deprived, or living under chronic stress is not in a neurological state that supports learning. This is not a social observation. It is a biological one.

Chronic stress floods the developing brain with cortisol. In sustained doses this does not merely impair concentration — it structurally affects the development of the prefrontal cortex, the part of the brain responsible for reasoning, planning, and self-regulation. The damage is not metaphorical. Sleep deprivation impairs memory consolidation, emotional regulation, and executive function — and teenagers are structurally sleep-deprived by school start times designed around adult convenience rather than adolescent neurology. Poor nutrition directly impairs cognitive performance and mood stability. Physical movement is not a break from learning — it is neurologically necessary for it.

A system that ignores all of this and then measures the results is not assessing children's potential. It is measuring the consequences of its own neglect.

The consequences of that neglect are now visible in the data. Mental health conditions are the fastest-growing cause of youth inactivity in the UK — a finding confirmed by a government-commissioned review in 2026. But the mental health crisis among young people is not a separate phenomenon from the educational system. It is partly produced by it.

Chronic assessment anxiety, the absence of psychological safety, persistent experiences of failure in environments that offer no alternative pathway, and the systematic neglect of the body-mind connection — these are not just barriers to learning. They are causes of the deterioration that the system then labels as a welfare problem and passes to another department. A 2026 review found that 30% of disabled and mentally ill young people aged 16–24 were not in education, employment or training, compared to 8% of their peers. The system creates the conditions for that outcome and then expresses surprise at the number.

Foundation Two — Belonging Is Not Optional

Attachment, stability, and the experience of being genuinely known by at least one consistent adult are not social luxuries. They are preconditions for the psychological safety that allows a child to take intellectual risks, make mistakes, and learn from them without being destroyed by the experience.

Social neuroscience has established that belonging is a basic human need with direct neurological correlates — not a nice-to-have that follows once academic performance is secured. A child who does not feel safe, seen, and connected cannot access the higher cognitive functions that learning requires. The threat response — fight, flight, freeze — is incompatible with curiosity and open exploration. You cannot be simultaneously defensive and genuinely intellectually open.

The factory model of education — large classes, rotating subject teachers, no continuity of relationship, assessment designed to produce comparative rankings rather than genuine feedback — is almost perfectly designed to prevent the conditions that connection requires. This is not deliberate. It is the accumulated consequence of a system built for industrial-scale processing, not human development.

Foundation Three — Purpose Drives Everything

The old developmental model suggested a strict hierarchy: meet physical needs first, then safety, then belonging, then self-worth, and only then the pursuit of meaning and purpose. Modern understanding is more fluid, and more useful.

Purpose and meaning do not sit passively at the top waiting to be reached. They feed back downward through everything below them. A child deeply engaged in something that genuinely matters to them — something they have chosen, something connected to who they are — can tolerate difficulty, sustain attention through frustration, and persist through failure in ways that no amount of external motivation can replicate. Purpose is not the reward for getting the foundations right. It is an active ingredient in building them.

This means the system must work in both directions simultaneously. Bottom up — meeting the physical, safety, and relational foundations without which nothing else is possible. And top down — finding and nurturing the genuine interests, curiosity, and sense of purpose that give a child a reason to engage with the whole enterprise. The current system does neither reliably. It assumes both are already present and proceeds to test for content regardless.

The question is not what should we teach. It is what does this particular child need in order to develop — and are we providing it? Everything else is detail.

4.3 The System Was Not Designed — It Accumulated

It would be convenient if the failure to meet these three foundations were the result of ignorance. Some of it is. But the deeper explanation is structural.

The current education system was not designed by anyone. It accumulated. It began in the nineteenth century as a mechanism for producing a literate, numerate, compliant industrial workforce and for sorting children into appropriate stations. The examination system, the subject silos, the rows of desks, the teacher at the front, the bell that ends one topic and begins another regardless of where any child is in their understanding — none of this was designed around human development. It was designed around institutional management, social stratification, and economic utility.

Each subsequent reform has adjusted the surface without touching the architecture. The core machinery — content transmission, examination, ranking — has persisted because it serves the interests of those who were selected by it and who now run the institutions that perpetuate it. Not by conspiracy. By the ordinary, self-reinforcing inertia of any system that selects for its own continuation.

That explanation does not make the failure acceptable. It makes it explicable. And understanding it correctly is essential — because a reform that does not address the architecture will fail as every previous reform has failed, regardless of how well-intentioned or well-funded it is.4.4 What This Means for the System We Are Proposing

The three foundations described in this section are not aspirations to be pursued once the academic curriculum is sorted. They are the preconditions without which no curriculum — however well designed — will reliably work. They must therefore be treated as the first design requirement of any reformed system, not as supplementary provision for children with identified needs.

This has practical consequences at every level — for how schools are structured, how teachers are trained and deployed, how time is allocated, how the physical environment is designed, how assessment works, and how the system responds to the reality that children arrive with vastly different levels of these foundations already in place.

KEY POINTS — Section 4

- The desire to learn is biological. Every child arrives with it. The current system does not create motivation — it destroys it. Reform begins with understanding that.
- Foundation One: body and brain are one system. Hunger, chronic stress, sleep deprivation, and lack of movement do not merely hinder learning — they structurally impair the developing brain. The mental health crisis among young people is not separate from the educational system. It is partly produced by it.
- Foundation Two: connection is a developmental necessity. Psychological safety — the experience of being genuinely known and not at risk of humiliation — is a neurological precondition for the kind of open, curious, risk-taking engagement that real learning requires.
- Foundation Three: purpose drives everything. A child with a real reason to learn can tolerate difficulty in ways external reward cannot produce. Finding that reason is one of the system's core jobs. Currently it is not a job the system recognises at all.
- The system was not designed to fail. It accumulated — layer by layer, over 150 years, for purposes that had nothing to do with human development. Surface reform will not work. The architecture must change.
- The three foundations are not supplementary provision. They are the first design requirement. Everything else — curriculum, technology, assessment — is built on top of them.

The Case for Human Education

Why the right to a fully human life demands a wholly different education system

PART B — THE METHODOLOGY

Section 5 — Why Every Previous Reform Failed

The starting point for any serious attempt at education reform starts with understanding why every previous attempt has failed. Not as a counsel of despair — but because the failure pattern is consistent, the reasons are structural, and unless this attempt is explicitly designed to avoid those reasons it will fail for exactly the same ones.

The pattern is identical each time. A new government arrives with a mandate and a white paper. Reforms are announced with confidence. Pilots are run selectively, usually in favourable conditions. Implementation is partial. A new government arrives. The reforms are renamed, redirected, or quietly abandoned. The essential architecture — examination-driven, compliance-measuring, industrially organised — persists unchanged.

Six structural reasons explain this pattern. All six must be addressed explicitly in the design of any reform that intends to be different.

5.1 The Reform Horizon Problem

Political cycles run four to five years. Genuine educational change — change that shows up in the developmental outcomes of children and the life trajectories of young adults — takes fifteen to twenty-five years. These timescales are structurally incompatible unless the reform process is explicitly insulated from electoral politics.

No government will invest seriously in a programme whose results will be claimed by a successor. No opposition will sustain commitment to a reform it did not initiate. Without a cross-party, statutorily independent governance body with a mandate that survives individual governments, every reform will be abandoned before it can demonstrate results.

5.2 The Measurement Trap

Every previous reform has measured what was easy rather than what mattered. The system optimised for the measure and destroyed the thing the measure was supposed to represent. Exam results rose while genuine understanding fell. Ofsted grades improved while teacher morale collapsed. Attendance figures were managed while the reasons for absence were ignored.

Goodhart's Law — when a measure becomes a target it ceases to be a good measure — has operated without exception in UK education for fifty years. The governance body must build a measurement framework that is resistant to this dynamic, which means measuring things that are difficult to game, tracking them longitudinally, and holding the north star constant regardless of short-term political pressure.

5.3 The Implementation Gap

Policy designed at national level hits a system of approximately 24,000 schools, each with its own culture, leadership, history, and constraints. Top-down mandates produce compliance theatre rather than genuine change. Schools learn to perform the reform for inspection purposes while continuing to operate as before. The gap between the policy as written and the policy as lived is routinely vast and routinely ignored by those who commissioned it.

The governance body must work with the grain of school-level autonomy rather than against it. The most effective interventions in education are not mandates but the creation of conditions — incentives, tools, evidence, support — that make it easier for schools to do the right thing than the wrong thing.

5.4 The Interest Capture Problem

Every oversight body, exam board, academy chain, teacher union, and university education department has institutional interests in the existing system. Reform that threatens those interests gets diluted, delayed, or defeated — not through conspiracy but through the ordinary operation of institutional self-preservation.

The governance body must be constitutionally protected against this capture. Its membership must exclude those whose institutional interests are served by the status quo. Its mandate must be explicit and statutory. Its independence must be structural — not dependent on the goodwill of those it is meant to hold accountable.

5.5 The No-Learning Problem

Previous reforms did not build in genuine mechanisms for learning from failure. When pilots did not work they were quietly discontinued. When reforms produced unintended consequences they were rebranded rather than analysed. When evidence contradicted the policy direction it was ignored or suppressed.

The governance body must treat acknowledged failure as more valuable than manufactured success. Every pilot that does not work tells the system something it needs to know. The credibility of the entire enterprise depends on the governance body's willingness to report failure as prominently and as rigorously as it reports success.

5.6 The This-Time-Is-Different Problem

Every previous reform claimed to have learned from its predecessors. None built the structural safeguards that would actually make it different. Good intentions are not sufficient. Personal commitment is not sufficient. A new minister with a genuine passion for education is not sufficient.

What is sufficient is a governance architecture that makes capture, abandonment, and the measurement trap structurally difficult — not just undesirable. The rest of Part B is that architecture.

KEY POINTS — Section 5

- Six structural reasons explain why every previous UK education reform has failed: the reform horizon problem, the measurement trap, the implementation gap, the interest capture problem, the no-learning problem, and the this-time-is-different problem.
- None of these is accidental. All six are predictable consequences of attempting to reform a large, complex, entrenched system through short-cycle political processes.
- The governance architecture proposed in the following sections is specifically designed to address all six. Good intentions are not sufficient. The structure must make failure difficult.

Proceeds to Section 6 — The Strategic Governance Body

Section 6 — The Strategic Governance Body

The central structural proposal of this document is not a curriculum. Not a technology platform. Not a new inspection framework. It is a governance body — permanent, statutory, independently constituted, and designed from the outset to avoid every failure mode identified in Section 5.

This body is the mechanism through which the voyage of discovery described in Part A becomes a governed, accountable, democratic process rather than a drift. Everything else in Parts B and C depends on it existing and functioning as specified.

> ***The purpose of the governance body is not to run education. It is to ensure that the system learns — continuously, openly, and at the pace that genuine improvement requires.***

Why a governance body — and why nothing else will work

The case for this governance body does not begin with the design of the institution. It begins with the failure of everything that has been tried without one.

The UK has attempted education reform through normal ministerial channels for fifty years. Every government has had its white paper, its flagship initiative, its rebranding of the previous government's abandoned programme. The result is the system described in Section 5 — and a NEET crisis costing £125 billion per year, rising, with no credible mechanism for sustained improvement in sight. That is not the cost of an imperfect system doing its best. It is the quantified cost of fifty years of short-cycle reform producing long-cycle failure.

Before proposing the governance body, it is worth examining what the alternatives are — because the proposal only earns its place if the alternatives are genuinely inadequate.

Alternative 1 — The Singapore model: sustained ministerial direction

Singapore has achieved remarkable and sustained educational improvement through a close tripartite relationship between the Ministry of Education, the National Institute of Education, and schools — with systemic changes spreading consistently across the system over five decades. This works because Singapore has had continuous one-party government since 1965, allowing long-term education policy planning that is simply unavailable to a democratic multi-party system with regular changes of government. The Singapore model requires political continuity the UK does not have and cannot manufacture. It is not transferable.

Alternative 2 — The Finland model: independent national agency alongside ministry

Finland operates a statutory National Agency for Education that sits alongside but independently of the Ministry, developing curriculum objectives and methods and monitoring system effectiveness — established in 1991 and progressively extended to a ten-year

strategic planning horizon. This is close to what this document proposes, and it validates the concept. The critical difference is that Finland's system operates within a high-trust, politically stable, socially cohesive environment with strong cross-party consensus on education as a national priority. The UK equivalent requires stronger statutory protections to survive the political turbulence that Finland's system does not face. The Finnish model is the right direction. The UK version needs more structural protection.

Alternative 3 — More of the same: ministerial reform within existing structures

The third alternative is to continue as now — ministerial white papers, reform cycles, and the hope that this government's initiative will succeed where the last fifty years have not. The evidence against this option is not theoretical. It is the system we have. High levels of ministerial turnover, poor coordination across departments, and the structural inability to set long-term strategy are documented features of the British civil service approach to education. They produced the NEET crisis. They will not solve it.

The question is therefore not whether the governance body is politically difficult. It is what the alternative is. There is no credible fourth option.

The British precedent — this has been done before

The governance body proposed here is not a constitutional innovation. It is the next application of a well-established British institutional model — one the UK has used successfully for precisely this kind of problem: long-term technical decisions that electoral cycles are structurally incapable of managing.

The closest and most directly relevant precedent is the Committee on Climate Change, created by statute in 2008. Composed of experts chosen for individual merit rather than institutional affiliation, with a clear statutory mandate, direct accountability to Parliament, and sufficient powers to hold government to account, it has operated successfully across multiple governments with entirely different political orientations. It is now cited as a model internationally — copied by eighteen countries in their climate legislation. Crucially, it advises against a long-horizon target, publishes its findings publicly, and makes the cost of ignoring its recommendations visible and permanent. That is precisely the function the Education Governance Body must perform.

The Bank of England's operational independence, granted in 1997, established the broader principle: that certain decisions with long-term consequences — ones no government can be trusted to make impartially about itself — are better placed in the hands of an expert body accountable to Parliament rather than to the minister of the day. David Cameron explicitly invoked this precedent when arguing for the Climate Change Committee: 'In economic policy, everyone can see that independence for the Bank of England has worked. We now need a Bank of England moment when it comes to climate change.' The Bank analogy motivates the structure. The CCC is the working model.

The Office for Budget Responsibility adds a cautionary note worth heeding. It has made fiscal forecasting more honest and harder to politicise — but it has also faced political friction whenever its analysis proved inconvenient to the government of the day. That is a realistic preview of what an Education Governance Body flagging a government's policy as developmentally damaging would face. The statutory protections described in Section 6.1 are designed precisely for that friction, not as a theoretical safeguard but as a practical necessity.

The Office for Budget Responsibility performs the same function for fiscal policy. The research councils do it for scientific investment. The infrastructure planning inspectorate does it for major capital decisions. In each case, the British state recognised that certain decisions are too consequential, too long-term, and too vulnerable to short-cycle political pressure to be left to ministerial discretion alone.

Education meets every one of those criteria and exceeds them. The decisions required — what to test, how to measure genuine development, which technologies to approve, what the system should produce — are technical, long-term, and cannot be made impartially by governments whose electoral survival depends on short-term results. The annual cost of getting these decisions wrong is £125 billion. The annual cost of running the Committee on Climate Change is approximately £10 million. The entire setup cost of the proposed Education Governance Body is a fraction of one year's NEET bill.

The UK already does this for money and for climate. The question is whether it is willing to do it for children. That is not a rhetorical question. It is the central political test of this proposal.

6.1 Constitution and Independence

The governance body is established by primary legislation. Its existence, mandate, and independence cannot be altered by a single government acting alone — change requires parliamentary consensus. This is the fundamental structural protection against the reform horizon problem.

It reports to Parliament, not to the government of the day. Its annual reports are published in full, without ministerial editing, and are subject to parliamentary scrutiny. The Education Secretary receives its findings and recommendations but does not control its agenda or its conclusions.

It is constitutionally prohibited from being abolished, merged, or substantially redirected by any government within its first twenty years of operation without a two-thirds parliamentary majority. This is the minimum protection required to make a 25-year mandate meaningful.

6.2 Membership — Experience Over Credentials

A note on democratic accountability

Education is not a purely technical domain. It involves values, culture, identity, and moral questions that cannot be fully delegated to experts. Who decides what flourishing means? What values should the system transmit? How should competing conceptions of the good life be handled? These are not technical questions with expert answers. They are democratic questions that require democratic participation.

The governance body is designed to be independent of short-cycle political pressure — not independent of democratic accountability. The distinction matters. The Bank of England is operationally independent but democratically mandated and publicly accountable. The Committee on Climate Change advises and scrutinises but does not govern without accountability. The Education Governance Body operates on the same principle: independent in its evidence gathering, its pilot evaluation, and its long-term strategic direction; accountable to Parliament and to the public for the values that guide its work.

Concretely, this means the governance body's mandate must include a public participation mechanism — periodic national consultations on the fundamental question of what

education is for, with results that genuinely inform the body's strategic priorities. Not a tokenistic engagement exercise but a serious, structured process for ensuring that the values embedded in the system reflect democratic consensus rather than expert preference. The governance body discovers what works. The democratic process shapes what it is trying to produce. These are different functions and both are necessary.

The single most important design principle for membership is this: selection is based on demonstrated experience and proven outcomes, not on educational credentials, professional status, or institutional affiliation.

The education establishment — professors of education policy, retired senior civil servants, former ministers, the great and the good — is explicitly not the primary pool from which members are drawn. These people are, almost without exception, products of the system they would be asked to reform. They share assumptions, know each other, and have institutional interests in continuity.

The membership should include people who have:

- Run a school in a challenging context and demonstrably improved outcomes for disadvantaged children
- Built technology that actually works in educational environments at scale
- Designed and delivered training programmes with measurable results in business or public service
- Conducted applied neuroscience or developmental psychology research with direct practical implications
- Experienced the system from the receiving end — young people and recent school leavers, particularly those from disadvantaged backgrounds
- Run organisations through significant change management successfully
- Worked in international education systems that do things differently
- Parents — not as token representatives but as people with direct, current experience of what the system does to children

What the membership explicitly should not be dominated by: academics whose careers are built on education policy theory, people whose institutional roles depend on the current system, political appointees of any stripe, and anyone who has served on more than two previous government education advisory bodies.

6.3 Rotating Tenure and the Renewal Principle

All senior roles — chair, deputy chairs, committee leads — have a maximum tenure of five years with a mandatory one-year handover period. No exceptions and no extensions. This is in primary legislation, not guidance.

At any given time, at least one third of the governance body is in its first two years of tenure. This ensures continuous injection of new thinking and prevents the gradual calcification that afflicts every long-lived public institution.

The one-year handover is not optional. Institutional memory — the understanding of what has been tried, what has failed, what the evidence base shows, what relationships matter — must be deliberately transferred rather than lost with each rotation. Outgoing members are required to produce structured handover documentation and to spend the final year working alongside their successors.

No one who has served on the governance body may serve on the selection body, and vice versa. The two bodies must never share personnel or become mutually self-selecting. This is the structural guarantee against the body becoming a closed circle.

6.4 The Independent Selection Body

Members of the governance body are selected by a separate, independently constituted selection body — itself operating on the same rotating tenure principles.

The selection body applies explicit criteria weighted heavily toward practical, demonstrated experience. Academic qualifications are neither required nor preferred. What is required is evidence of having done something that worked, in conditions that were genuinely difficult, with outcomes that can be measured.

The selection process is transparent and published. Every appointment is accompanied by a public statement of why the person was selected, what specific experience they bring, and what gap in the governance body's collective knowledge and experience they fill.

The selection body is explicitly required to seek cognitive and experiential diversity — not just demographic diversity, though that matters, but people who come from fundamentally different professional cultures, who hold genuinely different views about how learning works, and who will therefore generate real debate rather than comfortable consensus.

6.5 Powers

The governance body has statutory powers in four areas. Without these powers it is advisory and will be treated as such.

Commission and fund

Statutory authority to commission and fund pilot programmes directly — without requiring ministerial approval for individual pilots. The governance body has its own budget, protected in legislation, reviewed by Parliament not by the Treasury on an annual basis.

Demand transparency

Power to require any commercial provider operating AI systems, data collection tools, or learning platforms in UK schools to disclose, in full and within a specified timeframe: the nature of the data being collected, how it is stored and used, the contractual terms under which it is provided to schools, and the evidence base for any claimed developmental benefit. Failure to comply results in suspension of the right to operate in UK schools.

Set standards

Power to set and enforce the parallel metrics framework across all state-funded schools. Power to define the standards that any AI tool must meet before deployment in schools. Power to approve or reject tools against those standards. Power to suspend approval if a tool is found to have caused harm or to have failed its stated purpose.

Report and flag

Direct reporting line to the Education Secretary and to Parliament. The governance body can flag, publicly and in writing, any government policy that it believes will damage the developmental outcomes the metrics framework measures. It cannot veto government policy. It can make the cost of bad policy visible and permanent.

6.6 The Relationship to Existing Structures

The governance body does not replace Ofsted, regional directors, local authorities, exam boards, or any existing oversight structure. It sits above them as the coordinating intelligence that all currently lack.

Ofsted's role is fundamentally reframed. It becomes the operational field assessment arm of the governance body — not an adversarial inspection regime but a collaborative evidence-gathering network. This reframing requires specific institutional mechanisms. The governance body writes the inspection framework that Ofsted applies — replacing compliance-based judgements with developmental outcome measures. Ofsted's statutory accountability shifts: inspectors report to the governance body on developmental evidence, not to the Secretary of State on compliance scores. Inspector training is redesigned around the parallel metrics framework. The adversarial dynamic changes when what is being measured changes — schools cannot perform developmental outcomes in the way they can perform inspection compliance. Ofsted inspectors become the governance body's field researchers: reporting what is actually happening on the ground, what is working, what is failing, what is emerging. Their reports feed into the governance body's evidence base rather than feeding directly into school accountability judgements.

This is not the abolition of school accountability. It is its maturation. Instead of inspectors checking compliance with yesterday's standards, they are contributing to the development of tomorrow's understanding. The relationship between schools and inspectors changes from adversarial to genuinely collaborative — because both are now part of the same learning system.

Exam boards, academy chains, local authorities, and teacher training providers all report to and through the governance body on matters within its mandate. They retain their operational independence. They lose their current ability to operate without strategic coordination or accountability for developmental outcomes.

KEY POINTS — Section 6

- The governance body is established by primary legislation with a 25-year mandate. It reports to Parliament, not to the government of the day. It cannot be abolished or substantially redirected without a two-thirds parliamentary majority.
- Membership is based on demonstrated experience and proven outcomes — not credentials, status, or institutional affiliation. The education establishment is not the primary pool. People who have done things that worked in difficult conditions are.
- All senior roles have a maximum five-year tenure with a mandatory one-year handover. At least one third of the body is always in its first two years. No overlap between the governance body and the selection body.
- The body has four statutory powers: commission and fund pilots directly; demand full transparency from commercial providers in schools; set and enforce the parallel metrics framework and AI standards; report and flag publicly any policy that damages developmental outcomes.
- Ofsted becomes the field assessment arm of the governance body — not abolished but reframed from adversarial inspector to collaborative evidence-gatherer feeding a living learning system.
- The constitution is specifically designed to prevent the six failure modes identified in Section 5. Structure, not good intentions, is the guarantee.

Section 7 — The Testing Imperative

Nothing is deployed at scale without being tested properly first. This is not a preference or a guideline. It is the founding discipline of the governance body and the condition on which its credibility rests.

The history of education technology and curriculum reform is littered with interventions deployed at national scale before they were understood, on the basis of enthusiasm, commercial pressure, or political convenience rather than evidence. Interactive whiteboards. Learning management systems. Academisation. Free schools. Each arrived with confident claims. Each was deployed at vast expense before the evidence existed to support that confidence. Each produced results that were, at best, mixed and, at worst, damaging.

The governance body exists precisely to break this pattern. The testing imperative is its core operating discipline.

7.1 Pre-Registration

Before any pilot begins, the governance body publicly registers three things: what outcomes are expected, by when, and measured how. This registration is permanent and public. It cannot be amended after the pilot begins.

Pre-registration prevents the post-hoc reinterpretation of results that has allowed failed interventions to be presented as successes. If the expected outcome was a measurable improvement in belonging scores at six months and the pilot does not produce that, the pilot has not succeeded — regardless of what other interesting things it may have produced. The governance body reports this in full.

7.2 Independent Evaluation

Every pilot is evaluated by evaluators appointed by the governance body who are independent of the organisations that designed, funded, or delivered the intervention. Not the company that built the tool. Not the university that designed the curriculum. Not the government department that commissioned the work.

Independence is structural, not just declared. Evaluators are selected through a competitive process, their appointment is transparent, their findings are published in full regardless of what they show, and they have no contractual relationship with any organisation whose work they are evaluating.

7.3 What Is Actually Measured

The parallel metrics framework — detailed in Section 8 — is the measurement instrument. Not exam results. Not Ofsted grades. Not completion rates or attendance figures in isolation.

The five streams measure what Part A says actually matters: the physical foundation, psychological safety and belonging, engagement and purpose, capability development, and life trajectory. Every pilot is evaluated against the streams most relevant to what it claims to

do. An intervention claiming to improve engagement is evaluated against the engagement stream. An intervention claiming to improve physical wellbeing is evaluated against the physical foundation stream.

The connection to Part A is the point. The measurement framework exists to answer one question: is this intervention helping produce the 80-year human described in Section 2? If yes, scale it. If no, stop it and learn from why.

7.4 Acknowledged Failure

When a pilot fails to meet its pre-registered outcomes it is published, analysed in full, and the findings are used to inform the next iteration. It is not quietly discontinued. It is not rebranded. It is not attributed to poor implementation without evidence that implementation was the problem.

The governance body allocates an explicit failure budget — a defined proportion of pilot resources is expected to produce interventions that do not work. This is not waste. It is the cost of learning in a domain where the right answers are genuinely unknown. A governance body that never reports failure is not running genuine pilots. It is running demonstrations.

7.5 Scale-Up Criteria

An intervention moves from pilot to wider deployment only when it meets pre-registered outcome thresholds across multiple independent pilots in genuinely diverse school contexts — different regions, different demographics, different levels of resourcing, different existing performance levels.

A result from one outstanding school in one affluent area is not evidence of scalability. A consistent result across twelve schools in diverse contexts begins to be. The governance body defines the specific threshold for each category of intervention before piloting begins.

7.6 The Abandonment Protocol

Equally important to the scale-up criteria is the abandonment protocol. If an intervention fails to meet its pre-registered criteria after a defined period it is stopped and the resources redirected. No zombie programmes kept alive by institutional inertia, sunk cost thinking, or the embarrassment of admitting failure.

The abandonment protocol applies to the governance body's own initiatives as rigorously as to externally commissioned work. If a governance body programme is not working, the governance body stops it. The constitution makes this a requirement, not a choice.

KEY POINTS — Section 7

- Nothing is deployed at scale without being tested properly. This is the founding discipline of the governance body and the condition on which its credibility rests.
- Pre-registration is mandatory. Expected outcomes, timescales, and measures are publicly registered before any pilot begins and cannot be amended retrospectively.
- Evaluation is structurally independent. No organisation evaluates work it designed, funded, or delivered.
- Measurement is against the five parallel metric streams — what Part A says matters — not against exam results or inspection grades.
- Acknowledged failure is required. An explicit failure budget is allocated. Failed pilots are published in full and used for learning, not concealed or rebranded.

- Scale-up requires consistent results across diverse contexts. Abandonment is as rigorously governed as scale-up. No zombie programmes.

Proceeds to Section 8 — The Parallel Metrics Framework

Section 8 — The Parallel Metrics Framework

The measurement trap — identified in Section 5 as one of the six reasons previous reforms have failed — cannot be escaped by measuring more things. It can only be escaped by measuring the right things and refusing to let those measures become targets that schools game.

This section is clear about a difficult reality: we do not yet have the measurement infrastructure to do this properly. Some of what needs measuring — genuine curiosity, authentic belonging, real capability development — is genuinely hard to measure without creating the very gaming problem the framework is designed to avoid. A belonging score that becomes a target ceases to measure belonging. An engagement metric that schools report becomes an engagement performance. Goodhart's Law does not disappear because we understand it.

What this section therefore proposes is a developmental roadmap rather than a fixed framework. It describes where we can start with available tools today, where we are trying to get to over twenty-five years, and how the adaptive learning platform described in Section 11 provides the measurement infrastructure that makes reliable, ungameable, automated assessment of what actually matters eventually achievable. The framework is a direction. The governance body builds the evidence base as it goes.

The goal is not more measurement. It is better evidence — generated automatically, resistant to gaming, and connected to outcomes that matter across a human life.

8.1 What We Can Measure Now — Hard to Game, Already Available

The most defensible near-term measures are those generated administratively rather than reported institutionally. Schools cannot easily manipulate what they do not control the collection of. These measures are imperfect proxies but they are defensible ones:

- Attendance and persistent absence rates — already collected, nationally consistent, and a genuine signal of engagement and wellbeing. A school with rising persistent absence is a school where something is wrong. This does not tell you what, but it tells you where to look.
- Exclusion rates — permanent and fixed term. A lagging indicator of pastoral failure. Already collected. Resistant to gaming because exclusions are legally recorded.
- Free school meal uptake and eligibility data — a proxy for the physical foundation measures. Not perfect but available immediately.
- NEET destination data at 16, 18, and 24 — the ultimate lagging indicator, already partially collected, needing extension and standardisation. If the system is working, this number falls. If it does not fall, something fundamental is wrong regardless of what other measures show.
- Platform-generated diagnostic data — from year one of the adaptive learning platform deployment. Automatically collected, anonymised at individual level, aggregated nationally. Not self-reported. Not teacher-observed. Behavioural. A child

who engages voluntarily beyond the required task, whose error patterns flag a specific learning difference, whose engagement drops suddenly — these are signals the platform generates without asking anyone to report anything.

8.2 What We Are Building Toward — The Measurement Layer We Are Building Toward

The measures we actually want — genuine curiosity, authentic belonging, real capability development across the eight dimensions defined in Part A — require measurement infrastructure that does not yet exist at scale. The twenty-five year programme builds it.

The external mentor network described in Section 11 provides a second non-institutional data stream within three to five years of deployment. Mentors report to the governance body, not to the school. Individual data is confidential to the child and mentor. Class and school-level aggregates feed the governance body's evidence base. Because the mentor has no institutional interest in the school's performance measures, their observations are structurally more reliable than anything the school itself reports. This is not a perfect solution but it is a substantially better one than asking schools to report on their own developmental effectiveness.

The platform's longitudinal data becomes genuinely powerful after five to ten years of consistent deployment. A system that has tracked a child's learning patterns from age seven to sixteen has a picture of genuine capability development that no examination has ever produced. Not what the child performed under exam conditions on a specific day, but what the child actually understands, how their understanding developed, where it consistently faltered, and what interventions made the difference. This is the kind of evidence the governance body needs to make decisions. It takes time to build. It cannot be shortcut. But it can be started now.

8.3 The Long-Term Target — Life Outcomes as the Ultimate Measure

The only meaningful test of whether the education system is working is what happens to the people who passed through it — not at sixteen but at twenty-five, thirty-five, forty. Are they flourishing? Are they economically active? Are they in good health, physical and mental? Are they engaged with their communities? Are they continuing to learn? These are the outcomes defined in Part A as the design target. They are the only meaningful measure of whether the system is hitting it.

Building this longitudinal evidence base is a fifteen to twenty-five year task. Education records linked to employment, health, and wellbeing data already exist in fragmented form across government departments. The governance body commissions the cohort studies, coordinates the data linkage with appropriate privacy protections, and begins generating the evidence base that will, in time, tell us whether what schools do actually connects to what lives look like decades later. No education system has ever systematically built this evidence base. This one should. It is the only way to know, rather than assume, that the system is producing what it claims to produce.

8.4 The Goodhart's Law Defence

Any measure that becomes a target ceases to be a good measure. This is not a problem that can be designed away. It is a permanent dynamic that requires permanent governance attention.

The governance body's active monitoring for gaming is therefore a core function, not a supplementary one. It requires specific mechanisms, not just vigilance. First, triangulation across multiple data streams that cannot all be gamed simultaneously — a school that manages its belonging survey scores cannot simultaneously manage its platform behavioural data, its exclusion rates, and its NEET destination data. Second, regular independent audits comparing reported metrics against the platform's automatically generated behavioural evidence — discrepancies between what schools report and what the platform shows are the primary gaming signal. Third, mandatory rotation of external mentors across school clusters to prevent the relationship capture that would bias reporting over time. When a measure starts moving in ways that suggest institutional optimisation rather than genuine improvement — when belonging scores rise while exclusions also rise, when engagement metrics improve while NEET rates do not fall — the governance body investigates, replaces the compromised measure, and publishes its findings. The framework is a living instrument. The governance body's credibility depends on its willingness to acknowledge gaming and act on it, even when politically inconvenient.

This is the discipline that separates a genuine measurement programme from the performance theatre that current accountability systems produce. It requires an independent governance body with the authority and the will to make it happen. It is one of the strongest arguments for the governance model proposed in Section 6.

KEY POINTS — Section 8

- The current section is clear about a difficult reality: we do not yet have the measurement infrastructure to do this properly. The framework is a developmental roadmap, not a fixed specification.
- Near-term measures should be hard to game: administratively collected data (attendance, exclusions, NEET destinations) and platform-generated behavioural data that is automatically collected rather than institutionally reported.
- The external mentor network provides a structurally independent data stream because mentors report to the governance body, not to the school, with individual confidentiality protecting candour at the child level.
- The platform's longitudinal data becomes genuinely powerful after five to ten years — a picture of real capability development that no examination has ever produced.
- Life outcomes at 25, 35, 40 are the only meaningful ultimate test. The governance body commissions the cohort studies and data linkage that builds this evidence base over fifteen to twenty-five years.
- Goodhart's Law is a permanent dynamic requiring permanent governance attention. Active monitoring for gaming, willingness to replace measures when gamed, and public reporting of findings are core governance body functions, not optional extras.

Proceeds to Section 9 — What Can and Cannot Be Changed

Section 9 — What Can and Cannot Be Changed

The triage that every previous reform has avoided. Before the governance body can act it must understand the landscape it is acting within — what is genuinely changeable, on what timescale, at what cost, and what is fixed in the medium term regardless of political will or funding.

Three tracks run simultaneously from day one. They are not sequential. Track One begins immediately. Track Two begins in year one and runs for fifteen years. Track Three begins in year one as policy development and delivers structural change from year fifteen onward. All three are always active.

Track	Horizon	What it covers
One — Act Now	Years 1–3	No structural change required. School start times. Universal free meals. Mind and body in Initial Teacher Training. AI transparency requirements. Parallel metrics framework mandated. Quiet time daily. External Mentor Network piloted. Achievement awards redesigned.
Two — Build While Running	Years 3–15	Capability-based assessment piloted. AI longitudinal tracking in pilot schools. Teacher training redesigned. Ofsted framework progressively rewritten. Project-based learning at scale. Fast-track programme for 16-year-olds the system has failed.
Three — Long Game	Years 15–25	Qualifications framework reformed on basis of accumulated evidence. Academy structure rationalised. Independent sector brought within transparency framework. Persistent Personal Mentor deployed as public infrastructure. Modular lifelong learning architecture.

This document does not ask teachers to become developmental psychologists, AI specialists, pastoral counsellors, and curriculum designers simultaneously. It asks them to do what the best teachers have always done — build genuine relationships with children, respond to what they find, and teach with the kind of human presence and professional judgement that no platform can replicate. The adaptive learning platform takes the marking, the tracking, and the diagnostic reporting. The external mentor network takes the pastoral identification and support burden. The AI assessment tools take the administrative reporting. What remains for the teacher is the irreducibly human work — and more time to do it. The teacher training redesign is about ensuring teachers are equipped for that work, not about adding to an already unsustainable burden.

9.1 What Cannot Be Changed Quickly

Naming what cannot be changed quickly is as important as identifying what can. Previous reforms have repeatedly promised structural transformation on timescales that were never achievable, then failed to deliver, then been abandoned. This document does not make that mistake.

- The school estate — approximately 24,000 buildings, most more than thirty years old. Rebuilding or substantially reconfiguring the physical estate takes decades and requires capital investment at a scale no government has yet committed to. Any reform that requires a fundamentally different physical environment at scale is a long-term aspiration, not an immediate intervention.

- The teacher workforce — 500,000 teachers cannot be retrained simultaneously. The pipeline takes years to change. Any reform dependent on a fundamentally different kind of teacher is a ten to fifteen year project at minimum. Track Two begins the redesign. Track Three delivers it.
- The examinations framework — deeply embedded in university admissions, employer expectations, and international recognition. Cannot be dismantled quickly without cascading damage. This is the hardest single change in the entire system and requires fifteen years of accumulated evidence from the parallel metrics framework before the political conditions for reform exist.
- The independent sector — cannot be directly regulated without primary legislation and significant political capital. The transparency framework is the minimum viable near-term intervention. Full inclusion is a Track Three ambition.
- The academy structure — 10,000 academies with individual funding agreements cannot be homogenised quickly. Reform must work with this fragmentation, not against it, in the short and medium term.

9.2 The Agile Principle

The governance body operates on agile principles — not the waterfall model of traditional public sector project management that specifies everything in year one, commits the budget, lets the contract, and delivers the original specification five years later regardless of what has been learned in the interim.

Agile means: multiple small pilots running in parallel, each with clear success criteria, each evaluated independently, each either scaled or abandoned on a defined timescale. Sprint cycles every six months — the governance body reviews all active pilots and makes explicit decisions: scale, modify, or stop. No single large programmes. No bets on one approach. No zombie projects.

This is not a preference for a particular management methodology. It is the only approach that is epistemically defensible in a domain where the right answers are genuinely unknown and where the cost of getting it wrong at scale is borne by children.

9.3 Commissioning and Contracts

The governance body does not build things itself. It commissions them — from businesses, research institutions, and educators working in partnership. The contracting model is as important as the governance model. The current public sector approach to education technology procurement is a reliable producer of expensive mediocrity.

The governance body's contracting principles are simple and non-negotiable:

- Small contracts, multiple competitors. No single provider gets the whole thing. Three to five organisations build competing versions of any significant intervention. They are tested against each other. The best version wins the right to further development. The others are paid fairly for their work.
- Performance-based contracts. Payment is linked to measured outcomes against the parallel metrics framework — not to delivery of a specification but to evidence that the intervention actually moves the developmental indicators it claims to affect.
- Open specifications. The governance body publishes what it wants built. Any qualifying organisation can bid. Specifications are written by people who understand what is needed — not by civil servants who have never built anything.

- Public ownership. Anything built under governance body contract is publicly owned. Code, content, design. Companies are paid for their work. They do not own the result and cannot withdraw access or charge for continued use.
- Small company access. The contracting model is specifically designed to make small, innovative organisations viable bidders. This is where the most original ideas come from and they are currently systematically excluded from public sector procurement.

KEY POINTS — Section 9

- Three tracks run simultaneously from day one: act now with no structural change required; build while running over five to fifteen years; structural change over fifteen to twenty-five years.
- What cannot be changed quickly is named directly: the school estate, the teacher workforce, the examinations framework, the independent sector, the academy structure. Previous reforms have failed by promising structural transformation on impossible timescales.
- The governance body operates on agile principles — multiple small parallel pilots, six-month sprint reviews, explicit scale-or-stop decisions. No single large programmes. No waterfall commitments.
- Contracting is by competing organisations against open specifications, with performance-based payment tied to developmental outcomes, and public ownership of everything built. The current procurement model is not used.

Proceeds to Section 10 — The 25-Year Pathway

Section 10 — The 25-Year Pathway

This is not a plan. It is a framework for generating the plan through a governed process. The distinction matters. A plan specifies what will happen. A framework specifies who will decide, how they will decide, what evidence they will use, and what the destination is. The governance body generates the plan. This document describes the framework within which it operates.

Two measures are held constant across all 25 years as the north star. Every decision the governance body makes is tested against them. They are simple, meaningful, and resistant to gaming:

- What proportion of young people leaving the system at 18 demonstrably possess the eight capabilities defined in Part A?
- What proportion are NEET — not in education, employment, or training — at 24?

Currently close to one million young people aged 16 to 24 are NEET in the UK — 12.8% of that age group, rising, and placing the UK in the bottom six of all OECD countries. The economic cost alone is estimated at £125 billion per year. These are the baseline numbers. The governance body's mandate is to move them — measurably, consistently, and permanently.

10.1 Years One to Five — Establish, Audit, Begin

The governance body is constituted and the selection body appointed in year one. The ground truth audit is commissioned — the definitive baseline of the current system described in Part C. The parallel metrics framework is mandated across all state-funded schools. Track One interventions begin: school start times, free meals extension, AI transparency requirements, mind and body in Initial Teacher Training.

The External Mentor Network is piloted in a geographically and demographically diverse set of 200 schools. The first AI tools are evaluated against the governance body's standards framework. The first six-month sprint review takes place. Some pilots will already be showing promise. Some will already be failing. Both are reported openly and publicly.

By year five the governance body has: a functioning metrics framework producing baseline data across all schools; a portfolio of active pilots at various stages of development; the first evidence of which Track One interventions are moving the north star measures; and a clear picture of what the most important Track Two priorities are.

10.2 Years Five to Fifteen — Evaluate, Scale, Reform

The evidence from Track One pilots is accumulated and published. What works is scaled. What fails is stopped, documented, and learned from. The metrics framework is continuously refined as correlations with life outcomes begin to emerge.

Teacher training redesign is implemented. The new Initial Teacher Training specifications — incorporating neuroscience, body-mind connection, trauma-informed practice, and developmental psychology — produce the first cohort of newly qualified teachers trained in what Part A says matters. This cohort enters the workforce from approximately year four onward and grows as a proportion of the workforce through this period.

Assessment reform is piloted. Continuous capability profiles are tested in volunteer schools alongside terminal examinations. The evidence from these pilots begins to build the case — or not — for the Track Three qualifications reform.

By year fifteen the governance body has: fifteen years of parallel metrics data showing what school-level conditions predict genuine developmental outcomes; evidence from assessment pilots about whether capability profiles work; a teacher workforce that is progressively more developmentally informed; and the political evidence base for Track Three structural change.

10.3 Years Fifteen to Twenty-Five — Structural Change

The qualifications framework is reformed on the basis of fifteen years of accumulated evidence. This is the hardest single change in the system and it requires political will sustained across multiple parliaments. The governance body's role is to make that will possible by making the evidence impossible to ignore.

The Persistent Personal Mentor — described in Part C — is deployed as public infrastructure, informed by everything the human mentor network has learned over fifteen years. The modular lifelong learning architecture begins to take shape. Education that genuinely does not end at 18 becomes a practical reality rather than an aspiration.

By year twenty-five the system is not transformed. Transformation is too strong a word for what is achievable in any timeframe without sustained political will across multiple governments. What is achievable is a system that is measurably better at producing the 80-year human than the one it replaced — with the evidence to prove it, the governance infrastructure to sustain it, and the renewal mechanisms to keep improving it.

There is a longer argument still. Education reform is not only about the children currently in schools. It is a generational investment whose returns compound. The children educated well today become the parents, employers, policymakers, and citizens of tomorrow. Parents who were themselves given the developmental foundations described in Part A — who can think clearly, know themselves, manage their own wellbeing, and build genuine relationships — create better environments for their own children's development. Employers who understand human capability rather than credential compliance build organisations that develop rather than diminish the people in them. Policymakers who can reason carefully and tolerate genuine uncertainty make better decisions about complex systems. The cycle, if the foundations are right, moves upward — across health, mental health, economic participation, social cohesion, and the quality of democratic life. It will never be perfect. No generation inherits a clean slate. But the direction matters enormously, and the current direction is not upward. This is what a 25-year commitment to getting the foundations right is actually for.

> ***The goal is not a perfect system. The goal is a learning system — one that knows what it is trying to produce, measures whether it is producing it, and changes when the evidence says it should.***

KEY POINTS — Section 10

- This is a framework for generating the plan, not the plan itself. The governance body generates the plan through a governed process of evidence, testing, and rigorous evaluation.
- Two north star measures are held constant: proportion of 18-year-olds demonstrably possessing the eight capabilities, and NEET rate at 24. Everything is tested against these.

- Years one to five: establish the governance body, baseline the system, begin Track One interventions, pilot the External Mentor Network, commission the ground truth audit.
- Years five to fifteen: scale what works, stop what fails, reform teacher training, begin assessment pilots, build the evidence base for structural change.
- Years fifteen to twenty-five: qualifications reform, Persistent Personal Mentor as public infrastructure, modular lifelong learning. Structural change built on fifteen years of evidence.
- The goal is not a perfect system. It is a learning system that knows what it is trying to produce, measures whether it is producing it, and changes when the evidence says it should.

One significant gap in this framework document is the absence of an overall funding model. The adaptive learning platform is costed. The governance body's setup and running costs are comparable to the Committee on Climate Change. But teacher salaries, school buildings, mentor networks, early years provision, and the full cost of the system this document describes are not modelled here. This is a deliberate choice — a framework document that specifies funding in advance would be doing precisely what Section 5 identifies as a failure mode: committing to a plan before the evidence exists to justify it. The governance body's ground truth audit, commissioned in year one, is the appropriate mechanism for establishing what the reformed system actually costs and where that funding comes from. The investment case is clear. The specific figures require the evidence base to be built first.

Proceeds to Part C — The Inquiry Agenda

The Case for Human Education

Why the right to a fully human life demands a wholly different education system

PART C — THE INQUIRY AGENDA
The Twenty-Five Year Journey

Section 11 — The Immediate Opportunity

Part C opens not with a distant aspiration but with the most immediately actionable proposal in this entire document. It requires no structural change to the school system. No new legislation. No curriculum reform. No change to how teachers teach. It can be designed, built, piloted, and deployed nationally within thirty months.

It is also, by a considerable margin, the single intervention most likely to improve learning outcomes for the most children, in the shortest time, at the lowest cost per pupil. The technology exists. The evidence base exists. The public ownership model exists. The curriculum content foundation exists. The only thing that has not previously existed is someone specifying it clearly enough that it can actually be commissioned.

This section does that.

> ***The governance body's first and most urgent commission: build the world's best publicly owned adaptive learning platform for every state school pupil in England. Not a supplement to education. A transformation of it. Achievable within thirty months.***

11.1 What Already Exists — and What Does Not

Oak National Academy is a publicly funded, arm's-length body of the Department for Education. Since 2020 it has built something genuinely valuable: a professionally designed, nationally sequenced, open-licensed curriculum content library covering every national curriculum subject from Key Stage 1 to Key Stage 4. That intellectual work is done. The subject matter, the learning objectives, the pedagogical sequencing — designed by curriculum experts, aligned to the national curriculum, freely available under open government licence. This is a significant public asset.

Oak also provides proof of concept that the UK education system will adopt online curriculum tools. It is used in some form by the majority of schools. Teachers report meaningful workload reduction. That is important political and cultural evidence.

Oak has built the curriculum. The governance body builds the platform. These are different things requiring entirely different capabilities. Conflating them would be a strategic mistake.

The new platform is built from scratch, on modern architecture, by teams selected for capability. Over time — as the new platform demonstrates its superiority in every measurable dimension — Oak's natural role becomes what it does best: maintaining and updating the national curriculum content library that informs the new system's design. That is

a genuinely valuable function. It is not a technology function. The new platform will, in time, make Oak's current online roll entirely redundant. That is not a problem. It is the point.

11.2 What the World Has Already Built

While UK state education has been offering children a video and a worksheet, the commercial sector has been building something categorically different. The evidence is substantial and consistent.

DreamBox Learning uses adaptive algorithms to personalise mathematics lessons in real time for primary and secondary students. Students using DreamBox for 60 minutes per week improved their mathematics scores by nearly 60% beyond expected growth — meaning a student expected to gain 10 points on a standardised assessment gained approximately 16, a modest absolute gain but consistent across diverse school contexts. Century Tech combines AI with neuroscience and learning science to deliver tailored content across core subjects, providing real-time diagnostic data to teachers and enabling targeted intervention. ALEKS uses a sophisticated knowledge-mapping system to identify precisely what a student knows and what they are ready to learn next — students using it showed 15% higher completion rates in mathematics courses. Arizona State University deployed adaptive learning technology in its Introduction to Biology course and saw a 24% increase in pass rates and a 20% decrease in dropout rates.

These are not experimental results from a single study. They are consistent findings across diverse contexts, age groups, and subjects. Adaptive learning platforms demonstrably outperform static content delivery. The evidence is not contested.

The adaptive learning market was estimated at £3 billion globally in 2024 and is projected to reach £25 billion by 2034. The world is building this at enormous scale. UK state education is offering a worksheet.

The question is not whether this technology works. It does. The question is whether the UK builds a publicly owned version that every child can access free, or whether it waits until the commercial version achieves lock-in and the disadvantage gap becomes permanent.

11.3 The Concept — What We Are Actually Proposing

This is not a MOOC in the traditional sense. It is not a digitised textbook. It is not Oak with better graphics. It is a fundamentally different conception of what learning through a multimedia platform can be — one that takes seriously what we know about human attention, motivation, memory, and the neuroscience of learning established in Part A.

The core insight is simple. Gaming companies routinely produce immersive, narrative-driven, visually extraordinary experiences that hold the attention of children for hours. The same techniques applied to the curriculum would transform engagement entirely. Not as a gimmick. As a delivery mechanism that works with human neurology rather than against it.

A blood cell travelling through the human body, encountering each organ system, understanding through experience what each one does. Archimedes in his bath, working through the problem of the crown — the child discovering the principle of displacement alongside him, understanding why the mathematics works because they watched the problem demand it. The Fibonacci sequence revealed through the petals of a flower and the spiral of a shell, making visible why certain numbers appear everywhere in nature and others never do. The strange necessity of $\sqrt{-1}$ — arriving not as an abstract rule but as the answer to a question that has no other solution, the moment a child understands that mathematics

sometimes has to invent new kinds of number to stay true to itself. A young person standing in ancient Rome, talking to a senator, understanding why the Republic fell because they were there when it happened. A sci-fi flythrough of the solar system that pauses at each planet and asks what happens to gravitational force at this distance — and the child answers to see what the spacecraft does next.

This is not entertainment with a thin educational veneer. This is the oldest and most effective learning mechanism humans have — story, consequence, curiosity, revelation — delivered with the production quality that children now expect from the screens they use every day. The child answers the question to see what happens next. That is not a trivial motivation. It is the mechanism that keeps them engaged long enough to learn.

Alongside the cinematic layer sits the adaptive engine. The platform knows, from every previous interaction, what this child understands and what they do not. It does not deliver the same content to every child. It delivers what this child needs, now, at this moment, in the sequence that will build understanding most effectively from where they actually are — not where the curriculum assumes they should be.

> ***This is what education could have been for thirty years if the available technology had been positioned as an educational tool. The governance body's job is to build it now, own it publicly, and give it to every child in the UK as a free and remarkable learning resource.***

11.4 Built for the Teacher, Accessed by the Student

Every previous attempt to introduce technology into schools has failed partly because teachers experienced it as something done to them. Another system to learn. Another administrative burden. Another platform that replaced their professional judgement with someone else's algorithm.

This system is built around the opposite principle. It serves the teacher directly, reduces their workload immediately and measurably, and gives them information they could never previously have had. The teacher is not a passive user of the system. The teacher is the professional whose judgement the system is designed to extend.

Homework integration

The platform follows the teacher's course. Not a predetermined national sequence that the teacher must conform to — the teacher selects their units, their starting point, their pace. The platform aligns to that selection automatically. When the teacher sets homework, they choose from the platform's aligned module for that lesson, customising the questions if they wish, selecting the focus, adjusting the difficulty for their class. The platform delivers that homework to every child. It marks it. It tracks it. It records completion and performance.

The teacher receives, automatically, a dashboard showing every child's completion, every child's score, every child's specific error patterns. Not a summary. Every child. Every question. Which ones understood it first time. Which ones needed three attempts. Which ones show a pattern suggesting a specific misconception. Which ones did not attempt it at all and why — the platform records whether the child opened the homework and whether technical access was a factor.

The teacher walks into the next lesson knowing exactly where every child is. Not approximately. Exactly. They did not spend Sunday evening marking. They did not spend Monday morning trying to remember who handed what in. The system did it. They have the information. They teach.

The professional autonomy guarantee

The teacher's professional judgement is not replaced — it is extended. The teacher decides what is taught and when. The teacher decides which homework questions their class receives. The teacher can add their own content, modify the narrative context for their specific class, override the adaptive recommendations if their knowledge of the child tells them something the algorithm does not know. The system serves the teacher's judgement. It never substitutes for it.

Default settings are sensible enough that a teacher can use the system with zero configuration. The professional who wants deeper customisation can have it. Neither is required. The system works at whatever level of engagement the teacher chooses.

The political opportunity

Teachers have always deserved better tools than the profession has been given. A system that marks the homework, tracks the understanding, and delivers the diagnostic picture — while leaving every pedagogical decision in the teacher's hands — is not a threat to professional autonomy. It is the practical expression of it. A teacher who has used this for one term will not give it up. That is how cultural change happens in a system of 500,000 teachers — not through mandate but through a product so obviously useful that the people who try it become advocates. The teacher unions' consistent objection to educational technology is that it threatens professional autonomy. This system does the opposite — it is the answer to that objection, built into the architecture from day one. The strategic opportunity is significant: a teaching profession that spends less time on administration and more time on what only humans can do will be a more effective, more satisfied, and more sustainable profession. That is an educational argument, not a technology pitch.

11.5 The Diagnostic Revolution

The most transformative capability of the system is not the cinematic content. It is what happens in the background while the child is learning.

Every interaction with the platform generates data. Every question answered, every misconception revealed, every concept that landed and every one that did not, every pattern of errors that a human teacher could never track across thirty children simultaneously. Aggregated, anonymised, and fed to the governance body's diagnostic framework — automatically, continuously, without any teacher filling in a form.

Within two years of national rollout, the governance body will have, for the first time in UK education history, the following data:

- Which mathematical concepts are consistently misunderstood at which ages nationally — enabling curriculum adjustment in real time rather than waiting for exam results years later.
- Which children show consistent error patterns indicating likely dyscalculia — flagged at age six or seven rather than discovered at eleven after years of unnecessary struggle and accumulated damage to self-confidence.

- Which children show reading comprehension patterns consistent with dyslexia — identified through the pattern of responses to narrative content, not through a teacher noticing something is wrong.
- Which children from non-English-speaking homes are falling behind because of language rather than conceptual difficulty — enabling targeted language support rather than incorrect assumptions about general ability.
- Which concepts show sudden national dips in performance — potentially indicating a curriculum sequencing problem, a widespread misconception introduced by common teaching approaches, or an external factor affecting a cohort.
- Which schools show patterns suggesting a curriculum delivery problem rather than a pupil capability problem — the difference between a school that needs curriculum support and one that needs pastoral support.

None of this requires a specialist assessment. None of it requires a referral. None of it requires a parent to advocate loudly enough to be heard. None of it requires a teacher to notice something is wrong before it has already been wrong for years. The system sees the pattern. It flags it. The teacher receives the information. The governance body has the national picture.

The platform's diagnostic capability represents the document's primary contribution to the challenge of special educational needs and disabilities. Early identification — at seven rather than thirteen, through behavioural patterns rather than parental advocacy — is the single most important improvement the current SEND system requires. However, the platform identifies. It does not treat. Complex SEND needs — autism spectrum conditions, significant physical disabilities, profound learning difficulties, mental health crises — require specialist human support that no platform can replace. The governance body must commission a dedicated SEND workstream from year one, ensuring that the platform's diagnostic capability connects to a support system capable of acting on what it finds. Early identification without adequate specialist provision is worse than no identification — it creates expectation without delivery.

The child who would previously have struggled silently until secondary school, accumulated years of confusion and self-doubt, and eventually been assessed as having a specific learning difficulty at thirteen — that child is identified at seven. The intervention happens early. The damage does not compound.

This is not a marginal improvement to the current system. It is the elimination of one of the most persistent and damaging failures in UK education — the systematic late identification of children who needed help that was always available but never delivered in time.

11.6 The Platform Architecture Principles

The platform is built from scratch on modern architecture. It does not inherit any constraints from Oak's existing technical infrastructure. The following principles are non-negotiable design requirements, not aspirations.

- Open source codebase — publicly owned, publicly accessible, not dependent on any single vendor. The governance body owns the intellectual property. No company can withdraw access, increase charges, or impose conditions.
- Device agnostic — fully functional on any device including low-specification tablets and smartphones. Degraded but usable on slow connections. This is the technical

expression of the equity principle. A platform that requires a high-specification device is not a public platform.

- Offline capability — core learning modules downloadable for use without internet connection. This addresses the connectivity gap that COVID exposed and that remains real in rural areas and disadvantaged households.
- Data sovereignty — all pupil data stored on UK-sovereign infrastructure. No data leaves UK jurisdiction. No commercial entity has access to pupil-level data. The governance body's data ownership principle, built into the architecture from day one.
- Teacher-first dashboard design — the teacher interface is as carefully designed as the student experience. Simple, fast, actionable. Five minutes to understand everything the teacher needs for tomorrow's lesson.
- AI transparency — the adaptive engine's recommendations are explainable. The teacher can see why the system suggested what it suggested. The algorithm serves the professional. It does not replace their judgement with an unexplained recommendation.
- Accessibility by default — full compliance with accessibility standards, with specific design attention to dyslexia, dyscalculia, and EAL from the ground up, not retrofitted.

11.7 The Project Plan

The following is a realistic outline programme. Timings assume the governance body is constituted and has its initial budget from year one. The contracting model is as specified in Part B — competing teams, open specifications, performance-based payment, public ownership of everything built.

Phase	Timing	Budget estimate	What happens
1 — Foundation	Months 1–6	£5–8m	Commission platform architecture — open source, new build. Commission content production framework — cinematic standards, narrative design, game mechanics. Three to five competing development teams.
2 — Pilot content	Months 4–12	£15–25m	Build initial subjects in full across all key stages — subjects selected by the governance body based on pilot evidence and strategic priority. Three competing teams per subject. Best version selected per unit. Teacher dashboard and homework integration built alongside.
3 — Large pilot	Months 10–18	£8–12m	Deploy across 500 schools — diverse geography, demographics, deprivation levels, Ofsted history. Pre-registered outcomes. Independent evaluation throughout. Device programme for disadvantaged pupils included.
4 — Iterate	Months 16–22	£5–8m	Publish pilot results in full. Iterate content and platform based on evidence. Second pilot with revised version in different schools. All remaining subjects begin production.
5 — National rollout	Months 20–30	£20–30m	Every state school in the UK. Every subject. Every key stage. Free to every pupil. Automatic diagnostic reporting to teachers and governance body. Ongoing maintenance and content development: £10–15m per year.

Total programme cost to national rollout: approximately £53 to £83 million on current estimates. These figures are deliberately conservative and acknowledge significant uncertainty — particularly in Phase Two content production, where cinematic-quality output across all subjects and key stages is the most variable element. A more cautious range of £150 to £250 million would not be unreasonable, and the governance body should plan for this upper range. Even at £250 million, the programme costs less than two months of the annual NEET economic impact estimated by the Milburn review.

For context: the UK government spent £11.3 billion on a failed NHS IT programme. £1.7 billion on a single failed Universal Credit digital system. The transformation of UK state education through the most powerful learning platform ever built in this country costs less than one medium-sized government IT failure. Built with competing teams, open source architecture, performance-based payment tied to measured outcomes, and full public ownership — the programme design specifically addresses every reason previous government technology programmes have failed.

Divided by 9 million state school pupils, the total programme cost is less than £10 per child. The annual maintenance cost is approximately £1.50 per child per year. These are not large numbers. They are the cost of not having done this thirty years ago, finally being paid.

11.8 The Learning Environments — What Becomes Possible

The governance body determines which subjects are built first and in what sequence. That decision should be informed by the pilot evidence, not predetermined here. What can be said with confidence is that the range of subjects where immersive adaptive learning produces qualitatively better outcomes than any classroom can deliver is very wide.

In mathematics and the sciences, the platform can show what no textbook or whiteboard can — the invisible made visible, the abstract made concrete. A child can watch a mathematical principle emerge from a real problem that demanded it, understand why the number exists before being told what it is called, see the pattern that connects algebra to the spiral of a shell to the architecture of a virus. In the sciences a student can journey inside the systems they are studying — through the human body, across geological time, into the quantum scale — experiencing the principles they are learning rather than reading descriptions of them.

In history and the humanities, the platform can place a student inside the world they are studying. Not as a passive observer but as a participant making decisions, experiencing consequences, understanding causation because they were present when it operated. A student who has navigated the collapse of the Roman Republic understands political fragility in a way that no essay question has ever produced.

In languages, the platform can create genuine immersive environments that accelerate acquisition in ways classroom instruction cannot replicate. In computing, game mechanics and real programming challenges align so naturally that the distinction between learning and doing dissolves entirely.

The point is not which subjects are built first. The point is that this platform can deliver, in every subject, something the classroom alone has never been able to provide: a learning environment that adapts to the individual, responds to their understanding in real time, makes the subject genuinely compelling, and generates the diagnostic data that tells the teacher and the governance body what is actually happening. That is what the governance

body is commissioning. The curriculum teams determine the content. The production teams determine the form. The evidence determines what scales.

11.9 The Equity Imperative

The digital divide is real and must be addressed in the rollout plan explicitly. During the COVID school closures, access to online learning platforms for economically disadvantaged students fell to just over 40% — barely half the rate of their more advantaged peers. A national platform that 60% of the most disadvantaged children cannot access is worse than nothing. It would widen the gap it was supposed to close.

The rollout plan therefore includes a device programme as a non-optional component, not a supplementary aspiration. Every pupil in the lowest two deprivation quintiles who does not have an appropriate device at home receives one, funded from the programme budget. Every school in areas of poor connectivity receives a supplementary connectivity solution as part of the rollout. The offline capability built into the platform architecture is the technical backstop for situations where even this is insufficient.

The equity programme is estimated at £15 to £25 million for initial device provision, reducing significantly thereafter as devices are maintained and replaced on a rolling programme. This cost is included in the total programme budget. It is not optional and it is not negotiable. A public platform that is not genuinely accessible to every pupil is not a public platform.

11.10 What This Changes

When this platform is at national scale, the following things are permanently different.

Every child in the UK, regardless of school quality, teacher experience, geographic location, or family background, has access to the same quality of curriculum delivery. The child in a struggling school in a deprived area has access to exactly what the child in the best-resourced school in the country has access to. Not approximately. Exactly.

Children who are ill, excluded for the day, missing school due to family circumstances, or in hospital receive exactly the same learning continuity as children in school. The platform does not know or care that the child is not in the building. It continues the learning journey from where the child left off.

Children with unidentified learning differences are identified earlier than the current system has ever managed. Not at thirteen after years of struggle. At seven, eight, nine — when intervention makes the largest difference and the damage has not yet compounded.

Teachers spend less time on marking, record-keeping, and assessment administration than at any point in the history of state education. They spend more time on what only humans can do — the relationship, the encouragement, the explanation that goes beyond the algorithm, the professional judgement that knows this child in a way no system can.

The governance body, for the first time, has a real-time national picture of what children understand and what they do not — by subject, by concept, by age, by region, by demographic group. Not a picture constructed from exam results two years after the fact. A picture updated continuously, automatically, as children learn.

This is not a marginal improvement. It is the most consequential single intervention available to UK state education today — earlier diagnosis, better learning outcomes, reduced teacher burden, and genuine equity of access. It is achievable within thirty months. The governance body's first commission should be to build it.

KEY POINTS — Section 11

- The governance body's first and most urgent commission is to build a publicly owned, open-source, adaptive learning platform for every state school pupil in the UK. Not a MOOC. Not a supplement. A transformation of curriculum delivery, achievable within thirty months.
- Oak National Academy has built the curriculum content library — publicly owned, open-licensed, nationally sequenced. This is genuinely valuable. Oak's technical platform is not a foundation for what is proposed and should not constrain it. The new platform is built from scratch. In time it will make Oak's student-facing role redundant.
- The commercial sector has already proved this works. DreamBox users improved maths scores by 60% beyond expected growth. Adaptive learning platforms consistently outperform static delivery across diverse subjects and age groups. The evidence is not contested.
- The teacher proposition is the strategic key. The platform follows the teacher's course, delivers their homework, marks it automatically, and returns a dashboard showing every child's performance by concept before the next lesson. It is built for the teacher. It extends professional autonomy rather than threatening it.
- The diagnostic capability is transformative. Dyscalculia, dyslexia, language barriers, specific misconceptions — all identified automatically through the pattern of a child's responses. At age seven, not thirteen. Without a specialist referral. Without a parent having to advocate loudly enough to be heard.
- The programme costs approximately £53 to £83 million to national rollout — less than £10 per pupil. Less than one medium-sized failed government IT project. Built with competing teams, open source architecture, performance-based payment tied to measured outcomes, and full public ownership — the programme design specifically addresses every reason previous government technology programmes have failed.
- Equity is non-negotiable. A device programme for disadvantaged pupils is included in the budget. Offline capability is built into the architecture. A platform that 60% of the most disadvantaged children cannot access widens the gap it was supposed to close.
- When this is at national scale: every child has equal access to the same quality of delivery. Absent children continue learning. Learning differences are identified early. Teachers spend more time teaching. The governance body has a real-time national picture of what children actually understand.

Proceeds to Section 12 — What We Already Know Works

The Case for Human Education

Why the right to a fully human life demands a wholly different education system

PART C — THE INQUIRY AGENDA

Where we are, where we might go, and who controls the journey

Part C is the expedition map. It begins with what we can already see clearly — what the evidence from working systems tells us, what others around the world are doing, what we can build on with confidence from day one. It moves through what we can reasonably infer but have not yet proven. It ends at the frontier, where the known runs out and the governance body established in Part B takes over as the expedition leader.

One claim runs through every section of Part C, and it is the most important claim in the entire document. Technology must work for us. We do not work for it. The education system chooses its destination, chooses its tools, evaluates what it finds, and changes course when the evidence demands it. Nobody makes those choices for us — not commercial platforms, not foreign governments, not the velocity of technological change. The governance body established in Part B exists to ensure this. Part C is the agenda it must pursue.

Section 12 — What We Already Know Works

Part A made the case from developmental science and moral principle. Part C begins by showing that case in action — not in theory but in schools that exist, children who have passed through them, and evidence that has been independently gathered. The governance body does not begin its work on a blank canvas. It begins with a body of proof of concept that, while incomplete and not yet at national scale, is sufficient to establish that the direction of travel is right.

The consistent finding across every working example is the same. When you build an educational environment around genuine engagement, real-world application, human connection, and the development of capabilities rather than the reproduction of content, children flourish. Not some children. All children — including those the conventional system consistently fails.

11.1 XP School, Doncaster — Proof in a State School

XP School in Doncaster was founded in 2014 on the explicit premise that the existing model was wrong. It delivers its entire curriculum through cross-subject learning expeditions — term-long projects built around real-world questions. Students maintain portfolios, present their own learning publicly, and engage with local businesses and community organisations as genuine participants, not passive observers.

It is a state school in one of England's most deprived areas. It is rated Outstanding by Ofsted. Its Progress 8 score — the measure of how much students improve relative to their starting point — places it in the top quarter of schools nationally. Staff turnover is among the lowest in the country. The school has a waiting list.

What XP demonstrates is not that a different curriculum is possible in principle. It demonstrates that it works in practice, in a state school, with disadvantaged children, in a

deprived post-industrial town, against the same measures the conventional system uses to judge itself. The governance body should regard XP not as an interesting exception but as a validated model to be understood, dissected, and tested for scalability.

The key question XP does not yet answer is whether it scales. XP works partly because of its size, its culture, and the specific people who built it. Whether its core principles — real-world learning, cross-subject integration, portfolio assessment, genuine community connection — can be replicated across diverse schools without the founding team and the founding culture is precisely the question the governance body should be commissioning rigorous pilots to answer.

11.2 High Tech High, San Diego — Proof at Greater Scale

High Tech High in San Diego has been running project-based, interdisciplinary learning since 2000. It now operates across seventeen schools in the San Diego area. Its students report 82% college acceptance rates. Four-year college enrolment among graduates increased by 24 percentage points over its first decade. The approach works across demographic groups including students from low-income families who are the majority of its intake.

High Tech High also operates a graduate school of education that has trained hundreds of teachers in project-based learning methodology. This is the missing piece in the UK context — not just the model but the teacher training infrastructure that can spread it. The governance body should establish a formal relationship with High Tech High and commission a UK adaptation study as one of its first acts.

11.3 The Scandinavian Early Years Model — Proof in the Foundation Period

Finland does not begin formal academic instruction until age seven. The years before that are spent in play-based, largely outdoor, physically active, socially rich environments designed around developmental readiness rather than curriculum content. Finland has consistently outperformed the UK in international assessments from age ten onwards.

The careful use of the Finnish evidence — as discussed in Part A — requires acknowledging that Finland's recent decline in PISA scores is real and significant, and that multiple factors are involved including rising socioeconomic inequality and curriculum changes introduced in 2016. What the Finnish evidence does establish clearly is that delaying formal instruction does not damage academic outcomes and that the variation between schools in Finland remains only 9%, compared to an OECD average of 30%. Where the developmental foundations are properly laid, the outcomes are more equitable.

Estonia, Norway, Denmark, and to a significant degree Canada and New Zealand show comparable patterns. The evidence that the 0–7 period should be developmental rather than formally academic is consistent across multiple high-performing systems. This is not a contested finding. It is a settled one that the UK education system has chosen to ignore.

11.4 Shanghai's Deep Integration Teaching

One Shanghai school has been running a model called Deep Integration Teaching since 2013 — deliberately dismantling subject boundaries for interdisciplinary inquiry. It reports outstanding performance in creative thinking assessments. It also, strikingly, reports the lowest myopia rates in its district, nearly twenty percentage points below average — a direct

consequence of prioritising outdoor activity and physical movement as part of the learning day.

The myopia finding is more than a footnote. It is the body-mind connection evidenced in a real school context. A school that takes the physical foundation seriously enough to measure it finds that the developmental outcomes follow. The governance body should commission a formal study of the Deep Integration Teaching model and its applicability in UK contexts.

11.5 What These Examples Tell the Governance Body

Four things are established by the existing evidence base that the governance body can act on with confidence:

- Project-based, real-world, cross-disciplinary learning works — for disadvantaged children, in state schools, producing measurable outcomes by conventional metrics as well as developmental ones.
- Delayed formal instruction in the early years does not damage academic outcomes and produces greater equity. The developmental approach to 0–7 education is not a soft option. It is better.

The 0–5 period deserves particular attention. The evidence on early brain development is unambiguous: the foundations of cognitive function, emotional regulation, language acquisition, and social capability are substantially laid before a child enters formal schooling. The appropriate environment for this period is not a classroom. It is a well-resourced play environment — structured around physical activity, social interaction, creative exploration, and adult-supported discovery. No formal instruction. No desks. No assessments. The body-mind thread that runs through this entire document is most critical here: children of this age learn through movement, play, and genuine human relationship, not through sitting still and listening. The Scandinavian model demonstrates this works. The obstacle in the UK is not philosophical. It is structural. Genuine developmental provision for 0–5 requires adult-to-child ratios that the current system does not fund. You cannot deliver the relational, responsive, physically active environment that early development requires with one adult supervising twenty toddlers. This is a funding question, not a design question. The governance body must commission the early years review as one of its first acts and be prepared to make the funding argument honestly — because the evidence for return on investment in early years provision is among the strongest in the entire educational research literature.

- The body-mind connection is not theoretical. Schools that take physical movement, outdoor time, and physical health seriously produce measurably better cognitive and developmental outcomes.
- Teacher training is the critical variable. The governance body cannot scale better approaches without training teachers to deliver them. This is the infrastructure investment that makes everything else possible.

None of these findings is new. All of them have been available to UK education policy for decades. The governance body's first job is to act on what is already known — not to wait for more evidence on questions that are already settled.

KEY POINTS — Section 12

- XP School Doncaster proves the model works in a state school in a deprived area. Outstanding Ofsted rating. Top quarter Progress 8 nationally. The governance body should treat it as a validated model to be understood and tested for scalability.
- High Tech High San Diego proves it works at greater scale across diverse demographics. 82% college acceptance. The graduate school of education is the teacher training model the governance body needs.
- The Scandinavian evidence establishes that developmental 0–7 provision produces better equity of outcomes. This is a settled finding the UK has chosen to ignore.
- Shanghai's Deep Integration Teaching demonstrates the body-mind connection in practice — including measurably lower myopia rates from prioritising physical activity.
- Four things can be acted on immediately: project-based learning, developmental early years, body-mind provision, and teacher training redesign. The evidence is sufficient. Waiting for more is a choice not to act.

Proceeds to Section 13 — The Global Landscape

Section 13 — The Global Landscape

The UK is not alone in recognising that the existing model is inadequate. Across the world, educators, policymakers, researchers, and governments are grappling with the same questions this document raises. None of them has assembled a complete answer. Most are working on fragments. But the fragments, taken together, confirm that the direction this document proposes is not idiosyncratic. It is where serious thinking about education is converging.

The governance body is not starting from scratch. It is entering a global conversation and a global evidence base. Its job is to learn from what others are discovering, contribute what the UK's own pilots produce, and build the international partnerships that will make the evidence base richer faster than any single country could achieve alone.

12.1 The OECD Learning Compass 2030

The OECD has been running its Education 2030 project since 2016, with 35 countries engaged. Its central framework — the Learning Compass — asks what students will need to thrive in 2030 and beyond, and proposes student agency, wellbeing, and genuine competencies as the answer. The metaphor is deliberate and important: a compass rather than a map. Students learning to navigate by themselves through unfamiliar contexts rather than following prescribed instructions.

This framing is philosophically aligned with Part A of this document and with the 80-year human as the design target. The OECD's framework gives the governance body institutional credibility and a network of 35 countries with whom to share evidence and develop common approaches.

Ontario, Canada, offers a directly relevant precedent. The Education Quality and Accountability Office, established in 1996, operates as an independent statutory body in a multi-party democratic system with regular changes of government — closer to the UK context than either Singapore or Finland. It has maintained consistent assessment standards across multiple provincial governments with entirely different education philosophies. It demonstrates that independent educational governance can survive political turbulence in a Westminster-style system. The significant limitation of the OECD framework is also worth stating. Thirty-five countries are engaged in conversation. Very few are substantially changing their systems. The gap between the framework and the practice is wide. The governance body's contribution to this global conversation is not just intellectual alignment — it is the commitment to actually build and test what the framework describes.

12.2 South Korea — The Warning

South Korea is the clearest example in the world of what happens when you optimise an education system entirely for academic performance. PISA scores among the highest globally. Rates of youth depression and suicide among the highest in the OECD. A generation of young people who have been drilled to academic excellence and broken in the process.

South Korea's own analysts describe three possible futures for their education system to 2045. In the first, students continue in a cycle of extra classes, exams, rankings, burnout,

and depression with mental health deteriorating further. In the second, far-reaching reform places individual capacities and personal development at the centre. In the third, a hybrid attempts to balance both.

The South Korean case is not a cautionary tale about a different culture. It is the logical endpoint of the same incentive structure that governs UK education — taken further and faster. The UK's own rising rates of school anxiety, persistent absence, and youth mental health problems are early signals of the same dynamic. South Korea shows where that dynamic leads if left unaddressed.

12.3 China — Ambition and Caution

China's education reform programme is the most ambitious in the world by scale. Its 2024–2035 master plan explicitly targets the development of creative, innovative individuals capable of driving national technological leadership. In 2024 alone China added 1,673 new undergraduate programmes in strategic fields while removing 1,670 that did not align with economic development priorities. That is structural surgery at a pace no democratic system can match.

The Shanghai Deep Integration Teaching model described in Section 11 is one product of this reform culture. There are others. China is investing heavily in AI-supported personalised learning, immersive educational environments, and the integration of technology into teaching at a scale the UK cannot currently match.

The caution the governance body must hold alongside this is equally important. China's reform is driven by national competitiveness and state formation — producing talent that serves the party's vision of national development. That is not self-determination. It is sophisticated state formation wearing the language of personalised learning. The governance body can learn from China's methods. It must never adopt China's purposes.

12.4 The Consistent Global Pattern

Across every country where serious thinking about educational transformation is happening, the same pattern emerges. The fragments of the right approach exist everywhere. The complete system exists nowhere. XP School works. High Tech High works. The Finnish early years model works. The OECD has the right framework. South Korea has the clearest warning. China has the greatest ambition. Estonia has the best equity outcomes in Europe.

Nobody has assembled these fragments into a coherent national system. The governance body's opportunity — and its mandate — is to do what no country has yet done: take the existing evidence seriously, test it rigorously in the UK context, build the connections between the fragments, and scale what works.

> ***The UK does not need to invent a new education system from scratch. It needs the governance architecture to assemble what the world already knows into something that works at national scale. That architecture is Part B. The assembly instructions are Part C.***

KEY POINTS — Section 13

- The OECD Learning Compass 2030 framework is philosophically aligned with this document and provides international credibility and a 35-country evidence network. The limitation is that countries are talking but not yet building.
- South Korea is the warning: academic optimisation at scale produces extraordinary exam results and extraordinary rates of youth depression and suicide. The UK's own rising absence and mental health figures are early signals of the same dynamic.
- China's ambition is real and its methods are worth learning from. Its purposes — state formation and national competitiveness rather than self-determination — are not.
- The consistent global pattern: the fragments of the right approach exist everywhere. The complete system exists nowhere. The governance body's job is to assemble what the world already knows. {and rigorously test it at scale]

Proceeds to Section 14 — The Technology Landscape

Section 14 — The Technology Landscape

Technology is not the answer to the questions this document raises. The answer is human development. Technology is a set of tools that may, if chosen carefully and governed properly, help deliver that development more effectively than was previously possible. Understanding those tools fully — their current capability, their trajectory, their limitations, and their dangers — is essential preparation for the governance body's work.

This section does not predict the future of technology. It describes what exists now, what the trajectory suggests is coming, and where the genuinely open questions are. The governance body will need to update this picture continuously. The tools available in year fifteen of its mandate will be substantially different from those available in year one. The principles for evaluating them will not change

13.1 AI in Education — The Current State

Artificial intelligence is already deployed in UK schools at significant scale — through marking and feedback tools, adaptive learning platforms, administrative systems, and increasingly through general-purpose AI assistants that students use independently of any school programme. Most of this deployment has happened without systematic evaluation, without governance oversight, and without the transparency requirements that the governance body will introduce.

What AI can currently do well in educational contexts: generate feedback on written work at speed and scale; identify patterns in student performance data that individual teachers cannot track; adapt the difficulty and pace of practice tasks to individual performance; reduce administrative burden on teachers substantially; and provide accessible explanations of concepts across a wide range of subjects.

What AI currently cannot do reliably: assess open-ended creative work, complex ethical reasoning, or context-dependent judgement without significant human oversight; build the kind of genuine relational trust that psychological safety requires; identify the difference between a child who is struggling and a child who is deliberately underperforming; or replace the irreducibly human qualities of the best teachers — passion, moral example, the sense that this person believes in you specifically.

Khanmigo — Khan Academy's AI tutor — is the most developed AI educational tool currently deployed at scale. It grew from 68,000 pilot users in 2023–24 to over 700,000 in 2024–25. Students who used Khan Academy for an average of 30 minutes of additional mathematics practice per week saw greater-than-expected gains on standardised assessments. The limitation is that the efficacy studies are still ongoing — the evidence is promising but not yet definitive. The governance body should monitor this evidence closely and commission comparable UK-based studies.

13.2 MOOCs — What Fifteen Years Has Taught Us

Massive Open Online Courses launched around 2012 with extraordinary ambition — world-class education, free, globally accessible, democratising learning at a scale never previously possible. The reality was more complicated and more instructive.

Completion rates of 3–10% became the headline statistic used to dismiss the MOOC experiment. This interpretation fundamentally misunderstands what was actually happening. MOOCs operate in the reality of participants' everyday lives — no financial incentives, no protected time, competing priorities. The fact that significant proportions of enrollees completed courses under these conditions represents genuine commitment. More importantly, the completion rate misses the population that enrolled specifically to access particular content, got what they needed, and left. These learners — self-directed, purposeful, using the resource on their own terms — were not failures. They were exactly the kind of lifelong self-directed learners the governance body is trying to produce.

The experience of this document's principal author is illustrative. EDX and FutureLearn were used extensively and productively without formal enrolment, often in disagreement with the methodology but always extracting genuine value. Eric Lander's introduction to biology course on EDX was, by that account, the best piece of education received in a long life. That outcome does not appear in any completion statistic.

What the MOOC journey has established beyond serious dispute: short, structured, purposeful learning units work better than long abstract courses; employer-linked credentials drive completion more effectively than intrinsic interest alone; AI personalisation significantly improves both engagement and outcomes; and the commercial capture of the MOOC infrastructure has progressively narrowed its original democratising mission. Coursera and Udemy merged in a £2.5 billion deal in late 2025, creating a platform with 270 million registered learners and 19,000 enterprise customers. The democratising promise is being quietly buried under subscription models and employer-facing products.

13.3 The MOOC-Plus-AI Convergence

The combination of MOOC-quality content with AI personalisation and persistent memory is the most significant development in the technology landscape for education. It is already beginning to happen — Coursera's AI coach, Khanmigo on Khan Academy's content library — but the full convergence is not yet realised.

What the full convergence could deliver: every mainstream subject available as a high-quality, self-paced, adaptive learning experience, accessible to every child in the country regardless of school quality or geographic location. A child in a poorly resourced school in a deprived area with access to the same quality of content as a child at a leading independent school. The equaliser that the original MOOC promise described, now technically achievable.

The threat vector is equally significant. Commercial lock-in at planetary scale with AI personalisation creating dependency is the worst-case version of this convergence. A small number of platform companies owning the learning infrastructure of an entire generation — the data, the relationships, the dependency — is the commercial education nightmare that Section 15 addresses directly.

The governance body's response is the public alternative. A UK-owned MOOC infrastructure, built to open specifications, publicly owned, designed for genuine developmental outcomes rather than engagement metrics, available free to every child and every adult learner in the country. This is one of the most important long-term commissions the governance body should pursue.

13.4 The Trajectory — What Is Coming Regardless

Three developments are coming regardless of what the governance body does. The question is whether they arrive on terms the UK chooses or on terms chosen by commercial and foreign state actors.

First: AI systems that can maintain persistent memory of individual learners across years of interaction, building an increasingly accurate model of how that person thinks, what they know, where their gaps are, and how they learn best. This is technically achievable now in limited forms and will be widely deployed within five years.

Second: multimodal AI that can generate explanations, diagrams, simulations, and interactive environments on demand, calibrated to the individual learner's level and learning style. Not retrieved from a library of pre-made content but generated specifically for this person, at this moment, in response to this question.

Third: the progressive integration of AI into every aspect of cognitive work — writing, analysis, research, coding, design — which will continue to change what skills are valuable, what work looks like, and therefore what education needs to produce. This trajectory does not have a predictable endpoint. It requires an education system capable of producing people who can navigate whatever it produces — which is precisely the 80-year human argument of Part A.

KEY POINTS — Section 14

- AI is already in UK schools at significant scale, largely without governance, evaluation, or transparency. The governance body's transparency requirements are urgently needed.
- AI can currently do specific things well in education: feedback, pattern recognition, adaptive practice, administrative reduction. It cannot replace relational trust, genuine passion, or the irreducibly human qualities of the best teachers.
- MOOCs taught fifteen years of lessons. Completion rates were the wrong measure. The real lessons: short purposeful units work; AI personalisation helps; commercial capture narrows the mission. The Coursera-Udemy merger is the commercial capture completing itself.
- The MOOC-plus-AI convergence is the most significant development in the technology landscape. The governance body should commission a publicly owned UK alternative before commercial lock-in is complete.
- Three developments are coming regardless: persistent AI memory of individual learners; on-demand multimodal content generation; progressive AI integration into all cognitive work. The question is whether they arrive on UK terms or commercial terms.

Proceeds to Section 15 — The Frontier Ideas

The ideas in this section are not near-term proposals. They are the territory the governance body should be exploring and testing over a 15–25 year horizon — possibilities grounded in current technological trajectories but not yet sufficiently developed or evidenced to be commissioned at scale. They are presented to illustrate the kind of thinking the governance body should be pursuing, not to prescribe what it will find.

15.1 The Immersive Learning Environment

The civilisation model — deep immersive study of a complete human period or system, with the student as participant rather than observer — becomes considerably more powerful in an immersive digital environment. Technology already exists that can place a student inside

ancient Rome, the Industrial Revolution, or the human immune system, making decisions and experiencing consequences. The governance body should commission competing designs for specific curriculum areas, test them against the capability metrics, and scale what demonstrates genuine developmental benefit. The contracting model — small competing commissions, public ownership of what is built — is designed for exactly this kind of development work.

15.2 The Public MOOC Infrastructure

Every mainstream subject in the national curriculum should, over time, have a parallel MOOC — publicly owned, designed to open specifications, available free to every child and adult learner in the UK. Not replacing classroom teaching but sitting alongside it, extending it, equalising access to quality that currently depends on school quality and postcode. Over time, the MOOC infrastructure becomes the foundation for the immersive environments described above — content proven and refined in the MOOC, then developed into the immersive version as the technology matures and evidence accumulates. Section 11 describes the adaptive platform that is achievable now. The full MOOC infrastructure is the longer-term extension of that work.

15.3 The External Mentor Network

Every school cluster has an allocated external mentor, operating under the governance body's framework, working directly with children — potentially using Ofsted's field infrastructure as the delivery mechanism. The mentor's dual role: gathering developmental data through direct relationship, and providing the consistent caring adult connection that the developmental science of Part A identifies as a precondition for genuine learning. Individual data stays confidential to the child and mentor except where safeguarding requires otherwise. Class and school-level data feeds the parallel metrics framework. The mentor reports to the governance body, not to the school — the structural guarantee of independent data.

15.4 Longer-Term Possibilities

The following represent possibilities that depend on technological development, accumulated evidence, and governance conditions that do not yet exist. They are not proposals. They are the horizon the governance body should be watching.

- A Persistent Personal Mentor — an AI system that knows an individual learner over time, maps their genuine understanding, identifies what they do not know they do not know, and works from that map to build real capability. Fundamentally different from current AI tools which respond to what you ask rather than diagnosing what you need. If and when this becomes technically reliable, it must arrive as public infrastructure, designed constitutionally to build autonomy rather than dependency. The governance body's role is to ensure the public version exists before the commercial version achieves lock-in.
- Modular lifelong learning architecture — education that genuinely does not end at 18, with a portable credential system that recognises learning throughout life and connects it to genuine capability rather than course completion.
- Qualification reform — the replacement of terminal examination-based qualifications with continuous capability profiles, built on 15+ years of evidence from the parallel metrics framework and the adaptive platform. The political conditions for this do not currently exist. The governance body's task is to build the evidence base that makes those conditions possible.

KEY POINTS — Section 15

- Section 15 describes possibilities for a 15–25 year horizon, not near-term proposals. The governance body's role is to explore and test this territory, not to implement it before the evidence exists.
- The immersive learning environment and the public MOOC infrastructure are the nearest-term of these ideas. Both have current technical foundations. Both require the contracting model: competing commissions, public ownership.
- The External Mentor Network is achievable within the first five years and serves both a data-gathering and a developmental relationship function.
- The Persistent Personal Mentor is a 15–25 year possibility, not a near-term proposal. The governance body's task is to ensure that when it arrives, it arrives as public infrastructure designed to build autonomy — not as a commercial product designed to create dependency.

Proceeds to Section 16 — The Commercial Threat

Section 16 — The Commercial Threat

Every tool humanity has developed to extend human capability has eventually been monetised. The printing press. The telephone. The internet. In each case the tool arrived with democratising promise and was captured by economic logic that converted the promise into a product. This pattern is not accidental. It is structural. And it has never mattered as much as it matters now.

Because the tools being monetised this time are not external to the human being. They are not printing presses or telephones. They are systems designed to operate inside the cognitive and developmental life of every person from childhood. Systems that will know more about the inner life of a human being than any institution in history. Systems whose value to their owners increases precisely as their hold on the user deepens.

To monetise this is not a business model. It is the conversion of human development itself into a revenue stream. It is the logical endpoint of surveillance capitalism — not the harvesting of behavioural data as a byproduct of using a service, but the direct commodification of the process by which a human being becomes who they are.

15.1 What Is Already in UK Schools

The governance body does not inherit a blank canvas. It inherits a landscape in which commercial providers are already deeply embedded in UK schools, largely without transparency, largely without evaluation, and with data practices that most schools have neither the expertise nor the resources to scrutinise.

Google's infrastructure — Classroom, Docs, Chromebooks — is at near-total penetration in many school districts. A generation of children whose entire educational digital life exists within Google's systems, generating data for a company whose primary business is targeted advertising. Microsoft's Copilot is embedded across its entire product suite including the educational tools used in schools. Meta has explicitly invested in education technology. The combined reach of these three companies across UK schools is not yet fully documented — one of the governance body's first acts must be to document it.

Shoshana Zuboff, whose framework of surveillance capitalism is the most rigorous analysis of this dynamic, identified the education technology sector as one of the primary vectors for the extension of behavioural data extraction into childhood. That analysis was made before the current generation of AI tools existed. Its force has only increased.

15.2 The Incentive Architecture Problem

The people building commercial AI for education are not, in the main, malevolent. Many are genuinely motivated by the belief that they are making learning better. The problem is not the people. It is the incentive architecture within which they operate.

A company whose revenue depends on continued subscription and engagement cannot be constitutionally aligned with building learner autonomy. Autonomy reduces engagement. Reduced engagement reduces revenue. The company that is honest about this will always lose market share to the company that is not. The market cannot solve this problem. It produces it.

Meta. Google. Microsoft. And Musk — whose xAI and Grok are integrated into X, the platform he has deliberately shaped into an instrument of particular political and cultural influence, whose Neuralink is the logical endpoint of the Persistent Personal Mentor concept except implanted, proprietary, and owned by one individual with explicit civilisational ambitions. Each of these actors is building toward the same territory from different directions. None of them is building toward the destination this document describes. Steinbeck's machine that must be fed is already in motion. Nobody controls it. It simply demands.

15.3 Why the Governance Body Must Hold the Line

The governance body's commercial oversight powers — described in Part B Section 6 — are not a nice-to-have. They are the condition on which the entire enterprise rests. Without them, every pilot the governance body runs will be surrounded by commercial systems competing for children's attention and data on terms the governance body cannot control. Without them, the public MOOC infrastructure will be undercut by commercial alternatives that are better resourced, more engaging, and designed for dependency. Without them, the Persistent Personal Mentor will arrive as a commercial product before the public version is ready.

The transparency requirement is the minimum. Every commercial provider operating in UK schools must disclose, fully and publicly: what data is collected, how it is used, what the commercial model is, and what evidence exists for the developmental claims made. This is not onerous. It is the standard that any honest actor should be able to meet. Providers that cannot meet it should not be in schools.

The data ownership principle is the foundation. A child's developmental data belongs to that child, irrevocably and permanently. Not to the school. Not to the platform. Not to the state. The governance body's legislative framework must enshrine this beyond the reach of any future commercial pressure or political convenience.

> ***We use the tool. The tool does not use us. The governance body exists to ensure this principle holds in practice and not merely in aspiration. That is perhaps its most important single function.***

KEY POINTS — Section 16

- Every tool developed to extend human capability has been captured by commercial logic. AI deployed in human development is different in kind — not a byproduct of a service but the direct monetisation of how a child becomes who they are.
- Commercial providers are already deeply embedded in UK schools without adequate transparency or evaluation. Documenting this landscape is one of the governance body's first acts.
- The problem is not malevolent actors. It is an incentive architecture that makes building learner autonomy structurally incompatible with building revenue. The market cannot solve this. It produces it.
- The governance body's commercial oversight powers are not supplementary. Without them every other initiative is undermined by commercial systems operating on terms the governance body cannot control.
- Two non-negotiables: full transparency from every commercial provider in schools; developmental data belonging to the child irrevocably. These must be in primary legislation.

Proceeds to Section 17 — The Known Unknowns

Section 17 — The Known Unknowns

This document ends where it said it would begin — with honesty about what we do not know. Section 2 established the voyage of discovery as the governing metaphor and the intellectual discipline. Part C ends by mapping the frontier clearly: what remains unknown, what the governance body must find out, and what we may never know regardless of how good the inquiry becomes.

This is not a counsel of despair. It is the condition of all serious work in complex systems. The governance body that pretends to know more than it does will make the same mistakes as every previous reform. The governance body that is honest about the limits of its knowledge, and builds the mechanisms to extend those limits systematically over 25 years, has a genuine chance of producing something that lasts.

16.1 What We Could Find Out

These are answerable questions. The governance body's pilot programme and longitudinal research agenda should prioritise them because the answers will materially change what the system does.

- Does project-based, cross-disciplinary learning at scale produce better developmental outcomes than the conventional model? XP School and High Tech High suggest yes. The UK evidence base at national scale does not yet exist.
- How long does it take for improvements in 0–7 developmental provision to show measurable outcomes in later learning? The Finnish evidence suggests the benefits compound over years. The precise timescale and the mechanisms involved are not well understood.
- What does the capability profile of an 18-year-old who has experienced a genuinely developmental education look like, and how does it predict adult life outcomes at 30, 40, 50? No country has yet built the longitudinal evidence base to answer this definitively.
- Can AI assessment tools reach the reliability needed to supplement and eventually reduce the primacy of terminal examinations? The trajectory suggests yes within ten to fifteen years. The specific conditions under which AI assessment is reliable, and the domains where it cannot be trusted, are still being established.
- What does genuine persistent memory in an AI educational system do to learning outcomes and to learner autonomy over time? Nobody knows because no one has yet built and studied it properly.
- Does the External Mentor Network produce measurably better developmental data than teacher-reported measures, and does the mentoring relationship itself improve outcomes? Answerable by the pilot programme within three to five years.

16.2 What We Cannot Know Until We Build and Test

These are not answerable by research alone. They require building something, deploying it, watching what happens, and being honest about what we find.

- Whether the whole system coherently delivers what it promises. Individual components can be tested in isolation. The interaction effects — what happens when project-based learning, capability assessment, the mentor network, and AI-supported personalisation all operate together in the same school — can only be understood by building the whole and watching.

- What unintended consequences emerge. There will be some. There always are. The commitment is not to prevent unintended consequences — that is impossible — but to identify them quickly, report them honestly, and respond proportionately.
- Whether the Persistent Personal Mentor, however well designed, can replicate the thing that makes the best human teaching transformative — the sense that this person believes in you specifically, that their passion for the subject is infectious, that your encounter with them changed how you see the world. This may be permanently beyond what AI can do. Or it may turn out that AI can do something different but equivalently powerful. We do not know. We will not know until we build and test it seriously.
- Whether the governance body itself works as designed. The constitution is designed carefully. The incentives are designed to produce honest inquiry. But institutions are unpredictable. The governance body may calcify despite the renewal mechanisms. It may be captured despite the independence protections. It may find that the north star measures are the wrong measures. The commitment is to notice when this is happening and to change course — including, if necessary, changing the governance body itself.

16.3 What We May Never Know

These are permanently open questions. They are worth holding clearly rather than pretending they have been resolved.

- Whether any alternative system would have produced better outcomes than the one it replaced. The counterfactual cannot be run. A generation educated differently cannot be compared to the generation that would have been educated the same way. The best we can do is track the measures that matter and make the case from evidence — knowing that the case will always be probabilistic rather than definitive.
- How to separate the effect of the education system from the effect of everything else that shapes a child. Family, community, economics, culture, chance — all of these shape developmental outcomes as powerfully as any school. The governance body can measure what it can measure and be honest about what it cannot attribute.
- Whether political will can be sustained across 25 years and multiple governments. The constitution is designed to make abandonment difficult. It cannot make it impossible. The ultimate protection against the reform horizon problem is not legislative. It is political culture — a sufficient proportion of the population understanding why this matters and demanding its continuation. That culture is built by the evidence the governance body produces and by the lives it demonstrably improves.
- Whether we will back-track on some of what we propose here. We will. Some of the ideas in Part C will not work. Some will work in ways we did not predict. Some will produce results that contradict the assumptions of Part A and require us to revise those assumptions. The commitment is not to have been right in advance. It is to follow the evidence, report honestly, and change course when the evidence demands it.

16.4 The Closing Statement

This document began with a moral claim. Every human being has the right to be the author of their own life. That right requires capacity — the capacity to think, to know oneself, to relate to others, to adapt, to choose deliberately rather than accidentally. Education is the process by which that capacity is built. The current system does not reliably build it. This document proposes the governance architecture and the inquiry agenda through which a different system can be developed, tested, and sustained.

It does not claim to have all the answers. It claims to be asking the right questions at the right moment, with a process serious enough to find the answers that can be found and rigorous enough to live with the ones that cannot.

The technology is arriving whether we are ready or not. The commercial forces are building whether we govern them or not. The children are in schools whether the system serves them or not. The question is whether democratic societies — this one specifically — choose to shape what is coming on terms that serve human development, or allow it to be shaped on terms that serve commercial returns or simply fails yet again.

That choice is available now. It will not remain available indefinitely. This document is an argument that it should be made deliberately, urgently, and with the full weight of what is at stake for every child who passes through the system we build.

KEY POINTS — Section 17

- The governance body that is honest about the limits of its knowledge and builds mechanisms to extend those limits systematically over 25 years has a genuine chance of producing something that lasts.
- What we could find out: scalability of project-based learning, timescales for early years investment to show outcomes, AI assessment reliability, the effect of persistent AI memory on learner autonomy. These are the priority research commissions.
- What we cannot know until we build: whether the whole system coherently delivers, what unintended consequences emerge, whether the Persistent Personal Mentor can replicate what the best human teaching does, whether the governance body itself works as designed.
- What we may never know: the counterfactual, the attribution problem, whether political will will hold. The commitment is to follow the evidence and change course when it demands it — including back-tracking on what does not work.
- The choice between technology that serves human development and technology that feeds on it is available now. This document is an argument that it should be made deliberately, urgently, and with full understanding of what is at stake.

END OF MAIN DOCUMENT

Appendices follow — Evidence base, International case studies, UK ground truth data, Commercial AI landscape in UK schools, The Persistent Personal Mentor — extended concept paper, The Body-Mind Evidence Base, Governance models, The SFT concept document, The NEET data analysis.

The Case for Human Education

Addendum 1.0 — The Moral Foundation: The Discussion

This addendum is referenced from Section 1 of the main document. The main text asserts that every human being has a moral right to be the author of their own life. This addendum examines that claim — what it means, what challenges it, what the philosophical tradition says, and why it has become more urgent than any previous generation could have anticipated.

1.0 The Moral Foundation — The Discussion

1. The Premise and Its Challenge

The main document opens with a moral claim: every human being has a right to be the author of their own life. This is treated as foundational — the premise from which everything else follows. It is not naive about what it is doing. A premise is a starting point, not a conclusion, and starting points can be contested.

The most serious challenges come from two directions.

The communitarian challenge: individuals do not exist prior to their communities. We are constituted by our languages, cultures, relationships, and inherited social roles. The idea of a self-determining individual who authors their own life from some prior position of freedom is a liberal fiction. What matters is not self-authorship but belonging, contribution, and the health of the social fabric. On this view, education should form citizens capable of sustaining communities, not individuals pursuing private self-realisation.

The state formation challenge: the primary purpose of education is national. Educated citizens are productive workers, reliable taxpayers, effective soldiers, and participants in democratic governance. The state has a legitimate interest in shaping what education produces, and that interest may not align with individual self-authorship. Singapore, China, and historically Prussia all operated explicitly on this basis and produced remarkable results by their own measures.

These are serious positions with serious intellectual traditions behind them. The question is whether they are better than the premise they challenge.

2. What the Alternatives Actually Produce

State formation education produces citizens who serve the state's purposes efficiently — until the state's purposes diverge from human flourishing. At that point, citizens formed to comply have no internal resources with which to resist, evaluate, or redirect. They are defenceless precisely because their capacity for independent judgement was never developed. The history of the twentieth century is substantially a history of what happens when populations have been educated to defer.

Commercial capture is state formation's market equivalent. No deliberate ideology is required. The incentive structure of platforms optimised for engagement shapes human cognitive and emotional life as a byproduct of profit maximisation. Citizens formed to

consume, scroll, and respond to algorithmic stimulation are not free individuals. They are the product of a different formation process — one whose authors are shareholders, not philosophers or statesmen, and whose purpose has nothing to do with human flourishing.

The communitarian challenge is subtler and more serious, but it does not actually argue against developing human capacities. It argues about the relationship between individual development and community. A community of underdeveloped, dependent, uncritical individuals is not a healthy community. The communitarian vision of a flourishing community depends on flourishing individuals who are capable of genuine participation, genuine relationships, and genuine contribution. The communitarian and the liberal agree more than they disagree — the disagreement is about emphasis and sequence, not about whether human capacities should be developed.

None of the alternatives is a philosophy of human flourishing. They are departures from it. You cannot construct a serious philosophical argument that a human being flourishes more fully by having their capacity for self-directed reason suppressed.

3. What the Philosophical Tradition Says

The premise of this document is not a recent liberal invention. It is the sustained conclusion of the most serious inquiry into human flourishing across more than two thousand years of philosophy.

Aristotle — Eudaimonia and the Development of Capacities

Aristotle's concept of eudaimonia — usually translated as flourishing or happiness but more precisely 'living and faring well' — is built on a specific claim about human nature. Human beings have characteristic capacities — reason, language, practical wisdom, friendship, political participation — and they flourish when those capacities are developed and exercised. A human life in which those capacities are suppressed or never developed is, for Aristotle, a diminished human life regardless of its material comfort.

Aristotle was not a liberal individualist. He thought humans were political animals who could only flourish in community, and he had no objection to the state educating citizens for civic purposes. But his account of eudaimonia is incompatible with an education that suppresses the development of reason and practical wisdom. The state's obligation, for Aristotle, is to create conditions in which citizens can flourish — which means developing their capacities, not limiting them.

Kant — Persons as Ends

Kant's categorical imperative states that persons must be treated as ends in themselves, never merely as means. An education system that produces examination-optimised workers — whose purpose is economic productivity rather than their own flourishing — is treating persons as means to economic ends. This is not a contingent political preference that communitarians can reasonably reject. It is, for Kant, a categorical moral violation: the use of rational beings as instruments rather than as beings whose rational nature deserves respect and development.

The Kantian argument does not require liberal individualism. It requires only that we take seriously the moral status of the persons being educated. An education that treats children as inputs to an economic process is morally impermissible regardless of how productive that process is.

Mill — The Social Value of Developed Individuality

John Stuart Mill in On Liberty argues that the development of individuality — the cultivation of each person's unique capacities and the freedom to exercise them — is not only a personal good but a social one. Societies of developed, thinking, self-directed individuals are more capable, more resilient, more creative, and more just than societies of compliant, underdeveloped ones. The diversity of genuinely developed human beings is itself a social resource — the source of new ideas, new solutions, new challenges to established error.

Mill's argument answers the communitarian challenge directly: the community is not served by producing underdeveloped citizens. It is served by producing developed ones who can genuinely participate in it.

Sen and Nussbaum — The Capabilities Approach

Amartya Sen's capabilities approach, developed in the late twentieth century and extended by Martha Nussbaum, argues that human development should be measured not by income, GDP, or examination results but by what people are actually able to do and be — their real freedoms. An education that produces examination results but not genuine capability is a failure by this measure regardless of its economic outputs.

Nussbaum's list of central human capabilities — which she argues any just society must enable — includes practical reason (being able to form a conception of the good and engage in critical reflection about the planning of one's own life), senses, imagination and thought, emotions, affiliation, and play. These are not liberal preferences. They are universal human requirements derived from the study of what human beings need in order to live a distinctively human life. An education system that fails to develop them is not politically neutral. It is unjust.

4. The New Urgency — Non-Human Intellect

The philosophical arguments above were made in a world where the only competition for human cognition was other humans and human institutions. Aristotle, Kant, Mill, Sen — none of them were writing in a context where a non-human intellect of potentially vast and rapidly increasing capability was being integrated into the cognitive infrastructure of daily life.

That changes the stakes entirely. The traditional arguments for developing human self-direction were always strong. They are now existential.

A human being who has been educated to comply, reproduce, and defer — who has never developed genuine critical capacity, genuine self-knowledge, genuine ability to evaluate evidence and reach their own conclusions — is not merely unfulfilled in the Aristotelian sense. They are defenceless. They have no internal resources with which to evaluate what an AI system tells them, to resist manipulation at scale, to distinguish between what they actually think and what they have been shaped to think.

The Chinese state model is one version of this danger — a state that uses increasingly powerful AI systems to shape the cognitive and emotional lives of its citizens in directions the state determines. But the commercial version may be the more immediate threat in democratic societies. It is less visible, more seductive, and operates through apparent choice. At least in a state formation model there is a visible author of the shaping. In commercial capture, the shaping happens through a thousand small frictionless nudges, each one optimised to feel like a free decision. Steinbeck's machine that must be fed — no malice, no conspiracy, simply the structural logic of a system optimising for engagement and

revenue — is already building the infrastructure of cognitive dependency in children who have not been given the resources to recognise what is happening.

> ***The development of non-human intellect does not make the philosophical arguments for human self-determination obsolete. It makes them urgent in a way no previous generation needed to understand.***

Aristotle argued that the suppression of human rational capacity is a diminishment of human life. He was right, but in his world the diminishment was bounded — by the limits of human institutions, human lifespans, human power. In the world this document addresses, the cognitive infrastructure is being built at a scale and speed that no previous institution could approach. The window in which human beings can choose to develop the cognitive and moral resources needed to navigate this world is not indefinitely open.

This is why the main document says: the case for human self-determination is no longer only philosophical. It is existential. Not because previous generations were wrong to value it. But because the cost of failing to develop it has never been higher, and the systems that will operate on undeveloped human minds have never been more capable.

5. The Challenge This Document Issues

This document operates from the premise that human self-determination is the foundational purpose of education. It does not claim this premise is beyond challenge. It claims it is the best available premise — the one most consistent with the philosophical tradition, most consistent with what we know about human flourishing, and most adequate to the moment we are in.

If there is a better premise — one that does not require the suppression of human rational capacity, does not treat persons as means to external ends, does not leave human beings cognitively defenceless in the face of non-human intellect, and produces more flourishing rather than less — this document has not encountered it. The argument remains open. The governance body's public participation mechanism is precisely the space in which different conceptions of flourishing can be heard, debated, and incorporated into the system's evolving understanding of what it is trying to produce.

The question is not whether to have a premise. Every education system has one — explicit or hidden, chosen or inherited. The question is which premise produces the most flourishing and the least harm. This document's answer is human self-determination. The alternatives, examined honestly, are worse.

The Case for Human Education

Reference Appendix — Evidence Base and Source Citations

This appendix provides the evidential basis for significant factual claims made in the main document, organised by section. Where a claim is contested, qualified, or based on estimates rather than precise data, this is noted. References are numbered sequentially within each section.

Abbreviations: OECD — Organisation for Economic Co-operation and Development. PISA — Programme for International Student Assessment. DfE — Department for Education. ONS — Office for National Statistics. CCC — Committee on Climate Change.

Part A — The Foundation

Sections 1–4 — The Moral Foundation, Voyage of Discovery, Future, Human Development

NEET Statistics — UK

> *Nearly one million young people aged 16–24 are not in education, employment or training. The NEET rate is 12.8%, rising, placing the UK in the bottom six of all OECD countries. Annual cost estimated at £125 billion.*

[1] Powell, A. (2026). NEET: Young people not in education, employment or training. House of Commons Library Research Briefing No. 6705, 27 February 2026. — *Primary source for all UK NEET statistics cited in the document. Q4 2025 figure: 957,000 aged 16–24 NEET (12.8%).*

[2] OECD (2024). Youth not in employment, education or training (NEET). OECD Education at a Glance 2024. — *UK NEET rate of 14% (15–24 age group, OECD methodology) places UK 28th of 34 countries with data. Only six OECD countries had higher rates.*

[3] HM Government (2024). Get Britain Working White Paper. Department for Work and Pensions, November 2024. — *Government source for NEET policy response, Youth Guarantee, and associated cost estimates.*

[4] Milburn, A. (2026). Independent Review of Youth Inactivity — Discovery Phase Findings. Department for Work and Pensions, Spring 2026. — *Government-commissioned independent review. Confirms mental health and disability as primary drivers of rising NEET. Calls for preventive rather than reactive intervention.*

Intrinsic Motivation and Learning

> *The desire to learn is biological. Intrinsic curiosity declines through schooling under the current model.*

[5] Ryan, R.M. & Deci, E.L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist, 55(1), 68–78. — *Foundational study on intrinsic motivation. Establishes that external reward systems systematically undermine intrinsic motivation — the 'crowding out' effect.*

[6] Lepper, M.R., Greene, D. & Nisbett, R.E. (1973). Undermining children's intrinsic interest with extrinsic reward. Journal of Personality and Social Psychology, 28(1), 129–137. — *Classic study demonstrating that introducing external reward for an intrinsically motivated activity reduces subsequent intrinsic motivation.*

[7] Hidi, S. & Renninger, K.A. (2006). The four-phase model of interest development. Educational Psychologist, 41(2), 111–127. — *Evidence that interest (a form of intrinsic motivation) develops in phases and can be supported or undermined by educational environments.*

Neuroscience — Body-Mind Connection

> *Chronic stress, sleep deprivation, hunger, and lack of physical activity structurally impair the developing brain and learning capacity.*

[8] McEwen, B.S. (2008). Central effects of stress hormones in health and disease. European Journal of Pharmacology, 583(2–3), 174–185. — *Evidence that chronic cortisol exposure affects prefrontal cortex development in children.*

[9] Walker, M. (2017). Why We Sleep. Allen Lane. — *Comprehensive synthesis of sleep research. Documents impact of sleep deprivation on memory consolidation, emotional regulation, and executive function in adolescents.*

[10] Ratey, J.J. & Hagerman, E. (2008). Spark: The Revolutionary New Science of Exercise and the Brain. Little, Brown. — *Evidence that physical movement is neurologically necessary for learning, not merely supplementary.*

[11] Hoyland, A., Dye, L. & Lawton, C.L. (2009). A systematic review of the effect of breakfast on the cognitive performance of children and adolescents. Nutrition Research Reviews, 22(2), 220–243. — *Review of evidence on nutrition and cognitive performance in school-age children.*

Psychological Safety and Belonging

> *Psychological safety is a neurological precondition for genuine learning.*

[12] Edmondson, A.C. (1999). Psychological safety and learning behavior in work teams. Administrative Science Quarterly, 44(2), 350–383. — *Foundational work establishing psychological safety as a precondition for learning behaviour.*

[13] Baumeister, R.F. & Leary, M.R. (1995). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. Psychological Bulletin, 117(3), 497–529. — *Evidence that belonging is a basic human need with direct neurological and behavioural correlates, not a preference.*

Knowledge and Critical Thinking

> *Critical thinking requires substantive domain knowledge. The best-performing systems combine knowledge, reasoning, and developmental support.*

[14] Willingham, D.T. (2009). Why Don't Students Like School? Jossey-Bass. — *Cognitive psychologist's analysis of why knowledge is prerequisite to critical thinking. Memory and domain expertise free cognitive bandwidth for higher-order reasoning.*

[15] Ericsson, K.A. et al. (1993). The role of deliberate practice in the acquisition of expert performance. Psychological Review, 100(3), 363–406. — *Evidence on expertise development. Deep stored knowledge structures ('mental models') are the foundation of expert performance in any domain.*

[16] OECD (2023). PISA 2022 Results Volume I: The State of Learning and Equity in Education. OECD Publishing, Paris. — *Finland, Singapore, and Estonia consistently outperform the OECD average while combining rigorous knowledge acquisition with developmental approaches. Finland's between-school variance: 9% versus OECD average 30%.*

Part B — The Methodology

Section 5 — Why Every Previous Reform Failed

UK Education Reform History

> *The UK has attempted education reform through normal ministerial channels for fifty years without sustained structural improvement.*

[17] Barber, M. (2007). Instruction to Deliver. Politico's Publishing. — *First-hand account of the delivery unit approach under Blair. Identifies systemic barriers to sustained reform including political cycle pressure and institutional inertia.*

[18] Fullan, M. (2010). All Systems Go: The Change Imperative for Whole System Reform. Corwin Press. — *International comparative analysis of education reform. Establishes that surface reform without systemic change consistently fails.*

[19] Higham, R. (2014). Free Schools in the Big Society: The Stratification of School Choice under Libertarian Economics. Journal of Education Policy, 29(3), 321–338. — *Evidence of fragmentation effects of market-based education reform in the UK.*

Goodhart's Law in Education

> *When a measure becomes a target it ceases to be a good measure. This dynamic has operated consistently in UK education.*

[20] Goodhart, C. (1975). Problems of monetary management: The UK experience. Papers in Monetary Economics, Reserve Bank of Australia, Volume I. — *Original formulation of Goodhart's Law.*

[21] Muller, J.Z. (2018). The Tyranny of Metrics. Princeton University Press. — *Systematic analysis of how measurement systems in education and public services produce the gaming dynamic Goodhart identified.*

Section 6 — The Strategic Governance Body

UK Independent Body Precedents

> *The Bank of England independence model, the Committee on Climate Change, and the Office for Budget Responsibility demonstrate that the UK can successfully insulate long-term technical decisions from short electoral cycles.*

[22] Bank of England Act 1998. UK Primary Legislation. — *Established operational independence of the Bank of England Monetary Policy Committee. Interest rate decisions transferred from Treasury to independent expert body.*

[23] Climate Change Act 2008. UK Primary Legislation. — *Established the Committee on Climate Change with statutory mandate, independent budget, and direct parliamentary accountability. Model subsequently adopted by 18 countries.*

[24] Cameron, D. (2007). We will keep up the pressure on climate change. Speech to the Green Economy Conference, March 2007. — *Direct quotation: 'In economic policy, everyone can see that independence for the Bank of England has worked. We now need a Bank of England moment when it comes to climate change.'*

[25] Lockwood, M. (2021). A hard Act to follow? The evolution and performance of UK climate governance. Environmental Politics, 30(sup1), 26–48. DOI: 10.1080/09644016.2021.1910434 — *Assessment of the Committee on Climate Change's effectiveness across multiple governments. Identifies transparency, statutory mandate, and expert composition as key success factors.*

International Governance Models

> *Singapore's sustained ministerial model requires continuous one-party government not available to the UK. Finland's independent national agency validates the concept but operates in lower-turbulence conditions.*

[26] Ng, P.T. (2008). Educational reform in Singapore: From quantity to quality. Educational Research for Policy and Practice, 7(1), 5–15. — *Documents Singapore's tripartite Ministry-NIE-school governance structure and its dependence on political continuity since 1965.*

[27] Finnish National Agency for Education (2019). Finnish Education in a Nutshell. FNAE Publications. — *Describes the structure and mandate of Finland's independent national education agency, established 1991, with 10-year strategic planning horizon introduced 2019.*

[28] Sahlberg, P. (2011). Finnish Lessons: What Can the World Learn from Educational Change in Finland? Teachers College Press. — *Comprehensive account of Finnish education governance. Notes that cross-party consensus and social trust are prerequisites for the Finnish model.*

Section 7 — The Testing Imperative

Pre-registration and Independent Evaluation

> *Pre-registration of expected outcomes and independent evaluation are the minimum standards for credible education research.*

[29] Education Endowment Foundation (2023). The EEF Guide to Working with Schools. EEF Publications. — *UK standard for independent educational evaluation, including pre-registration requirements and publication of null results.*

[30] Slavin, R.E. (2008). Evidence-based reform in education: What will it take? European Educational Research Journal, 7(1), 124–128. — *Establishes standards for experimental evidence in education reform and identifies common methodological failures.*

Section 9 — What Can and Cannot Be Changed

Agile Methodology in Public Sector

> *Agile iterative approaches produce better outcomes than waterfall project management for complex public sector technology programmes.*

[31] Government Digital Service (2019). Government Service Standard. Cabinet Office. — *UK government standard for digital service development specifying agile, iterative approaches as the mandated methodology.*

[32] National Audit Office (2013). Universal Credit: Early Progress. HC 621. HMSO. — *Case study of waterfall project management failure in major UK government technology programme. Established the case for agile approaches in government IT.*

Part C — The Inquiry Agenda

Section 11 — The Immediate Opportunity

Oak National Academy

> *Oak National Academy has built a complete national curriculum content library, publicly owned and open licensed, used in the majority of UK schools.*

[33] Oak National Academy (2025). Usage Statistics 2024–25. thenational.academy. Available at: https://www.thenational.academy/blog/impact — *Usage data: 200,000 teachers used Oak between January and July 2025, with over 54,000 weekly users (53% rise year on year). 4.2 million pupil lesson engagements.*

[34] Department for Education (2025). Independent Review of Oak National Academy. DfE, September 2025. Available at: gov.uk/government/publications — *Independent assessment concluding that Oak's pupil-facing provision requires 'significant refresh'. Recommends development of a pupil impact strategy. Confirms completion of full national curriculum content library.*

Adaptive Learning Evidence

> *Adaptive learning platforms demonstrably outperform static content delivery. DreamBox users improved mathematics scores. High Tech High achieves 82% college acceptance rates.*

[35] Wang, H. & Woodworth, K. (2011). A randomized controlled trial of two online mathematics curricula. Society for Research on Educational Effectiveness. — *RCT establishing statistically significant positive effect of DreamBox Learning on MAP mathematics scores. Cited by US What Works Clearinghouse.*

[36] Center for Education Policy Research, Harvard University (2016). DreamBox Math Impact Study. Harvard CEPR. — *Establishes linear relationship between DreamBox usage (60 min/week) and MAP score improvement. 7.5 point gain at recommended usage level across 3,000 students.*

[37] Lenard, M. & Rhea, A. (2019). Adaptive Math and Student Achievement: Evidence from a Randomized Controlled Trial of DreamBox Learning. Working Paper. — *RCT across 12 schools, approximately 8,000 students. 0.12 standard deviation improvement on early elementary skills test.*

[38] High Tech High (2023). Outstanding Record in College Admissions. hightechhigh.org. — *82% college acceptance rate — highest in San Diego County. 24% increase in four-year college enrolment. 99% of 2024 graduating class met California A-G university requirements.*

[39] Ofsted (2023). XP School Inspection Report. June 2023. — *XP School Doncaster rated Good overall, Outstanding for behaviour and personal development. Previous 2017 inspection: Outstanding overall. Project-based learning model in state school serving deprived area.*

Digital Divide

> *During COVID school closures, access to online learning platforms for economically disadvantaged students fell to approximately 40%.*

[40] National Audit Office (2021). Supporting children's education through COVID-19. HC 136, July 2021. HMSO. — *Documents differential access to online learning during COVID school closures. Confirms correlation between deprivation and digital access.*

[41] Pensiero, N. et al. (2020). Learning inequalities during the Covid-19 pandemic. UCL Institute of Education. — *Evidence on differential learning loss by socioeconomic background during school closures. Establishes case for device provision as equity requirement.*

Section 12 — What We Already Know Works

Finland Early Years Model

> *Finland does not begin formal academic instruction until age seven. Between-school variation is 9% versus OECD average of 30%. Finland consistently outperforms the OECD average in mathematics, reading, and science.*

[42] OECD (2023). PISA 2022 Country Note: Finland. OECD Publishing. — *Finland scores 490 in reading (OECD average 476), 511 in science (OECD average 485). Between-school variance 9% versus 30% OECD average.*

[43] Finnish National Agency for Education (2024). PISA Results Reflect Broader Changes in Finnish Society. FNAE Blog. — *Finland's own assessment of declining PISA scores. Notes between-school variance remains low. Attributes decline to rising socioeconomic inequality rather than model failure.*

South Korea Warning

> *South Korea has among the highest PISA scores and among the highest rates of youth depression and suicide in the OECD.*

[44] OECD (2023). Health at a Glance 2023. OECD Publishing. — *South Korea suicide rates among young people. Contextualises mental health cost of high-pressure academic systems.*

[45] Kim, S. (2016). Education fever and after-school tutoring in South Korea. KEDI Journal of Educational Policy, 13(1). — *Documents the 'education fever' phenomenon and its social costs in South Korea.*

Section 13 — The Technology Landscape

Khanmigo and AI in Education

> *Khanmigo grew from 68,000 pilot users to over 700,000 users. Students using Khan Academy showed greater-than-expected gains on standardised assessments.*

[46] Khan Academy (2024). Khanmigo: AI for Education Impact Data. blog.khanacademy.org, 2024. Usage growth data from Khan Academy public announcements. — *Usage and outcome data for Khanmigo AI tutor. Growth trajectory and assessment outcome data.*

[47] Angrist, N. et al. (2020). Stemming Learning Loss During the Pandemic: A Rapid Randomized Trial of a Low-Tech Intervention in Botswana. National Bureau of Economic Research Working Paper 28271. — *Evidence on AI-supported learning interventions and assessment outcomes.*

MOOC Evidence

> *The completion rate misunderstanding: MOOC participants who enrol to access specific content and leave having done so are not failures.*

[48] Reich, J. & Ruipérez-Valiente, J.A. (2019). The MOOC pivot. Science, 363(6423), 130–131. — *Authoritative reassessment of MOOC completion rates. Distinguishes between goal completion (high) and course completion (low). Most enrollees achieve their actual goals.*

[49] Coursera (2024). 2024 Impact Report. Coursera Inc. Available at: coursera.org/about/impact — *Platform data on learner outcomes, employer partnerships, and credential value. Evidence of shift toward employer-linked credentials.*

Adaptive Learning Market

> *The adaptive learning market was estimated at approximately £3 billion globally in 2024 and is projected to reach £25 billion by 2034.*

[50] Global Market Insights (2024). Adaptive Learning Market Size and Forecast 2024–2034. GMI Research. Note: multiple independent market research firms produce comparable estimates for this sector. — *Market sizing report. Multiple independent forecasters cite similar trajectory.*

Section 16 — The Commercial Threat

Surveillance Capitalism

> *Commercial AI systems collect extensive behavioural data from children. The incentive architecture of commercial platforms is structurally misaligned with building learner autonomy.*

[51] Zuboff, S. (2019). The Age of Surveillance Capitalism. Profile Books. — *Foundational analysis of behavioural data extraction as the commercial model of platform capitalism. Identifies education technology as a primary vector.*

[52] Williamson, B. (2017). Big Data in Education: The Digital Future of Learning, Policy and Practice. SAGE. — *Specific analysis of data practices in education technology. Documents gap between declared and actual data use by commercial providers.*

Email: info@growthuk.uk