The Long Childhood

On the Convergence of Humanity

Krishna Pagadala[1]

June 2026

In memory of Richard A. Easterlin (1926–2024), who asked the right question.

In memory of Jane Goodall (1934–2025), who saw what the juvenile dependency period does.

In memory of Mahbub ul Haq (1934–1998), who knew that people are the real wealth of a nation — and left one step for the rest of us.

Abstract

Humans have the longest juvenile dependency of any species — nearly two decades of plastic brain growth embedded with adult teachers. This window is the channel through which one generation installs a cognitive architecture in the next deep enough to use, extend, and transmit. Wallace's theory explained how species differentiate under isolation; I document the inverse — what one species has done with its window, post-1960,

[1] Software engineer of twenty-five years' practice, self-taught in the substantive literatures over the same period; the framing tested here emerged across it (see §1.3). https://educationfirst.world

at population scale. By 2022, 154 of 185 countries — across every political system and continent — have crossed a shared developmental threshold: fertility below 3.65 births per woman and life expectancy above 69.8 years (the 1960 United States values). The convergence is the population-scale signature of loading schooling into the long childhood; its absence is the signature of failing to do so. The outcomes I track — fertility, longevity, offspring survival — are the classical fitness components of life-history theory; this convergence is the event I predict and explain. I address the crossing event in the post-decolonization window and its generational predictor; pre-transition local fertility patterns and post-crossing trajectory dynamics are outside my scope.

What is uniquely human is cultural transmission applied to tools that rebuild the environment, and to the specialisation of individuals across them. Fire was the first; agriculture, metallurgy, literacy, and medicine followed — each a culturally-transmitted tool that manipulates the environment at a scale no individual could invent alone. Formal schooling is the latest and highest-bandwidth payload ever delivered through this evolved channel; the convergence is what that delivery produces.

I argue that education is necessary and sufficient for convergence at population scale. **Necessity:** of the 154 countries that have crossed, every one through mass schooling. **Sufficiency:** every country that reaches the educational threshold crosses, at a pace set by state priority — one generation (~28 years) under singular-priority expansion, two to three generations (~60 years) under competing-priority expansion, never under no-priority regimes. Sufficiency holds because an educated population operates every other policy lever at policy speed — adopting new institutions, rewriting rules, taking up new technologies — while an uneducated population cannot, no matter how good the rules layer is. Institutions, markets, and rules are downstream of the population that builds and uses them (Section 5.5). Because the transmission operates through the *home niche* — the near-older humans embedded in the child's daily life across the juvenile dependency window — the one-generation floor is biological (the turnover time), not policy-responsive. An educated home niche carries norms, health behaviours, and the expectation of school completion forward to children who exercise these as household decisions a generation later.

At the household scale the claim is concrete: a household whose parents have completed lower secondary is a developed household, and a devel-

oped country is one in which enough households have crossed this parental threshold. The 185-country universe (1950–2015), of which 144 are observed during their lower-secondary transition (the expansion sub-panel: 1975–2015 country-years where child-year completion sits in $[10\%, 90\%]$), spans four generational depths — one century — and bears the mechanism out. A great-great-grandparent's school completion predicts their great-great-grandchild's child-survival outcomes a century later. Six empirical signatures are deduced from the biology — generational timing, multi-generational persistence, asymmetric disruption, income independence, universality, collective action — and tested against the panel; all six hold. The identification is biological and natural-experimental; the cross-country panel is the population-scale evidence that bears out those predicted signatures across the whole species.

Easterlin (1981) identified mass schooling as the cause of uneven development and asked why the whole world wasn't yet developed; his answer was largely set aside while the field pursued institutions, geography, and factor accumulation. Tested on its own terms with the field's own methods, the alternative collapses: log GDP per capita's predictive power decays within a single generation, and once education's contribution is residualized out, GDP predicts nothing. Provision without educated populations produces no durable outcomes. Institutions are built by educated populations, not the reverse.

Keywords: life-history theory; extended juvenile dependency; fitness components; vertical cultural transmission; developmental niche construction; cultural-transmission niche; home niche; literate cultural transmission; generational lag; mass schooling; convergence; human development; demographic transition

Contents

List of Figures

List of Tables

How to Read This Paper

Two readers are intended. The first is the policy maker who must act today: the ministry official, the foundation programme officer, the multilateral staff member who allocates finite political and fiscal capital between competing priorities and needs to know what the evidence requires of that allocation.

The second is the scholar who reads this as a primary text long after the present arguments have ceased to matter — the reader for whom this is one document in the record of how a species came to recognise the channel through which it became what it is.

Neither reader is served by a single linear path through a long argument. The policy maker drowns; the second reader finds the structure inferable but not declared. This guide is the declaration. Four entrances are offered. Three are organised by the discipline the reader arrives from; the fourth is for the reader whose first question is where the work came from. Each traverses the same argument; the difference is the order of approach and which sections carry the load.

From biology. The species-level claim is established in Sections 2 (the eighteen-year juvenile dependency), 3 (cumulative cultural transmission and the tools that rebuild environments), 4 (formal schooling as the highest-bandwidth payload through the channel), and 5 (the state as the agent that paces the channel). Sections 7 and 8 are the mechanism running in observable time — famines, regime collapses, and country trajectories where the channel is severed, held fixed across institutional variation, or pushed at maximum speed. Sections 13–14 extend the framework to coalitionary capacity and forward-looking convergence. The 185-country panel from Section 9 onward is the population-scale signature of the life-history claim; the biologist may read it as the signature without engaging the regression apparatus that demonstrates it to readers from other disciplines.

From human development. Section 1 fixes the developmental threshold in the tradition of Mahbub ul Haq and Amartya Sen, who already did the philosophical work of dethroning income. Section 7 is the natural experiments, Section 8 is the country trajectories, and Section 9 is their population-scale signature. Section 11 is the engagement with Smith, Sen, Deaton, and the institutional alternative. Section 12 diagnoses why the human-development apparatus has under-weighted its own central variable. Section 16 is what the path is for.

From the decision today. The reader who must allocate this quarter may begin at Section 6 (the six signatures I require of the data), proceed through Sections 7 and 8 (the natural experiments that do the identification work the panel alone cannot, and the country trajectories closest to the reader's situation), then Sections 9 and 10 (the universal panel test and the Soviet hollow-education falsification). Section 16 is the operational claim. Appendix A points the reader to the full reproducibility stack (scripts and 20-test econometric battery on GitHub) for satisfying a technical audience before acting. Sections 13–14 are optional for the operational reader; they argue the same decision delivers peace, but the operational instruction at §16 stands without them.

From the question of warrant. The reader whose first question is where this came from — before what it claims — may begin at §1.3 (how the framing emerged, on what data, across what period) and the Acknowledgments (the intellectual lineage the synthesis was built from, and the AI-assistance disclosure). The rest of the paper is then read against a synthesis whose origin and method are already disclosed.

Across all paths. The evidence runs in two tiers of history and one translation. The deep history (Chapters 2–5) — evolutionary anthropology — establishes the mechanism: what humans are, why the juvenile dependency window admits a payload no other species can carry. The recent history (Chapters 7–8) — the natural experiments and the country histories — shows the mechanism running in observable time. The 185-country panel (Chapter 9) translates that historical pattern into the regression form an economics reader requires; it shows the pattern is universal across the species — the population-scale evidence that supports what the biology and the histories establish. These are not one method applied many times but several pointed from different directions — deduction from the biology, comparison across natural experiments and country histories, regression across the panel. The proof is biological; that the others converge on the same channel, none of them able to see the assumptions of the rest, is why I trust it is the channel and not an artifact of how I looked.

This paper is written for readers who want to evaluate the evidence directly. No specialised training is required to read this.

The argument is one argument. The paths are entrances, not parts. A reader who follows only one entrance will see less than is here. A reader who must act this quarter should follow the third path and accept that the first two are why the third holds.

Glossary

I use several specialised terms drawn from evolutionary anthropology, cultural-transmission theory, and demographic analysis. Each is defined below with a section anchor for its load-bearing use.

Term	**Definition**
Long childhood / Juvenile dependency window	The roughly eighteen-year period of plastic brain growth during which the human is dependent on adult caregivers — the longest juvenile dependency of any species. The substrate every other term in this glossary loads into (Section 2.1).
Agency transfer	The threshold at roughly age fifteen — and for many individuals a year or two later — at which the cognitive architecture absorbed during the long childhood becomes the learner's own operating system rather than the parents' to direct. The window of categorical loading remains open until eighteen; agency transfer is the milestone inside it (Section 2.5).
Cultural-transmission (CT) niche	The set of near-older humans embedded in a child's daily life across the eighteen-year juvenile dependency window — parents, grandparents, older siblings, near-kin, and the surrounding adults of the cooperative-breeding unit (Hrdy 2009): teachers, neighbours, and other community alloparents. Channel and content are inseparable: the near-adults *are* the transmission (Chapter 3).
Home niche	The CT niche carried by the household. Symmetric: it transmits whichever CT regime — literate or illiterate — it inherits (Section 5.1).
School niche	The CT niche constructed by the state through schools, teachers, and curriculum. Adds a second niche alongside the home niche during the transition (Section 5.4).

Term	Definition
Home-niche modulation	The home niche conditions how much of the school niche a child can absorb. Fully illiterate home niches give no scaffolding; literate home niches compound with school (Section 5.1).
Categorical literacy (Level 1)	The categorical reorganisation of the individual brain that formal schooling produces. Literate and pre-literate cognition differ in kind, not degree (Section 4.2).
Literate CT (Level 2)	The categorical regime change at the society scale: cultural transmission itself comes in two incomparable forms. Educated-CT enables non-kin coordination at population scale; uneducated-CT does not, no matter how long it runs (Section 4.3).
Human ratchet	Each generation's schooled cohort becomes the next generation's home niche, raising absorption of the next school niche the state delivers. Self-amplifying, one-way, state-paced (Section 3.2).
Time-to-agency	The roughly twenty-eight-year per-cohort lag between educational investment and developmental outcomes — the interval required for educated children to become adults exercising household decisions. Biological constant per generation; country-aggregate crossing time is compositional and depends on the starting baseline (Section 8).

Term	Definition
β_g	Generational amplification coefficient: the within-country slope of child lower-secondary completion on parental completion at the one-generation lag, with country fixed effects. A one-pp rise in the parental cohort predicts a β_g-pp rise in the child cohort. Exceeds 1 at low baselines (the school niche extends reach above what the home niche alone carries); approaches 0 at high baselines (ceiling compression). The headline panel is post-1975, 144 countries in the expansion window $[10\%, 90\%]$ (Table 5); a long-run 1900–2015 panel on a 28-country subsample with self-determined education policy shows the same pattern at deeper baselines, where the post-1975 panel is weighted toward countries already approaching the ceiling. Per-country 28-year sliding windows (Figure 4) trace the within-country trajectory. Same regression specification, three samples.
Four radii of educational effect	Self and children (action), close relatives (action), polity (political pressure), humanity (talk). Each is a boundary draw of the same coalitional in-group mechanism (Tooby & Cosmides 2010); potency falls as the boundary widens (Section 5.2).
Three state regimes	No priority, competing priority, singular priority. The state's choice sets the pace at which the school niche reaches the home niche (Section 5.4).
Hollow education	School systems where reported completion runs ahead of cognitive depth, so the certificate is real but the literate-CT regime has not flipped. The Soviet republics 1960–1990 are the paper's headline case; convergence lags (Chapter 10).

1. The Convergence

Human childhood is the channel through which our species reproduces itself across generations. Education is what we load into it. That is my argument.

Education is fundamental to human well-being. Not important — fundamental. The distinction matters. Important means high on a list. Fundamental means the list does not exist without it.

Between 1960 and 2022, 154 of 185 countries crossed the two thresholds that define human development — the United States values of 1960 (formalised in §1.1). The crossing is what one species does when it loads its long childhood — the eighteen-year juvenile dependency window — with formal schooling at population scale. Figure 1 traces the accumulation. The share of the world's population living in a country that has crossed climbs slowly from 13% in 1961 to roughly 20% by 1993, jumps when China crosses in 1994, passes 50% in 2001, and reaches 80% by the late 2010s. One species, across every political system and every continent, is converging onto a single developmental threshold.

This convergence is the paper's empirical anchor. What varies between countries is the speed of arrival. Spain had wealth, an empire, and global power; Korea in 1953 had none of them. Spain took four-hundred-and-fifty years; Korea took thirty-five. Same juvenile window, same species, same biology. What separated them was the decision to extend schooling.

I argue that the mechanism is biological. Humans alone have a juvenile dependency period long enough — nearly two decades of plastic brain embedded with dedicated adult teachers — for cultural transmission to load the population with knowledge deep enough to produce the convergence (Konner 2010; Hrdy 2009). The pace of that loading is not administrative but biological: it is set by demographic metabolism (Lutz 2013), the rate at which more-schooled cohorts replace less-schooled ones (Section 5.5). Formal schooling is the latest and highest-bandwidth payload ever delivered through that channel. Once an educated population is in place, development is no longer a policy target pursued through a list of interventions; it is the population's own expression of what the juvenile dependency window permits. Chapters 2 through 5 establish the biology. Chapter 6 states the predictions the deep history entails.

Chapter 9 renders the pattern as a 185-country universe with the regression sample restricted to the 144-country expansion sub-panel — universal across regimes and continents. Chapters 7 and 8 are the recent history — natural experiments and country trajectories where the mechanism is observable in time. Chapters 13–14 extend my argument forward. Chapter 11 takes up the frameworks the biological claim must displace; Chapter 12 diagnoses why all of them miss the substrate at once. Chapter 16 states what remains: the decision that no policy calculation dispenses from.

By “education” I mean the educated population, not the school system. The school system is the delivery mechanism; the educated adult is the operative unit. A country can rebuild schools within a decade; it cannot rebuild an educated population in less than a generation (Cambodia, Section 7.2). The mechanism — the home niche, the school niche, and the state extending one to reach the other — is in Chapter 5.

1.1 Defining Development

I define a country as developed when it has simultaneously achieved a **total fertility rate (TFR) below 3.65** and **life expectancy (LE) above 69.8 years** — the 1960 United States values (World Bank WDI).

These two measures are chosen because they are intrinsically valuable — not as proxies for something else, but as the things themselves. The UN’s Human Development Index (UNDP 1990) already treats health and education as intrinsically valuable alongside income. This definition takes that logic to its conclusion: it retains the two outcomes that are ends in themselves and drops the one that is not.

Life expectancy is the most fundamental measure of human welfare. Every other outcome — income, freedom, knowledge, dignity — presupposes being alive. A country where people die at 45 has failed on the most basic term, regardless of its GDP or institutional quality. Life expectancy is not one metric among many; it is the precondition for all others.

The number of children a woman bears defines the shape of her entire adult life in ways that no other variable does. Seven children means two decades of continuous pregnancy and nursing, followed by a further decade of childrearing before the last child is self-sufficient (Section 5.3). The difference between seven children and two is not incremental; it is the difference between a life fully absorbed by reproduction and a life in which the education already received has room to be used — in work, civic participation, or rest.

These are also the two most direct population-level expressions of whether the cognitive transformation produced by education has occurred. Life expectancy tracks whether people have the planning horizon, health knowledge, and behavioural capacity to sustain their own survival. TFR tracks whether women have moved from biological fate to conscious reproductive choice — precisely the cognitive autonomy that underlies every other dimension of human flourishing.

Hans Rosling's Gapminder (Rosling 2018) plotted every country's path on the TFR–life-expectancy plane and dissolved the two-box "developing vs developed" framing: populations travel together along a single diagonal, from high fertility and short life to low fertility and long life. The two-threshold definition used here pins Rosling's diagonal to a specific point: the position occupied by the 1960 United States. Rosling's own account of what drives the movement was multivariate; I argue one variable does it: education.

Every competing metric of development is either a component of life expectancy and fertility — child mortality is a component of LE, nutrition is a cause of LE variation, sanitation prevents child death that shows up in LE — or a consequence of the cognitive transformation that moves them (institutional quality, income, democracy). A part can't outrank the whole. An effect can't outrank its cause. LE and TFR are not the best items on a list of development metrics — they are what the list is measuring.

No country considered developed — by any definition, under any framework — has life expectancy below 69.8 or TFR above 3.65. Even Sen's capabilities — health, bodily integrity, political participation, practical reason — all presuppose being alive and having cognitive autonomy over one's reproductive life. Every capability on Sen's list, examined for its generative source, has educated cognition as its enabling precondition. On his own framework, LE and TFR are not two capabilities among many; they are the enabling conditions for the capability to have capabilities.

GDP is excluded: measuring development through income would assume the answer. For readers who include income in their definition: education predicts GDP per capita as strongly as it predicts LE and TFR (Table 8). Crossing dates under three threshold specifications shift by 10–30 years but the ordering holds. Countries with high expansion rates cross any threshold within 14–16 years. Excluding income makes the test harder. Inside the outcome, income would carry the result mechanically. Outside, it competes against education on equal footing.

Why the 1960 USA? It is the natural baseline: 1960 marks the start of mass decolonization, and the USA was the unambiguous global hegemon. A country that crosses both thresholds has surpassed the richest nation on earth on the metrics that matter for human welfare. The USA in 1960 was mid-transition — its TFR of 3.65 was the baby boom peak, and it did not complete its own fertility transition until 1972. A TFR below 3.65 is itself an achievement: in 1960, only the developed world — Europe, Japan, and a few settler states, roughly a quarter of humanity — had fertility that low. The remaining three-quarters of humanity averaged a TFR above 6 — the human baseline.

The countries that had already crossed both thresholds by 1960 were those with the longest prior histories of mass education: Scandinavia, the Netherlands, France, Switzerland, the United Kingdom, Australia, and New Zealand. Japan — despite a TFR of 2.0 — had life expectancy of 67.7, just below the threshold; it crossed in 1964. By 2022, 154 countries have crossed both thresholds, representing 80% of the world's population. Figure 1 traces this accumulation over time.

1.2 What the Convergence Requires

Two intellectual traditions — one asking what development is, the other asking what humans are — converge on the same mechanism.

In development economics, Haq and Sen broke the identification of development with income when they built the Human Development Index (UNDP 1990). Sen's capabilities framework (1999) made explicit that what matters is what people can do and be, not what they earn — and he pointed to Kerala, Sri Lanka, and China as the right anomalies, welfare outcomes far above what their incomes predicted. Easterlin (1981) identified mass schooling as

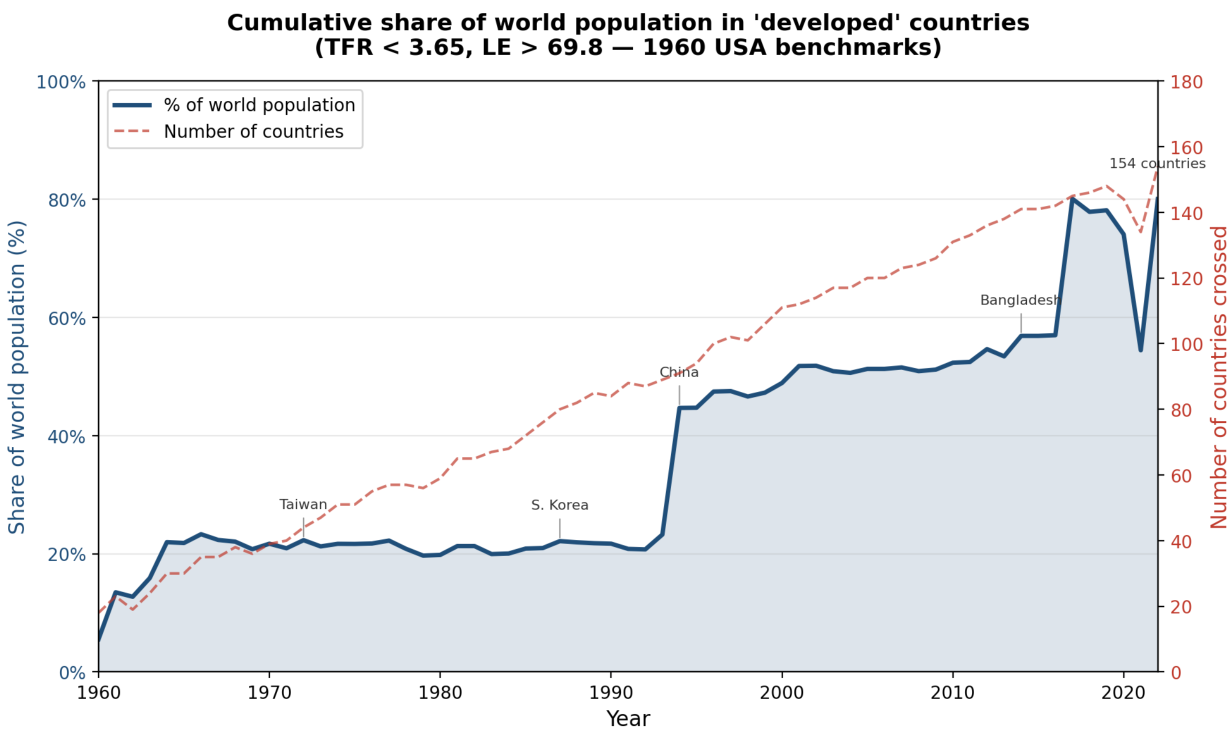

Figure 1: Cumulative share of world population crossing both development thresholds (TFR < 3.65, LE > 69.8), 1960–2022. Slow climb from 13% (1961) to ~20% (1993); vertical jump in 1994 when China crosses; 50% by 2001; 80% by 2017–2022. *(Source: World Bank WDI, WCDE v3.)*

the binding constraint and asked why it had not reached the whole world. Lutz and Kebede (2018) showed that the income–life-expectancy relationship tightens when education replaces income on the horizontal axis. The income–mortality link was education proxying through income all along.

The Human Development Index (UNDP 1990) registered this without naming it. Haq and Sen could have placed any indicator alongside life expectancy and income — sanitation, calories, freedom indices, infrastructure, institutional quality — and chose education. The choice predates the panel evidence by three decades; they reached it by induction from the cases they could see. I supply the mechanism the choice was already tracking: education is co-equal with health and income on the index because it is upstream of both, and because it is the only input that meets the structural criterion (§5.5).

In evolutionary biology, Konner (2010) documented that the human juvenile dependency is the longest of any species — roughly eighteen years of plastic-brain exposure to dedicated adult transmitters, most of it observational rather than formal; schooling is a late, high-bandwidth intensification of the same channel. Hrdy (2009) showed that a dependency this long was only viable because humans evolved as cooperative breeders: mothers, fathers, grandparents, older siblings, and alloparents share the burden of provisioning children that no pair of parents could carry alone. The household I measure is the one Hrdy identifies. Boyd and Richerson (1985) formalised parent-to-child cultural transmission as the mechanism by which human populations accumulate knowledge across generations — at higher bandwidth than the genetic channel, because transmission is continuous across the dependency window rather than a single reproductive event. Hawkes (2003) closed the loop: post-reproductive human longevity, unique among primates, evolved to support grandchild survival, which is why longevity and fertility move together as one developmental signature.

These findings form a chain. Hamilton (1964) formalised inclusive fitness: behaviour that raises the reproductive success of kin is selected even when it lowers the actor's own. Boyd and Richerson extended the logic from genes to culture: knowledge, norms, and the decision to educate are themselves heritable and subject to the same selective pressures, with the adults embedded in the child's daily life the highest-fidelity route. The home niche is inclusive fitness operating through cultural rather than genetic inheritance — kin-directed investment in descendants, carried by the extended juvenile dependency period and supported by the post-reproductive longevity that evolved alongside it.

The two traditions converge in this paper's central claim: that education is the necessary and sufficient cause of human development at the population scale, that income, institutions, and provision are downstream of the same biological channel, and that the channel operates through the cultural-transmission niche (home and school together) across generations. In what follows I state the mechanism and test it.

1.3 Why the Mechanism Has Been Missed

Every major approach to human development — income-first, institutions-first, provision-first, governance-first — was built by people standing on educational foundations they could not see. Adam Smith (1776) built the *Wealth of Nations* on the division of labour while taking the educated workforce that made it possible for granted — a foundational omission Section 3.4 develops. Every framework since has made the same move: taking the products of education as the causes of development.

Japan, with negligible fossil fuels, industrialised within a generation of the Meiji education reforms; Nigeria, sitting on one of the world's largest oil reserves, did not. The policy consequence is dispersion: resources spread across interventions, with education funded as one line item rather than recognised as the mechanism that makes the others work.

Smith's blindspot had a twin. The same framework that took educated labour as given also took abundant coal as given (Wrigley 2010) — but both were outputs of the same channel: centuries of culturally-transmitted tools (schooling, metallurgy, mining, steam engineering), each rendered invisible by the prosperity it produced. Energy and the channel run as a loop, not a one-way street: cultural transmission builds the tools that capture energy, and a stable energy surplus is what lets the channel widen from a literate few to a whole literate population. Coal was both — an output of the channel and the enabler of its next turn. But the surplus that flips a society to mass literacy is thin and specific — synthetic nitrogen and transport, a few per cent of the energy budget, enough to free a child's day from gathering fuel and working the land — not the industrial build-out itself. That loop is the species' deep-history engine; my empirical work is about the operative layer because that is where cross-country variance lives. The channel — from fire to formal schooling — runs under every case I analyse and every prediction I make (Section 2).

A note on this work's provenance. The framing — one species, one mechanism, the convergence as a species-level signature — emerged from a long prior period of reading evolutionary anthropology, human development theory, and adjacent literatures while plotting and re-plotting public country-level data, until the patterns the graphs were already showing resolved, under the combined frames, into a single shape. The data was not the bottleneck; the panels these tests use are public and have been throughout the period in which the question has been asked. Once the framing stabilised, the empirical test in Sections 8 through 9 was built as a software project against public datasets, with the verifier designed in from the start. The six signatures the panel tests took their shape across the long period of reading and re-plotting that preceded writing; once the framing stabilised, they were the operational form of what the mechanism of Sections 2 through 5 requires the data to show if the framing is real. Testing and writing took roughly three months. Workflow specifics are in the Acknowledgments.

2. The Longest Juvenile Dependency

2.1 The Dependency Window

The human juvenile dependency period is the longest of any species. A chimpanzee reaches competence and full subsistence at around seven years. A human takes roughly eighteen — more than twice as long, spent with the body still growing, the brain still being built, and the learner embedded in a household of adults whose daily behaviour is the curriculum (Konner 2010). Across mammals, dependency length tracks brain growth and the complexity of the behavioural repertoire that must be acquired before independence; no other species has stretched the window this far.

The length is not a byproduct of slow biology. It is the channel itself. Eighteen years of plastic brain, embedded with an adult caregiver, is what allows a cultural inheritance of thousands of years of accumulated technique, norm, and knowledge to be installed deeply enough that the learner will transmit it in turn. A shorter window — five years, ten — could transmit a skill. Only eighteen can transmit a cognitive architecture.

The window was built on a substrate of cooperative breeding. Human infants are the most helpless of any primate, and human mothers could not have borne the energetic load alone; fathers, grandmothers, siblings, and non-kin caregivers share care across the window, and the allomothering pattern is itself part of the species adaptation (Hrdy 2009). The grandmother contribution is evolutionarily load-bearing: human females live decades past menopause because post-reproductive women raised grandchildren's survival and transmission odds enough to select for the long post-reproductive lifespan (Hawkes 2003). Dependency, cooperative breeding, and the long post-reproductive adult life are three parts of a single architecture.

A child schooled to lower-secondary completion can, at twenty-five, be trained as a factory worker or in a skilled trade — a welder, a mechanic, an electrician, a heavy-equipment operator. A thirty-year-old who left school at eleven cannot. Specific tasks can be taught; the cognitive architecture that absorbs an industrial training programme or a trade apprenticeship cannot, at population scale, be installed once the window has closed. The skilled workforce of 2050 is sitting in primary schools today. If they are not in school today, there will be no skilled workforce in 2050 — only adults whom no training programme can retrofit.

2.2 How the Window Was Built

A brain that learns for eighteen years is metabolically extraordinary. The human brain represents roughly 2% of body mass and consumes about 20% of resting energy — an energetic share unmatched among mammals. Herculano-Houzel's (2016) comparative work on neuron counts across species shows that the human brain is not anomalous in structure — it is a primate brain at the predicted size. What is anomalous is the energetic budget that sustains it. No other primate can afford a brain this large because no other primate can extract the calories.

Wrangham (2009) identifies the substrate of the substrate: cooking. Cooking raised the net caloric yield of food enough to fund the expensive brain and, critically, to shorten gut transit time so that the energy budget could be redirected from digestion to cognition. Cooking is a culturally-transmitted technique; it requires the channel it later expands. Once the threshold was crossed, the architecture became self-reinforcing: cooking funded the brain, the brain extended the dependency window, the window allowed cooking technique (and everything that followed) to be transmitted with higher fidelity and deeper absorption than any pre-cooking ancestor could manage.

Cooking did more than fund the brain; it freed the day. A great ape spends close to six hours a day chewing raw, fibrous food; a cooking human spends under one (Wrangham 2009). That recovered time is the budget every later elaboration is paid from — the tool sequence, the lengthening childhood, the alloparental care that provisions it, and the cooperative temperament self-domestication selected for. Fire bought the hours; what the species did with them was build the channel and the sociality that runs it. This is the biology, and it was settled in the Pleistocene; nothing that follows in recent history is a change in it.

Every organism is an energy-extraction system — the substrate of life is shared across every species (Smil 2017). What makes the human architecture different is not the energy extraction but the capacity to accumulate and transmit techniques for extraction across generations (Wrigley 2010; Smil 2017). That capacity requires the eighteen-year window. Without it, the architecture of cumulative cultural transmission does not begin.

2.3 What the Window Is For

The window is long enough for cumulative cultural knowledge to be absorbed so deeply that it becomes the learner's own baseline — not content recalled but cognition restructured. Other species transmit behaviour; only humans transmit an installed operating system. What the learner gains by eighteen is not a body of facts (most facts will be forgotten within a decade) but a cognitive architecture shaped by sustained exposure to the accumulated tradition. That architecture: planning horizon, categorical literacy (the brain reorganisation formal schooling produces; see Section 4.2 for the individual level and Section 4.3 for its society-level counterpart, literate CT), symbolic reasoning, and the dispositional expectation that their own children will pass through the same channel.

This is not a description of deficiency. Eighteen years of dependency is not an unfortunate delay before the human becomes useful; it is the channel that makes the adult useful at all. Every capacity the developed adult will later exercise — household planning, health behaviour, reproductive choice, civic participation, transmission to the next child — arrives through what was installed in the window.

The dependency window and the period of cognitive plasticity are the same window because they are the same evolutionary commitment. Parental care of vulnerable young is one of the deepest patterns in the animal kingdom: octopus mothers guard their eggs to death, fish brood and mouth-carry their young, crocodilians defend hatchlings, birds feed and train fledglings, mammals nurse and teach. Across these lineages the same logic recurs — adults invest resources during a juvenile period in which the young cannot yet survive alone. Some lineages, birds and mammals most extensively, upgraded this template by linking the juvenile period to learning: the loading phase during which the young acquires the skills it will use as an adult. This split defines a dividing line. Where the adult repertoire is innate — the octopus that never meets its mother, the megapode that digs itself out of the mound, the cuckoo raised by a host it will never resemble — no juvenile window is needed; nothing has to be transmitted. Where the repertoire must be learned, the window is neither optional nor refundable: a juvenile deprived of the model it must learn from becomes an adult that does not recover the form, however long it is given afterward. Humans extended the mammalian variant further still; the extension itself is what biologists call *neoteny*, a slowed developmental schedule that keeps juvenile traits — including the plastic brain — open longer. Dependency supplies the energetic, nutritional, and social budget that funds the human extension; plasticity is what that budget pays for. Humans pushed almost the entire adult repertoire into the learned column; the length of the childhood is the measure of how much.

Closing both at the same age stabilises the loaded architecture: the adult brain consolidates rather than restructures, because consolidation is what makes the loaded operating system reliably available for action across the rest of life. Plasticity and dependency end together because the architecture is one architecture — the loading phase and the stable-action phase of a deep animal design, with humans running the loading phase longer than anyone else. The policy consequence — that the adult educational stock is fixed once the cohort passes the window — I take up in Section 5.5.

2.4 The Window Supports a Continuous Dose

Educational dose within the dependency window is continuous, not thresholded. The learner who enters at five and leaves at eighteen receives thirteen years of the cognitive architecture; the learner who leaves at nine receives four. The biological substrate does not switch on at nine years of schooling and off at twelve; the substrate is the window itself, and schooling loads into the window at whatever depth the population supports.

Evidence from populations given the full option confirms this. In Singapore, a country that has run the dose to the end of the dependency window for decades, lower secondary completion in the cohort aged 20–24 is effectively universal (~100%), upper secondary completion is ~96%, and tertiary completion is 73%. That is roughly three-quarters of the cohort continuing into formal education past the age at which agency transfers from parent to child.[2] When the option is given, most of the cohort takes it. Singapore is not unique. Tertiary completion in the 20–24 cohort in 2020 sits at 73% in Taiwan, 57% in Sweden, 54% in Republic of Korea, 47% in Norway, and 32% in Japan (WCDE v3, both sexes) — between a third and three-quarters of the cohort continues into tertiary education across cultures and political systems wherever the option is open. This is what a biological architecture for continuous cultural loading looks like at the population level: deeper loading is chosen wherever it is available.

The nine-year (lower-secondary) measurement I use is the empirical floor at which the mechanism reliably crosses the development threshold (Section 5.6); it is not the ceiling of what biology supports.

[2] WCDE v3, both sexes, age 20–24, 2020 vintage.

2.5 The Threshold of Agency Transfer

The eighteen-year window has an internal milestone that matters for measurement. By roughly age fifteen, three things have substantially changed. Reproductive biology is online: menarche in modern populations sits at twelve to thirteen, down from a pre-industrial sixteen to seventeen (Eveleth & Tanner 1990; Walker et al. 2006). Pair-bonding, peer identity, and the dispositional baseline carried into adulthood are forming (Konner 2010). And the categorical reorganisation of cognition that schooling produces (Section 4.2) has largely consolidated.

The cross-cultural record points to the same threshold. Hunter-gatherer boys reach independent productivity by roughly sixteen (Kaplan et al. 2000; Hill & Hurtado 1996; Marlowe 2010); adolescent girls take on substantial allomothering well before reproductive age (Hrdy 2009; Kramer 2011). In agricultural societies, marriage at or near puberty was the default until the twentieth century (Goody 1976), with the northwest-European late-marriage exception (Hajnal 1965) the well-known counter-case. Across subsistence regimes, humans organise reproduction and adult role-taking from the early-to-mid teens.

The claim is not that agency transfers on a fifteenth birthday, nor that the CT regime is locked at fifteen; for many individuals the transition lands a year or two later, and later schooling and adult experience continue to shape the regime. The weaker and sufficient claim is that by roughly fifteen, enough is in place that fifteen serves as a reasonable floor for measurement.

2.6 Why No Other Species Develops

Many species transmit culture — elephants, orcas, chimpanzees (Chapter 3, §3.1). What none of them transmits is a cognitive architecture for restructuring the environment at scale.

The constraint is the dependency window. A chimpanzee's seven years (§2.1) are long enough to absorb termite fishing, social hierarchy, and the foraging repertoire of the group. They are not long enough to absorb writing, arithmetic, the accumulated technical tradition of metallurgy or medicine, and the cognitive disposition to add to that tradition. The human window — with cooperative caregiving and post-reproductive grandmother contribution stabilising it — is the minimum architecture that cumulative culture requires. No other species has the window; no other species develops.

Chapter 3 turns from the window itself to what has been loaded through it: the tool sequence that has rebuilt the human environment at every scale.

3. Cultural Transmission and the Tools That Rebuild Environments

The window of Chapter 2 is what humans load; this chapter is what gets loaded. The content of the channel is a tool sequence that has rebuilt every environment humans occupy. No other species has a tool sequence of this kind, because no other species has the window that makes a tool sequence possible.

3.1 Cultural Transmission Across Species

Cultural transmission — learned behaviour passed across generations — is not uniquely human. Elephant matriarchs transmit migration routes and waterhole knowledge across decades of drought cycles that no single elephant would experience. Orca pods have culturally distinct hunting techniques and vocal dialects that persist across generations, differentiated from neighbouring pods by transmission, not genetics. Chimpanzees have regionally varied tool traditions; Goodall's (1986) four-decade record at Gombe documents termite fishing, nut cracking, and leaf medicine varying by community, each learned by juveniles watching adults. De Waal's (2013) primate work extends this to social cognition, coalitional behaviour, and proto-normative enforcement — sophisticated transmission, carried within an ecological niche.

The cooperative dimension of human cognition is continuous with this. Hare and Woods (2020) document how self-domestication produced the cognitive substrate for cooperative learning that humans require: reduced reactive aggression, greater tolerance of strangers, sustained joint attention. Domesticated foxes show the same syndrome when selected over generations for tameness. The cognitive architecture that lets a human child sit in a classroom for thirteen years learning from a non-kin adult is built on a primate sociality that was reshaped over evolutionary time toward cooperative attention.

What every non-human species transmits is behaviour — a technique refined within an ecological niche and passed by observation. What none of them transmits is a cumulative-culture ratchet. The next subsection traces it.

3.2 The Human Ratchet

The distinctive feature of human cultural transmission is that each generation builds on the last without losing prior gains. Tomasello (2014) calls this the ratchet: cumulative cultural evolution that advances and does not slip back. Chimpanzees learn termite fishing; their descendants do not invent termite-fishing 2.0. Humans learn writing; their descendants invent printing, then typewriters, then computers, then the internet, with each layer stacked on the last. The ratchet requires both directions of fidelity — a child who absorbs the tradition deeply enough to transmit it, and a cohort in which enough children absorb deeply enough that the tradition survives any individual's death or failure to transmit.

The ratchet runs through the eighteen-year window of Chapter 2. A five-year window cannot support a ratchet because five years is not long enough to absorb the accumulated tradition before the learner is asked to transmit it. Each tick of the ratchet is a generation; each tick requires the window; without the window there is no ratchet; without the ratchet there is no cumulative culture.

3.3 The Tool Sequence

Fire is the seed case (§2.2). Every subsequent enlargement of the human population rode the same channel: agriculture displaced foraging, animal traction and wind displaced human muscle, fossil fuels displaced all of it (Wrigley 2010; Smil 2017). Each step is a culturally-transmitted tool that manipulates the environment at a scale no individual could invent or reinvent alone, and no individual masters all of them. Cultural transmission distributes tools across a specialising population; the population-level tool-stock grows because each child acquires a different toolkit through the same channel.

Writing compresses time. A technique refined over three generations can be written down and read by someone in a different continent in a different century. Arithmetic compresses space: a single symbolic system allows trade, taxation, navigation, and surveying to coordinate across populations that never meet. Metallurgy, chemistry, medicine, and engineering each load a body of accumulated technique through the same eighteen-year window into a specialised fraction of the cohort. The tool-stock is no longer held in any individual brain; it is held in the population-level transmission system, and each generation's contribution is the sum of every specialisation in it.

Formal schooling is the latest and highest-bandwidth layer. What took centuries to diffuse through observational transmission, and what parents alone could never transmit beyond their own trade, schooling loads into differentiated cohorts within twelve years. The species-level adaptation — cumulative, specialised tool-making through extended juvenile dependency — is the same; only the loading rate and the breadth of specialisation have changed. Mass schooling became materially possible because the fossil-fuel economy produced the surplus to sustain schools, teachers, and the extraction of children from productive labour for the duration of their juvenile dependency. Scotland in 1696 could afford Knox's parish schools. Most of the world in 1696 could not. Chapter 4 turns to schooling specifically — among everything the window has carried, the layer that changed the loading rate itself.

3.4 Specialisation Requires Loaded Labour

Smith (1776) opened his treatise with the pin factory: ten workers specialising produced thousands of pins per day where one unaided worker produced only twenty. The division of labour was the engine of the wealth of nations. Smith observed specialisation and built a wealth theory around it. He did not ask where the literate, numerate, trade-aware workers came from.

Scotland's 1696 Education Act had been producing them for eighty years before Smith was born. Knox's network of parish schools, extended through the eighteenth century into the Scottish system of free primary education, had done something more than teach a share of the population to read: it had flipped Scotland's cultural-transmission regime. By Smith's lifetime Scotland was operating inside literate CT (Section 4.3) — schooling was widespread and normative, not universal and not uniform, but enough of the population lived inside the new regime that contracts could be written, quantities measured, and specialised roles held across the non-kin relationships that specialisation requires. The pin factory worked because the ambient regime supported it, not because every worker had identical attainment. What Smith saw as the cause of wealth — specialisation — was a downstream consequence of literate CT having replaced illiterate CT as Scotland's transmission floor.

Every industrial revolution since has followed the same pattern. Populations without prior cultural-transmission depth do not specialise into productive factories; populations with it do. The Meiji reform put schools before factories. Korea's post-war expansion put universal primary education in place before the first heavy industrial plant was built. Where the substrate is absent, the factory produces nothing but frustration; where it is present, the factory runs. Chapter 9 will show what this looks like at the 185-country scale.

4. Education as the Highest-Bandwidth Payload

Chapter 3 traced the tool sequence that has loaded through the eighteen-year window and ended on its latest layer, formal schooling. I ask what it does in the window that observation, imitation, and apprenticeship cannot.

4.1 Why Schooling Is the Highest-Bandwidth Layer

Twelve years of structured cognitive immersion delivers what centuries of observational diffusion could not. Before mass schooling, literacy spread through apprentice-trader-scribe chains, monastery scriptoria, and the slow diffusion of devotional texts through communities that built religious reading aloud; the doubling time for literacy rates in pre-schooling Europe was on the order of centuries. Within a generation of universal schooling, populations went from near-zero functional literacy to near-universal. In the same window they acquired numeracy, the capacity to read a contract, the capacity to follow a written health protocol, the capacity to hold a symbolic role in a bureaucracy, and the disposition to expect their own children to do the same.

The bandwidth advantage is structural. Observational transmission loads one technique at a time into one learner at a time, and the learner acquires only what is within sight of an adult already fluent. Schooling loads a cognitive architecture into an entire cohort in parallel, using specialised teachers whose job is transmission itself. Knowledge that previously required direct access to a practitioner — arithmetic, mechanics, natural history, the operation of a written legal system — becomes reliably accessible to every child who passes through. The channel did not become faster; the architecture of the channel changed. Schooling is the first transmission technology in human history that scales transmission itself.

4.2 Categorical Brain Reorganisation

What schooling delivers is a categorical reorganisation, and the categorical claim operates at two levels: in the individual brain and in the society's cultural transmission itself. I take up the first (the brain) here; Section 4.3 takes up the second (the regime). Both are load-bearing, and neither reduces to the other.

At the individual level, what schooling delivers is not content but a reorganisation of the cognitive architecture that processes content. Herculano-Houzel's (2016) comparative-neuroanatomy work establishes the substrate: the human brain carries the largest neuron count of any primate, and the energetic budget that this neuronal density requires is what the long dependency window and cooperative breeding evolved to sustain. The substrate is evolved; what schooling does is put it to use in ways that never develop without schooling.

Dehaene's (2009) neuroimaging work on literacy shows that learning to read physically restructures the brain's visual processing pathways. A specialised region in the left fusiform cortex — the visual word form area — becomes tuned over years of reading to respond preferentially to letter strings over faces, objects, or other visual stimuli. Adults who never learned to read do not have this specialisation; their visual cortex processes text no differently from any other pattern. The reading brain is a physically different brain from the non-reading brain, produced by sustained exposure to a specific cognitive task during the dependency window. The 2010 extension of this work (Dehaene et al. 2010) traced the timing: the reorganisation begins early and continues to deepen through the schooling years.

Dehaene's (1997) parallel work on numeracy shows the same pattern for quantitative reasoning. Specific parietal-cortex regions activate in numerate adults during symbolic numerical tasks that remain dormant in populations without formal arithmetic training (see also Gordon 2004; Frank et al. 2008 for corroboration from the Pirahã natural experiment: an adult population with no numerical vocabulary cannot reliably distinguish sets of 5 from sets of 7). The capacity to manipulate quantities symbolically — which underlies every aspect of modern economic, scientific, and administrative life — is a culturally-installed cognitive technology that the brain does not develop without schooling.

This is what *categorical literacy* names in this paper: the categorical reorganisation of the brain through formal education, through which literacy, numeracy, and symbolic reasoning become capacities the adult exercises automatically across every domain of life. The reorganisation is not a metaphor. It is a physical restructuring of cortical pathways that schooling produces and whose absence leaves the brain in its unschooled state.

The unit of acquisition is the categorical jump. Each load-bearing proposition schooling installs is a kind-flip from one cognitive state to another, with no halfway position between them. Germ theory: bacteria cause disease, or they do not. The Earth orbits the Sun, or it does not. The natural numbers extend without bound, or they stop somewhere. Atoms compose matter. Evolution organises biological diversity. Geology proceeds on a timescale that dwarfs human history. The laws of motion describe how bodies move. Each is a proposition no observation, imitation, or apprenticeship installs; each requires sustained text-mediated instruction; and once installed, each reorganises the categories through which the adult interprets every domain it touches. There is no halfway position between germ theory and miasma, between heliocentrism and a sky that turns over a stationary Earth, between numbers that continue without bound and a horizon at which counting stops.

The jumps compose. Each rung enables the next, and what schooling delivers across years of exposure is not the deepening of one category but a stack of composed categories rising on the rungs already in place. A child who can compute the area of a circle as the limit of inscribed-triangle areas has acquired and composed five separate jumps: triangle, as a categorical geometric object; area, as a measurable quantity attached to a bounded region;

the natural numbers extending without bound; the limit, as the value an unbounded sequence approaches; and the composition itself — the recognition that as the inscribed polygon's sides go to infinity its area approaches the circle's. None of the rungs is reachable from below it, and each enables the rungs above it. The same shape governs every cumulative tradition. Calculus rests on algebra rests on arithmetic rests on counting. Molecular biology rests on germ theory rests on chemistry rests on atoms. Modern reasoning about infection rests on germ theory plus dose-response plus the abstraction of an invisible cause producing a visible effect. The structure is recursive — categorical capacities, composed onto categorical capacities, all the way up.

What "more years of schooling" delivers, then, is more rungs of the stack — additional categorical capacities, each enabled by the ones beneath it. This is what depth means at the individual level: not graded improvement on one axis, but the count of composed rungs that actually installed by the time the child leaves school. Where the development-economics literature speaks of cognitive depth (e.g. Hanushek's knowledge-capital line, taken up in Section 10.5 and Section 10), the variable being indexed is stack height. A mother whose schooling installed germ theory, dose-response, and the hygiene chain makes different decisions about a sick child than a mother whose schooling stopped before those rungs composed; the difference is categorical at each rung and additive across them. Hollow credentials are credentials issued without the rungs installing; the certificate carries no information about how many rungs of the stack the child actually crossed.

Tuning is graded; capacity is not. Selectivity in the visual word form area deepens over years of reading and varies across readers (Dehaene et al. 2010); the same is true of parietal numerical areas. The categorical claim is not that a switch flips inside any one cortex at some literacy threshold. It is that this evolved-and-graded substrate sustains a categorically composed stack of capacities the unschooled brain cannot construct. Across mammals, neuron counts vary continuously, yet recursive language is human-only — graded substrate, categorical capability. The evolved substrate enables the stack; schooling installs the rungs.

The illiterate counterpart at the individual level is not absence of learning. It is dense, narrow learning. The unschooled adult typically commands extensive practical knowledge: kin relations and the social fabric they imply, ecology and weather and the seasonal rhythm they organise, plants and their uses, animal behaviour and the techniques of domestication and hunt, the crafts a village can demonstrate within itself. Real, often deep within its domain, accumulated across a lifetime of observation and apprenticeship. What this learning cannot do is install the jumps just enumerated. The categorical capacities schooling delivers — propositional, abstract, text-mediated, composed across rungs — are unreachable through observation, imitation, and apprenticeship alone, no matter how long those channels run. The difference between the literate and the illiterate adult at the individual level is not that one knows more and the other less. It is that one carries a composed stack of categorical capacities the other cannot construct.

The next section takes up the corresponding claim at the society scale: that the regime sustained by a population of stack-carrying adults is itself a categorically different cultural-transmission regime, not merely a denser version of the illiterate one.

4.3 Literate CT

The second level of the categorical claim is at the society. Cultural transmission (CT) itself comes in two categorically different forms, and the difference between them is not a difference of degree. No amount of illiterate CT, however long it runs or however densely it loads a population, produces what literate CT produces.

Illiterate CT transmits through observation, imitation, and apprenticeship. Its coordination is bounded by kin and line-of-sight. Its specialisation extends to the trades a village can observe within itself — smith, weaver, miller, herbalist — and stops there. Tools improve, but they do not ratchet: a loom a craftsman builds is carried forward in the next craftsman's head, not encoded in a form the next village can pick up from a page. Knowledge that cannot be demonstrated in person does not travel.

Literate CT transmits through cumulative symbolic encoding. A specialist class — teachers — exists whose job is transmission itself. Coordination extends across non-kin relationships at population scale because the typical counterparty can read a contract, follow a written protocol, and hold a symbolic role in a bureaucracy. Specialisation extends into tens of thousands of distinguishable roles, each trained through structured exposure to accumulated text. Tools ratchet: each generation's contribution is encoded, handed forward without in-person demonstration, and built on by the next. The result is what the historical record shows — pin factories, rail networks, national health systems, semiconductor fabs — none of which an illiterate CT society, however long it runs, ever constructs.

This coordination scale is not fixed; it ratchets with the depth of literate CT, by the same logic as the tool sequence (Section 3.2). Each cohort schooled deeper widens the circle of strangers a person can weigh, contract with, and hold a common role beside. What ratchets is the *scale* of coordination, not the disposition behind it: the cooperative temperament is the ancient, self-domesticated baseline, and education widens whom it can reach, not how benign it is (Chapter 13).

The two regimes are not points on a continuum. They are incomparable kinds of CT. A small literate minority inside an illiterate CT society — monks copying manuscripts in ninth-century Europe, court scribes in Pharaonic Egypt, administrative elites in Mughal India — does not constitute literate CT and does not produce what literate CT produces. The difference lies in whether schooling has become the ambient norm for the next generation, not in how many individual literate brains can be counted in the current one.

The distinction has a deep-history test. Before the flip, a literate elite can build great complexity on an illiterate-CT base — and it runs back. The complexity such an elite builds rests on the elite itself and on a precarious, weather-dependent surplus: when crisis came, the literate class and the knowledge it carried were lost, and the complexity collapsed — the Bronze Age palace economies, Rome into the loss of Latin literacy, the Maya (Tainter 1988; Morris 2010). Over ten thousand years the record is excursion and collapse, excursion and collapse, because the regime never flipped: elite literacy is not literate CT, and what it builds falls back toward the illiterate baseline. What the flip changes is not that complexity becomes possible but that it becomes irreversible — knowledge carried by a whole literate population cannot be lost with an elite, and a surplus that does not fail with the harvest does not starve it. Literate CT is the first regime whose gains hold.

What constitutes the flip is schooling being both widespread (carried by a critical mass of households) and normative (the default expectation for the next generation, not a rare or elite achievement). The household scale gives the claim its natural form: a household in which both parents have completed lower secondary is a developed household. The cognitive capacity to make the fertility, health, and reproductive decisions that yield 1960-US outcomes is installed in the parents; what a developed *country* is, in this account, is a country in which enough households have crossed this parental threshold that the national averages for TFR and LE cross too. Individual attainment inside literate CT varies widely — some primary, some tertiary, some barely literate — and that variation is normal, not evidence against the flip. What has changed is the ambient norm and the transmission pathways the norm sustains.

By definition the asymptote is a country at 100% — every household carrying the parental criterion, every child raised in a developed household. Countries cross the TFR and LE thresholds well below that ceiling because education's effects are community-level, not strictly household-level: shared neighbourhoods, public-health systems run by educated workers, vaccination and sanitation that reach the whole polity, and a labour-market premium for literacy that reorganises norms even where particular parents did not complete. The threshold-crossing record below measures how much of that community-level lift country-scale convergence carries at the moment the averages cross.

Empirically the flip shows up in the threshold-crossing record. Ninety percent of countries that cross the life-expectancy threshold ($LE > 69.8$) do so with lower-secondary completion at or above 42% in the 20–24 cohort; the median is 65% ($n = 84$ after excluding the USSR, Europe, Cambodia, and countries already above threshold by 1960). The fertility threshold (TFR < 3.65) crosses at lower lower-secondary completion (p10 $= 31\%$, median 54%; $n = 88$): TFR responds at the short biological lag to the mother's own schooling and so registers earlier in the expansion; LE responds at the longer lag — its dosing, sanitation, and care-seeking run through the children surviving the under-five window — and so crosses later, by which point completion has risen higher. The grandparent cohort, two generations earlier, sat at a median of 7% lower-secondary completion at TFR crossings, with a tenth percentile near zero — the flip does not wait on a long historical runup. What it waits on is the ratchet moving far enough, fast enough, that lower-secondary completion becomes the home-niche norm for the next generation's parents.

The operational threshold I track — lower-secondary completion across the 20–24 age cohort — is developed at panel scale in Chapter 9; this section names the theoretical object the empirical threshold is a proxy for. Where state-reported completion runs ahead of the developmental threshold, as in the Soviet republics (Chapter 10), the admin number is treated as inflated and the country is excluded ex ante, not relabelled post hoc.

4.4 How the Two Levels Compose

The two levels compose through the household. Level 1 (brain reorganisation) is installed in individuals; the household, with both parents carrying it, is the categorical unit. Level 2 is what educated households produce at country scale: the population aggregate of household categoricals, compounded by community-level spillovers (Section 4.3). The categorical kind-difference between literate and illiterate CT is a comparative and historical claim, not a within-panel one. The within-panel signature is the population-aggregate transition band; the kind-difference is what the band is a transition into. These are not two claims in tension but one phenomenon counted two ways. The aggregate transition band is the additive face the panel measures; the kind-difference is the leaderless turn that band records — the same convergence the panel reads as a column of numbers and that Section 5.4 describes as a flock's turn. One is how the emergence is counted; the other is what it is.

The falsifier of the regime claim is the hollow-education case in Chapter 10, where countries whose administrative reporting outruns observable schooling are excluded *ex ante* on independent reporting-integrity evidence, not *post hoc* on convergence failure. A country with high reported completion, clean reporting, and no convergence would falsify the regime claim; the panel contains no such case. Both regimes self-reproduce through the home niche; the asymmetry is that literate CT is also self-*amplifying* — each schooled cohort raises absorption of the next school niche (§3.2). Illiterate CT reproduces flat; literate CT ratchets forward and holds. Complexity built by a literate elite on an illiterate-CT base — never a flipped regime — rises and runs back instead (Section 4.3).

The dependence runs one way. Categorical literacy without literate CT — a few categorically literate individuals inside an illiterate CT society — does not produce convergence. Literate CT without categorical literacy is not possible: the regime is what a sufficient density of categorically literate brains, embedded in a home-niche norm, constitutes. This is why the paper's unit of analysis is the country-cohort rather than the individual. The categorical claim in its load-bearing form is a claim about societies; the individual-level claim is what the societal claim is constituted from.

4.5 Duration Over Fidelity

The content is the vehicle; the reorganisation is the payload. Heyes (2018) calls literacy and numeracy *cognitive gadgets* — culturally installed capacities that require duration of exposure to develop in any given individual, not fidelity of instruction. The install is slow and deep, not fast and precise.

The vehicle is disposable. A carpenter does not remember tenth-grade biology; a doctor does not remember trigonometry; a lawyer does not remember the periodic table. If the mechanism operated through content, education would depreciate as content was forgotten. It does not: the 1950s-schooled cohort, trained with slide rules and pre-computer physics, made the same health, fertility, and transmission decisions in the 1980s that a 1980s-schooled cohort makes now. The reorganisation was installed by the duration of schooling, not the specific techniques it used. The bandwidth is in the years, not the curriculum: a country that teaches twelve years of almost anything worth structured study will produce categorically literate adults; a country that teaches two years of the optimally-designed curriculum will not.

The Hanushek-Woessmann "knowledge capital" line argues the opposite — that test scores, not years, capture the human-capital content that drives growth (Hanushek & Woessmann 2008, 2012); within that framework, populations that achieve many years of schooling with low test quality have gained nothing. The panel partitions the claim by outcome. For fertility, completion dominates; for child survival and longevity, test scores dominate in a cross-sectional horse race — but those same test scores are themselves best predicted by parental-generation completion 10–28 years earlier, the parental lag the mechanism demands. Tests are not an alternative to years; they are compounded years measured one generation downstream. Section 9.5 develops the partition; Chapter 10, §10.5 develops the downstream-of-completion finding.

4.6 Why Poorly-Nourished Populations Learned to Read Anyway

Every successful educational expansion in the historical record was made by populations that were, by contemporary standards, poorly nourished. Scotland's 1696 parish school population, Meiji Japan's 1872 cohort, Korea's postwar 1950s learners, Cuba's 1961 literacy brigades — none of them had anything like modern nutrition. They learned to read, did the arithmetic, crossed the categorical-literacy threshold, and transmitted to their children.

The reason is structural. The brain energetics that sustain the long dependency window were established roughly two million years ago (§2.2); the neural substrate for categorical literacy is what evolved over that time. Contemporary nutrition affects the ceiling of performance and the speed of acquisition but not the substrate itself (Grantham-McGregor et al. 2007). A poorly-nourished child in 1696 Scotland learned to read because the substrate was already in place; the reorganisation was installed, even if contemporary nutrition would now deepen attainment and accelerate acquisition.

The practical consequence is that hunger, insecurity, and poor nutrition are part of the human baseline, not preconditions that must be resolved before education can begin. Every historical case of successful educational expansion ran through poor nutrition, not after resolving it. The substrate evolved in populations whose nutrition was highly variable and often poor by modern standards; it was not built for abundance. The decision to educate does not wait on the resolution of any of these conditions; resolving these conditions is what educated populations then do.

5. Generational Transmission and the State

5.1 The Generational Transmission Mechanism

A child absorbs culture from whoever is embedded in their daily life across the dependency window. I call this set of near-older humans — parents centrally, but also grandparents, older siblings, near-kin, and surrounding adults in the cooperative-breeding unit Hrdy (2009) describes — the **CT niche**. Channel and content are inseparable: the near-older humans are the transmission, their practices are the content, and the child's brain is shaped by the specific niche it spends its window with. What the child becomes is what the niche carries, sustained for eighteen years.

I claim two CT niches matter at population scale. The **home niche** is the inherited cultural-transmission regime carried by the household's near-older humans. The **school niche** is the regime the state constructs through schools, teachers, and curriculum. In pre-state populations only the home niche exists, and the child grows up inside whatever regime the household carries. After mass schooling begins, children live in two niches simultaneously — they absorb from both. Their cognitive trajectory depends on how much of each they absorb.

Each niche carries one of the two regimes developed in Section 4.3: *literate CT* (written, translocal, externalised, decoupled from any single speaker) or *illiterate CT* (oral, local, embodied, tied to the speaker's presence). What matters for a child's outcomes is which regime their total absorbed exposure tracks.

I argue the home niche does more than carry its own regime forward: it modulates the child's capacity to absorb the school niche. A child whose home niche is fully illiterate enters school without language scaffolding, without books, without literate role models, without a bridge between home and school: they absorb less of the school niche per year of enrollment. A child whose home niche carries partial literacy — an older sibling who finished school, a mother who reads, a grandparent who keeps ledgers — enters school prepared, and absorbs more. A child whose home niche is fully literate has home and school reinforcing each other; absorption is near-complete. This is what the literature describes as the schooling–learning gap (Pritchett 2013; Hanushek & Woessmann 2015): enrollment is exposure to the school niche; literacy is absorbed CT; absorption depends on both niches, not one.

I claim this architecture produces an intrinsic ratchet. Each generation's young adults who pass through school become the next cohort's near-adults in the home niche. The next cohort's home niche is therefore more literate than the previous one, their absorption of the school niche is greater, and the generation they produce is more literate still. The ratchet is self-amplifying, and absent catastrophic state failure it is one-way: literate near-adults cannot become illiterate, and they gate forward what they carry. The state's role, I argue, is to set the pace of clicking. No state contribution: the ratchet stalls. Modest state effort under competing priorities: roughly one notch per generation, crossing in ~60 years. Singular priority: multiple notches per generation, crossing in ~28 years. Section 5.4 develops the regimes; the cases (Chapter 8) trace them.

I take the ~28-year generational cycle to be the biological anchor for the paper's lags: the cycle time for one full generation of children to pass through the school niche and become the near-adults of the next cohort's home niche. The empirical anchor is the cohort-weighted mean age at childbearing, $\overline{\text{MAC}} = 28.8$ years across the expansion-phase panel (Section 9.3); the 28-year cross-generation step is a slight underestimate of the empirical 29-year cycle, but the lag-sweep profile is smooth enough that the results carry through. The anchor is not the interval between policy and response. Policy is the click of the ratchet; the cycle is how long the cohort absorbed under the click takes to become parents themselves. A woman completing lower secondary at age 15–18 has children reaching school age 20–30 years later. Taiwan's development crossing arrived ~20 years after educational expansion began: one generation. Kerala's arrived at ~65 years: three generations of slow ratchet-clicking before enough of the home niches carried literate CT for household decisions to produce population-level development outcomes. Across longer windows the same signature holds: the grandchild's attainment reflects the grandparent's through two sequential cycles (~50 years), the great-grandchild's through three (~75 years).

How does the literacy of the adults produced by the ratchet translate into fertility, life expectancy, and child survival? I locate the routing in the decisions made in the household they head, jointly between both parents, with mothers carrying heavier weight on proximate care. Fathers co-determine schooling enrollment (especially for daughters), resource allocation, and fertility negotiations — without them, educated mothers cannot always act. Mothers dominate proximate care: feeding, sanitation, early-symptom recognition, health-seeking, and the spacing of births. Caldwell's classic result — that maternal education is the single strongest predictor of child survival, independent of income (Caldwell 1979; Cleland & Van Ginneken 1988) — reflects this routing. The composition-by-level result (Section 9.5) operationalises it by outcome — and the routing differs across the three outcomes. TFR moves through two channels on the same biology: the mother's own literacy in her reproductive window (basic literacy makes family size a choosable parameter — lower-secondary depth carries the active within-country signal, because primary saturates earliest while lower-secondary holds the variation during the transition), and the grandparent's literacy acting as an independent floor across baselines (Section 9.7: $\beta_{gp} = -0.048$ at low parental baselines, $\beta_{gp} = -0.011$ in the full panel; both highly significant; the parent's coefficient is everywhere larger but the grandparent's never collapses to zero). LE moves through the mother's own primary literacy at lag 0; in a same-sample horse race the deeper levels rise alongside primary but add no independent LE signal, because 60–80% of their within-country variation is already in primary. U5MR has the most layered structure: the cohort's own primary literacy during the childrearing window, plus an independent grandparental channel carried at the upper-secondary depth — the generation that, twenty-eight years earlier, produced the doctors, nurses, sanitation engineers, and public-health teachers whose work shapes the environment in which the next cohort's children grow up.

I claim the architecture explains three properties of the transmission. Non-depreciation: what the child absorbs persists for life because what survives is not content but the cognitive reorganisation that content exposure produced (Section 9.5); literate near-adults cannot revert to illiterate CT. Amplification: at low baselines, where home niches carry illiterate CT, the state's school niche moves children who have no home-niche literacy at all, and each year of schooling clicks the ratchet by a large increment — the $\beta_g > 1$ result, where

β_g is the ratio of the child generation's education gain to the parental baseline (Section 9.5). At high baselines, home and school niches are both literate; additional state effort produces smaller increments. Catastrophic reset: because literate CT lives in the near-adults of the home niche and the teachers of the school niche, destroying the educated population leaves the next generation's plastic brains with a home niche stripped of literate CT and a school niche with no-one to teach. Cambodia (Section 7.2) is this test: schools rebuilt within a decade, the home niche gutted of literate adults, progress stalled for a generation. Knox did not invent generational transmission; he added a competing-priority school niche alongside the home niche, starting — at a pace slower than any state of the past seventy years — a literate-CT ratchet inside transmission machinery as old as the species.

I take the human baseline — the state before formal education reaches a population — to be high fertility, high child mortality, low life expectancy, hunger, and disease. This is not a description of poverty. It is a description of the species when only the home niche exists and that niche carries the illiterate CT regime. Hunger and insecurity are part of the baseline, not preconditions that must be resolved before education can begin (§4.6; Chapter 8).

I trace the pathway: absorption from the school niche restructures cognition; restructured cognition restructures intent (family size becomes a choice, child survival becomes an expectation); the fertility transition follows; fewer children releases resources per child; those resources raise the literacy of the next generation's home niche, which lifts absorption of the next generation's school niche. Each pass clicks the ratchet.

5.2 From Action to Talk: How Education Reaches Beyond the Household

Education's effects start in the household but do not stay there. The four radii shown in Figure 3 extend education's effects outward from the Household decisions node of the generational loop (Figure 2), with decreasing durability as the radius widens.

The home niche operates through action — the household decisions of educated adults. But education's effects extend beyond the household, and the

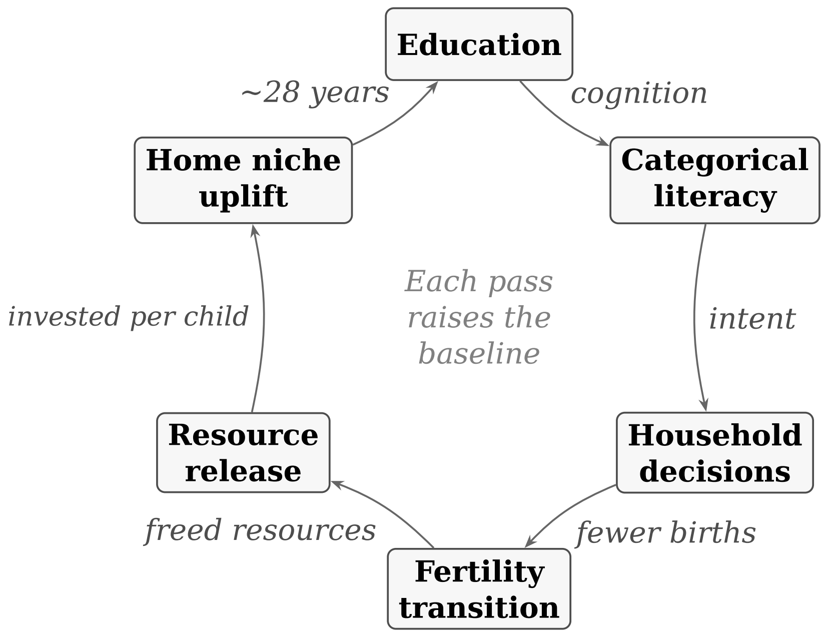

Figure 2: The generational transmission loop. Absorption from the school niche installs categorical literacy in the individual brain; at population scale, the same absorption sustains literate CT as the society's transmission regime. Categorical literacy shapes household decisions; household decisions drive the fertility transition; fewer children per household releases resources; those resources raise the literacy of the next generation's home niche, lifting absorption of the school niche. Each pass through the ~28-year cycle clicks the ratchet one notch.

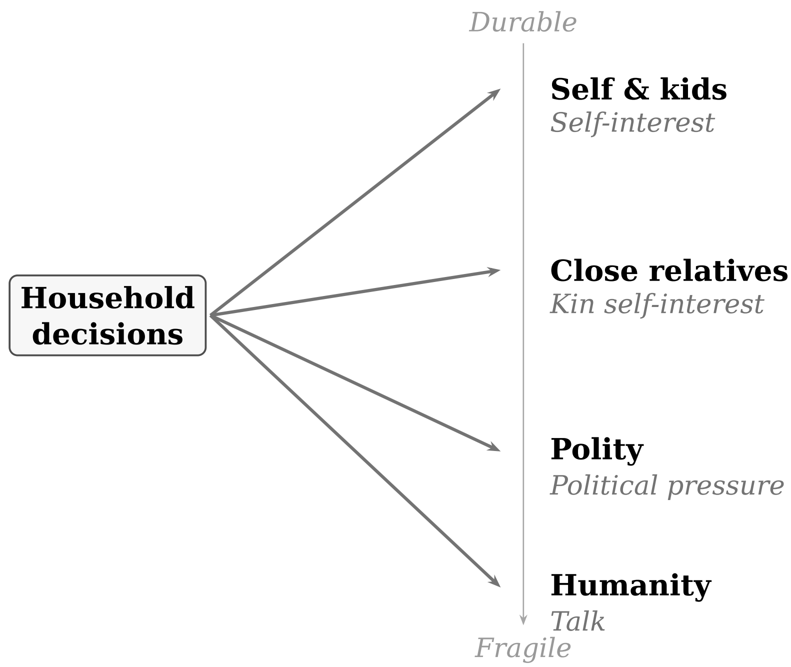

Figure 3: Four radii of educational effect — each a boundary draw of the same coalitional in-group mechanism, with durability decreasing as the boundary widens: action at the kin scale, political pressure at the polity scale, talk at the humanity scale.

nature of that extension changes as the radius widens. Two things widen it simultaneously: education expands the boundary of who counts as "my people" — coalitional psychology (Tooby & Cosmides 2010) allows flexible group boundaries that education stretches from kin to polity to humanity — and it releases surplus, as educated parents have fewer children and nearly all survive, freeing time, money, and attention. The arc widens (who you invest in) and the constraint loosens (what you can invest).

The result is four radii of effect, each a boundary draw of the same coalitional in-group mechanism, each weaker than the last:

1. **Self and children.** Genetic in-group. The home niche is action: the parents raise the child in front of them, and the 18-year biological channel guarantees delivery. Runs on self-interest.

2. **Close relatives.** Kin in-group. Action on siblings, nieces and nephews, cousins. Runs on self-interest extended by kinship.

3. **Polity.** Institutional in-group. Education shifts citizens' demands on the polity — voting, organising, professional norms, policy advocacy. The

most likely channel is political pressure: concrete, but mediated through institutions rather than direct action. This is where the competing-priority pathway (Section 5.1) originates.

4. **Humanity.** Species in-group. Just talk: cross-border exhortation, conferences, advocacy. No biological or institutional channel of guaranteed delivery; the outcome depends on whether someone elsewhere acts on the talk.

The axis that predicts durability is motivational: the home niche compounds because it runs on self-interest. Every non-education intervention requires someone to keep serving past the point where self-interest would redirect them. Revolutionary commitment, donor altruism, state fiscal discipline: all are fragile because all run against the grain of normal human motivation. Education runs with it. Non-education intervention buys time; educational investment changes trajectory.

5.3 Demographic Structure and the Fertility Transition

Fertility declines continuously as female education rises, requiring no threshold to trigger. The evolutionary mechanism is a shift in reproductive strategy, not an override of the reproductive drive: in high-mortality environments, the optimal strategy is quantity — have many children because some will die. Education changes both inputs to this calculus. Child survival rises (fewer births needed to achieve the same number of surviving children), and the planning horizon extends — a literate parent can see that investing more in fewer children produces better outcomes than hedging with numbers. The drive to maximise offspring success is unchanged; what changes is the strategy that serves it.

The steepest fertility decline occurs as the cohort crosses into lower-secondary completion (Section 9.5). What this depth of literacy gives the mother is capacity over preferences that already exist — the capacity to read a contraceptive instruction, understand a health protocol, act on a desire not to be pregnant continuously. This is the mother's own-literacy channel; in the within-country panel the active variation sits at lower-secondary depth, because primary is already approaching ceiling for most of the panel period and the year-on-year movement in literacy useful for contraceptive action runs through lower-secondary expansion. On top of it, the grandparent's literacy adds an independent signal across baselines (Section 9.7) — never as large as the parent's own, but never zero — extending the kin/community radius back one generation and giving the household additional capacity to act even where the parent's own schooling is thin. Yet the demographic headwinds are also steepest here: where total fertility rates are typically above 4, population growth is rapid, and the gains from each educated cohort are diluted across a larger and faster-growing next generation. As completion rises, fertility falls, child survival improves, and household resources per child increase. The home niche is a household mechanism. States can create conditions for these decisions; they cannot make them.

The routing is mother-weighted but household-located: both parents act on household decisions; mothers carry heavier weight on proximate care because of the biology of pregnancy, nursing, and the daily allocation of time during the dependency window. Cross-sectional microdata identifies this within a single country at a single time. At the country-panel scale, where state-led schooling expansion moves male and female completion together, the household, not the individual woman, is the empirical unit (Section 9.5).

Human biology enforces the scale of these decisions. A woman who begins bearing children at fifteen and has seven will still be pregnant or nursing at thirty-five — and the last child, born then, still needs daily care when she is in her mid-forties. High fertility is not merely correlated with constrained lives; it is the constraint.

Even in populations where life expectancy at birth is low, women who survive to childbearing age typically live into their fifties or beyond; the average is dragged down by infant and child death, not by adults dying at forty. But survival does not mean freedom. The grandmother hypothesis in evolutionary biology holds that human females live decades past menopause precisely because post-reproductive women raised grandchildren's survival odds. The long post-reproductive lifespan is itself part of the transmission architecture. Cultural transmission selected for extended juvenile dependency; extended juvenile dependency selected for long post-reproductive lives. Together they are the biological signature of the human developmental niche.

Fertility decline is not a proxy for women's empowerment — it is the first expression of it. Education gives capacity over preferences that already exist: contraceptive availability alone does not produce the shift; education does. The further dimensions of empowerment that the literature measures — labour force participation, household bargaining power, political voice, economic independence — follow after fertility falls, as smaller families free time, resources, and attention. Education produces the first move; the rest cascades from it.

5.4 The State: Reach, Not Mechanism

The unit acted upon is the child during the eighteen-year window. The state is not a substitute for the home niche. It constructs the second niche — the school niche — that, combined with the home niche, produces the transition state in which children absorb from both simultaneously. Its reach is the number of children whose childhoods are loaded by both niches together. Its pace is how thickly it builds that niche: how many children enrolled, for how many years, at what quality.

Lutz and Kebede (2018) document the empirical signature of this at the global scale: state-led educational expansion shifts the fertility curve down and the life-expectancy curve up at a speed determined by how quickly the state reaches households whose home niche still carries the illiterate CT regime. The state matters because it sets the pace of the ratchet at the population scale; it does not determine whether, once children have absorbed the school niche, the adults they become will carry literate CT into the next generation's home niche. That is guaranteed by the biology of the dependency window.

What follows from that reach the state does not direct. The convergence is emergent, in the precise sense a starling flock's turn is emergent: it forms from millions of private household decisions, each made for its own reasons, with no one coordinating them and no one — not the state, not the households — perceiving the whole they compose. No household joins a movement; no committee convenes the fertility transition. The state thickens the school niche and sets the pace; the pattern that follows is the population's own, leaderless and unwilled.

This emergence is older than development. Self-organised collective order is the ancient mode of a social species: the same evolved sociality that lets a child learn from a non-kin adult (Section 3.1) runs the population's coordination with no coordinator, and it ran the illiterate world — high fertility, high mortality — as leaderlessly, for millennia, as it runs this one. Biology supplies the emergent form, common to both regimes; literate CT supplies only the direction — which equilibrium the flock turns toward. The convergence is what an evolved social species does once the regime flips, not an anomaly needing its own cause.

The pattern is unwilled, but the decisions it is built from need not be; Section 16 is the case for making them deliberately.

The operative question, in any country that has not crossed the development threshold, is not whether the state will invest in education — all states claim to — but whether it invests with *singular priority*. Three regimes recur:

- **No priority.** The state does not build schools, does not train teachers, does not ensure access. Only the home niche exists, and it gates forward whatever CT regime it inherited. Pre-industrial populations lived in this regime for millennia.
- **Competing priority.** The state builds schools and trains teachers, but education competes with health, nutrition, security, governance, and other demands for budget and political attention. The school niche exists but is thin: modest enrollment, few years, variable quality. Absorption per child is limited, the ratchet clicks one notch per generation, and the development threshold is crossed at a multi-generational pace. Chapter 8 traces this through Sri Lanka's earlier phases, India, and much of sub-Saharan Africa.
- **Singular priority.** The state makes education the unconditional focus and builds the school niche ahead of need — schools before children exist to fill them, teachers before students exist to teach, standardisation before local pressure demands it. Absorption per child is high because the school niche is thick. The ratchet clicks several notches per generation, and the development threshold is crossed within one. Knox's Scotland, Meiji Japan, post-war Korea, Taiwan, and revolutionary Cuba all ran this regime.

The household operates within whichever school niche the state provides.

Bangladesh from the 1990s illustrates the home-niche/school-niche interaction. The state-led push for girls' secondary education reached households whose home niche carried illiterate CT. The girls who absorbed the school niche then carried literate CT into the next generation's home niche, making the household decisions — fertility, health, investment — that produced the next generation's gains (Bora et al. 2022; Section 8.3). Mullainathan and Shafir (2013) describe the bandwidth tax of the illiterate home niche: evaluating delayed, abstract returns requires the cognitive reorganisation that only absorption from a literate niche produces. The state's reach constructs the school niche; the home niche, once uplifted by a first-generation school-niche absorber, compounds the gain across generations.

5.5 Why Education Is the Limiting Factor

Childhood plasticity is bounded. The dependency window (Section 2.1) loads the cognitive substrate across roughly eighteen years; once it closes, schooling reached is schooling kept. Adult re-entry into formal schooling exists — literacy classes, mid-life degree completion, returning students — but at frequencies far below what would shift a country's educational stock at population scale. The exceptions remain exceptions.

By age twenty-five, education is fixed for life for a given cohort. The country's adult educational distribution at any time is the sum of cohort distributions already past that age. It can be added to only by demographic turnover — births of new children, ageing of existing cohorts, deaths of older ones — the process Lutz (2013) names *demographic metabolism*. Education stock changes at the speed of that metabolism: a country shifts its adult educational mix only as fast as new, more-schooled cohorts displace older, less-schooled ones in the population.

This is not symmetric with other policy levers. Institutions can be reformed in years; commercial regulation can be liberalised in a parliamentary session; tax codes, property regimes, and trade rules can be changed before lunch. The educational endowment of the adult population cannot. It is the slowest-moving variable in the policy ledger and, for that reason, the binding one. The ceiling on a country's developmental trajectory is set by the educational stock it has — not by the institutions, the markets, or the rules layer that change underneath an unchanged input.

This is why development under competing priority feels slow to those who want it. The mix of educated and uneducated adults is what it is at any moment, and the metabolism that would refresh that mix runs at biological speed. Reformers reach for the levers they can move — legal codes, fiscal policy, market regulation — and find that the underlying trajectory does not change, because the input the levers act upon is unchanged. The asymmetry is also why education is sufficient as well as necessary: an educated population can change institutions, rewrite rules, and adopt new technologies at policy speed; an uneducated population cannot, no matter how good the rules layer is.

Education (Section 16) is the only intervention that alters the metabolism's input; singular priority is its maximum-speed setting. A government that gets the rules wrong can fix them next year. A government that gets the schooling decision wrong loses a generation. The institutional rebuttal in Section 11.9 is a consequence of this argument, not an additional one.

This is also why no other variable has, or can have, education's influence. The juvenile dependency window (§2.1) is the only window in which the cognitive substrate is loadable, and education is the only intervention that loads it. Institutions, markets, technology, capability provision — every other lever the discipline names — act on adults already formed, on a substrate already set. They can be applied, refused, optimised, or wrecked; they cannot remake the input. The asymmetry is not a measured effect that further evidence might overturn. It is the structure of the species. A variable that acts on the substrate during the window is in a different category from a variable that acts on the output once the window is closed; that is the category education occupies alone.

The asymmetry runs inside education as well. The distance from no formal education to any formal education is categorical: it is the difference between a brain that has been reorganised by the dependency window's loading and one that has not. The distance from a bad school to a great school is marginal by comparison. The curriculum wars argue over the second contrast and leave the first contrast unmade.

The panel's measurement convention reflects the same logic. The 20–24 cohort age captures education at the point the schooling decision has resolved for that cohort: schooling is substantially complete, upper-secondary either achieved or not, and adult re-entry rare enough that it does not materially shift the distribution (Section 5.6). The country's permanent educational stock for that generation is what the 20–24 cohort distribution shows.

5.6 Why I Measure at Lower Secondary Completion

I measure development at lower secondary completion — nine years of schooling, finished at roughly age fifteen — because that is the point inside the long childhood at which enough has been loaded for the resulting cohort to start acting through it. Three independent considerations converge on that floor. The argument is biological first; data availability follows.

First, the regime floor. By roughly fifteen, the biological, reproductive, and cognitive thresholds of Section 2.5 are substantially in place. The cultural-transmission regime is by then enough to be acting through household decisions — fertility, child survival, the home niche the next generation inherits. Lower-secondary completion is the operational marker: the point at which reading for unfamiliar content, engagement with formal institutions, and numeracy beyond counting become normal household activity rather than skills the household has but does not use.

Second, the empirical floor. Lower-secondary is the level below which full convergence on all three outcomes (TFR, LE, U5MR) does not complete, with further completion compounding the signal continuously beyond it (Lutz & Kebede 2018). TFR crosses at lower lower-secondary completion than LE; the threshold-crossing record and the reason are in §4.3.

Third, the data is available where the threshold is. WCDE v3 supplies long-run lower-secondary completion estimates back to 1875 across the panel; higher-secondary and tertiary series cover narrower windows and fewer countries. The measurement convenience matches the threshold — fortunate, not foundational.

The nine-year floor is not a ceiling on what biology supports. Singapore, Taiwan, Korea, and the Nordic countries run the dose to the end of the dependency window (Section 2.4); where the option is given, most of the cohort continues. The window is eighteen years; fifteen is the floor for when enough of the next generation's regime is in place to act.

6. The Prediction

The claim this chapter tests is that education is the necessary and sufficient population-scale input for convergence.

A causal model, if real, entails empirical patterns that the data must show (Pearl 2009). These predictions follow from the biology set out in Chapters 2–5, not from the data; they are deduced from the mechanism, not fitted to outcomes. Chapter 9 is the formal check. If the mechanism is real — if extended juvenile dependency creates a biological channel for cultural transmission, and formal education is the most powerful payload ever delivered through that channel — then six empirical patterns must hold. The outcomes across which those patterns must appear — fertility, longevity, offspring survival — are the classical fitness components of life-history theory (Stearns 1992).

1. **Generational timing.** Because the home niche is loaded through the juvenile dependency window (Level 1), which closes around age 18 (Section 2.1), and educated children reach household agency — first marriage, first pregnancy, first independent health decision — in the early-to-mid twenties, education at time T should register in development outcomes at $T+\sim 28$ (the biological generation cycle, $\overline{\text{MAC}} = 28.8$ years across the expansion-phase panel, Section 9.3) — and, where the 20–24 cohort at T is already inside its own reproductive and productive window, at shorter lags as well: log GDP registers contemporaneously (lag 0, the educated worker's current output), TFR at $T+5$ (the reproductive peak), and under-5 mortality at the childrearing window (lag 10–15). Life expectancy at birth registers at that same childrearing window (lag 12): it is mechanically dominated by child mortality, so the schooling that moves it is the same schooling that moves under-five mortality, not a separate adult-longevity clock a generation back. The lag signatures are biological and outcome-specific; the independent grandparental channel is separable only for TFR and U5MR (Section 9.3).

2. **Multi-generational persistence.** Because home-niche loading (Level

1) routes through the dependency window in every generation (Chapter 5), a parent's educational depth installs cognition that the child installs in turn, so education's predictive power should decay slowly across generational depths (50, 75, 100 years). Income, which has no transmission channel equivalent in bandwidth or duration to the dependency window, should collapse within one generation.

3. **Asymmetric disruption.** Because the payload lives in the educated adult (Level 1), not in the institutions, income, or infrastructure that delivered the schooling (Level 2) (Sections 5.1 and 5.4), destroying the educated population should reset generational progress, while destroying the Level-2 shell while leaving the educated population intact should not.

4. **GDP independence.** Because GDP is the national-accounts face of the Level-2 (educated-CT) regime — services and the industrial conversion of materials done by educated populations — income carries no separate channel to development outcomes. Once education's signal is stripped from GDP, what remains should have no independent predictive power (Section 9.5).

5. **Universality.** Because both Level 1 (home-niche loading through extended juvenile dependency) and Level 2 (cumulative cultural transmission through population-scale schooling) are species-level capacities (Chapters 2–3), not institutional products, the pattern should hold across political systems, cultures, geographies, and colonial histories.

6. **Collective action.** Household cognitive capacity (Level 1) aggregates into societal coordination capacity (Level 2): the same parents who plan family size and seek timely care also substitute foods under shortage, coordinate migration, pressure authorities for relief, sustain compliance with public-health regimes, and run the institutions that respond to shocks — the Level-2 categorical claim (Section 4.3). Societies that do these things are societies whose households can. The capacity is one mechanism with many testable forms. The strongest test is famines averted under severe food shocks (Section 7.1). Adjacent forms: the fertility and mortality transitions running together, the routine functioning of democratic accountability where mass education has arrived, the defensive capacity that has kept educated populations from being colonised

in the modern record. Where demographic pressure is also present, the same capacity projects organised force outward (Chapter 13). The capacity to avert famine is not a separate faculty from the capacity to run the demographic transitions or to defend against external projection (Section 5.1, Chapter 12) — it is the same capacity at scale.

Educated populations should display this capacity across its range; uneducated populations facing equivalent pressures should not. Where the asymmetry runs in the modern direction — educated populations no longer under demographic pressure, populations still mid-transition under continued pressure — the pressure cannot translate into projection against the educated population, because the projecting capacity is what is missing (Chapter 14). Exceptions to the prediction should require external force physically suppressing the educated population's capacity to act.

Chapter 7 tests them in natural experiments where the design is cleaner than the cross-country regression can deliver; Chapter 9 then shows the population-scale signatures in the 185-country panel.

6.1 Necessity and sufficiency, stated forward

Necessity, stated forward. I claim that lower-secondary completion is necessary for the joint crossing of TFR$<$3.65 and LE$>$69.8. In the panel of seventy-three non-oil joint crossers with available data, I find no country that crossed with lower-secondary completion in the 20–24 cohort below 36% at the year of crossing — 35% rounded, the necessity floor. I therefore predict that none of the thirty-one non-converged countries will jointly cross both demographic thresholds with lower-secondary completion below 35%. Single-threshold crossings — LE-only via oil rents, TFR-only via local shocks — are not counterexamples; the claim is about the joint crossing. A country that lands both thresholds below 35% breaks my necessity claim.

Sufficiency, stated forward. I claim that lower-secondary completion is sufficient for the joint crossing when the 20–24 cohort expands at 1.25 percentage points per year or more from below 10% and the channel runs undisrupted. Among the rate-subset countries that meet these conditions in the historical panel, lags from passing the 35% floor to joint crossing run from sixteen years (Cuba, 1958→1974) to thirty-three years (Indonesia, 1984→2017), with South Korea at thirty (1957→1987) and China at twenty-six (1968→1994). I therefore predict that any future country meeting both conditions will cross both demographic thresholds within thirty years of passing the 35% floor — thirty is my forward bet, tighter than Indonesia's historical thirty-three. A country that meets both conditions and does not cross within thirty years breaks my sufficiency claim. Cases where the channel is disrupted — by war, occupation, sustained net emigration of the educated, or hollow content — do not test sufficiency until the disruption ends; they fall under the hollow-education logic of Chapter 10.

7. The Natural Experiments

The panel establishes the patterns. The natural experiments isolate the mechanism under designs cleaner than cross-country regression can deliver — shocks that isolate one variable while holding the others constant, or within-country cases that eliminate confounders entirely. Each subsection below tests a prediction stated in Chapter 6.

7.1 The Famine Test (Prediction 6)

The famine test is the strongest empirical test of the collective-action prediction (Prediction 6). The shock test asks whether education survives crisis. The famine test asks the converse: does education *prevent* the most extreme form of crisis?

Every major famine since 1950 occurred in a low-education setting. Of 21 famines in the dataset, 19 occurred where lower secondary completion was below 50%. The median education level at the time of famine was 19.6%; the mean was 25.4%. The two exceptions — North Korea (1996, 100% completion) and Yemen (2018, 68%) — both required external force physically preventing food access: totalitarian state control of distribution in the first case, naval blockade in the second. Where the educated population retained agency, famine did not occur.

The near-miss cases make the pattern sharper. Eighteen country-scale near-misses faced shocks comparable to or worse than those that produced famines elsewhere — Cuba's 35% GDP collapse when Soviet imports were cut by roughly three-quarters (1993), post-WWII Japan and Germany, Armenia under simultaneous blockade and war (1993), Ukraine under Russian invasion and grain blockade (2022), Kerala under the 1966 monsoon failure — and none experienced famine. Their median education was 71.6%, against 19.6% for famine countries. The near-miss cases span the same decades as the famine dataset (1945–2022), so the comparison is not an era effect.

Educated populations do not starve. They reorganise supply chains, ration collectively, substitute crops, demand state response, and migrate strategically. Uneducated populations, facing the same shock, lack the household-level capacity to adapt and the political capacity to compel response. The mechanism is the home niche operating through the household: the same cognitive and behavioural repertoire that drives the fertility transition also drives famine prevention.

Bihar and Kerala, 1966. The within-country case eliminates every confounder except education. The 1965–66 national drought crashed India's grain production by 19%. Bihar experienced famine: 70,000–130,000 excess deaths (Dyson & Maharatna 1992). Kerala — India's most food-deficit state, 40% import-dependent, hit so severely that food riots erupted and the UK Parliament identified it as one of the "most severely hit areas" — had no famine. Same Constitution, same free press, same central government, same democratic institutions. Bihar's literacy was 22% (female: 9%); Kerala's was 55% (female: 39%). (The 1966 figures are literacy rates — the only education statistic available at state level for that year — not the lower secondary completion measure used elsewhere in the paper; the gap runs in the same direction under any available measure.)

Sen's claim that "no famine has ever taken place in a functioning democracy" (Sen 1999) cannot explain the Bihar–Kerala divergence. Democracy was constant across both states. Education was not. Myhrvold-Hanssen (2003) challenges Sen directly on this point, arguing that Bihar 1966–67 meets any reasonable definition of famine and that the Indian media, "although free and independent, did not provide reliable information." Drèze himself (quoted in Brass 1986) described the evidence that Bihar was a success story as "precious little."

What the Bihar–Kerala case isolates is that education is a community-level shield, not a national one. The state apparatus was the same across both; the Constitution was the same; the press was the same. What differed was how many households in each state collectively made educated decisions about food, migration, and political pressure. Where most households can and do, the population forms a shield that national institutions alone do not provide. Where most households cannot, no national apparatus substitutes for the missing mechanism. The same pattern extends into the deeper historical record. The Madras Presidency famines of 1876–78 killed an estimated five million people (Davis 2001) under the same imperial administration that was feeding its own educated population at home. Bengal 1943 and the earlier Irish famines of 1845–49 fit the same structure. The population that starves is the uneducated one regardless of which state nominally governs it; the educated population does not starve, even when its government is the one presiding over the starvation.

Kerala's own history confirms the mechanism. In 1943, Travancore (pre-independence Kerala) experienced famine under the same structural vulnerability — extreme food-import dependence — killing approximately 90,000 people. The difference between 1943 and 1966 was not democracy (Travancore had an elected legislature from 1932) but education: literacy had risen from approximately 30% to 55%. The educated population built the institutions — universal public distribution, food committees, political mobilisation — that prevented the same structural vulnerability from producing the same outcome.

Ireland, 1845–1849. Ireland is not an exception to the prediction; it is the deepest historical confirmation, where the cause of low literacy is itself diagnostic. The Great Irish Famine killed approximately 1,000,000 people and drove a further 1,500,000 into emigration over the famine decade (Mokyr 1983; Ó Gráda 1999). Mortality concentrated in the Irish-speaking Catholic west and south and was lowest in the Protestant east of Ulster, the same geographic gradient that ran in pre-famine literacy. The structural cause was not the potato blight, which struck the entire island, but the prior colonial suppression of education. The Penal Laws restricted Catholic schooling for most of 1695–1829; mass schooling for the Catholic majority began only with the National School system in 1831, 14 years before the blight, and reached the famine zones thinnest.

The pre-famine census recorded aggregate literacy at roughly 47% for the population aged five and above, but the distribution was sharply colonial: Protestant Ulster around 65%, Catholic Connacht and west Munster around 33% (Ó Gráda 1995). Mokyr's (1983) county-level analysis found literacy a significant negative predictor of excess mortality while potato acreage did not survive the same specification, with female literacy carrying the strongest signal; Ó Gráda (1999) confirms the gradient. The same Crown, the same Parliament, the same Poor Law, the same press — and a population whose education had been held below the community-shield threshold for 150 years by deliberate state policy. The prediction holds: famine struck where literacy was lowest. What the Irish case isolates, that the post-1950 dataset cannot, is the colonial mechanism that produced the underlying education gradient over 150 years.

7.2 Cambodia: Destructive Disruption (Prediction 3)

Cambodia is the direct test of destructive disruption: what happens when the educated population itself is destroyed, as opposed to shocks that damage institutions, income, or infrastructure while leaving the carriers intact (Section 7.3). Together this subsection and the next test Prediction 3 as a pair — destruction of the carrier resets progress; destruction of everything else does not. Schools are the delivery mechanism; the educated adult is the loading channel. The Khmer Rouge destroyed the second and left the first repairable — buildings came back within a decade, but rebuilding the channel required a generation of its own operation. The country experienced severe disruption under the Khmer Rouge (1975–1979), during which the education system collapsed. The data shows two distinct stalls. The first is the exogenous shock itself. The second — a generation later — is the mechanism's prediction: the home-niche baseline is the binding constraint (Section 5.1).

1975	Khmer Rouge takeover; completion 10.1%
1979	Khmer Rouge falls; educated population destroyed
1985	Completion 9.5% — **first stall** (system collapse)
1993	Paris Accords; UN transitional authority ends
1995	International reconstruction; completion jumps to 35.1%
2010	Plateau at 31–36% — **second stall** (parental baseline)
2011	Recovery as post-disruption cohort's children enrol

First stall (1975–1985). Completion was 10.1% in 1975, fell to 9.4% by 1980, and was still 9.5% in 1985 — a decade of zero progress.

Recovery and second stall (1990–2010). After the Paris Accords and UN transitional authority (1991–1993), schools were rebuilt with sustained international investment. Completion jumped to 35.1% by 1995 on the back of that reconstruction. Then it stalled again. From 1995 to 2010, completion plateaued at 31–36%, unmoved despite continued external funding. The buildings were there; the teachers were there; the money was there. Progress did not come. The 35% plateau is what competing-priority expansion looks like when it interacts with a parental education shadow: international reconstruction dispersed resources across health, infrastructure, governance, and education simultaneously, while the parental cohort remained frozen at ~10%. Korea proves singular priority can push far beyond any parental baseline — but Cambodia after 1991 was not Korea. The plateau reflects dispersed investment meeting a damaged generational base.

The first stall is the exogenous shock — the regime destroyed the education system. The mechanism predicts the second. It is the generational shadow: the children reaching secondary school in 1996–2010 were born approximately 1982–1996 — their parents were the cohort whose own education was frozen at ~10% during 1975–1985 — the pre-disruption level, preserved in households even as the school system collapsed. Countries that started at Cambodia's level in 1960 reached a median of 21% by 1985. The children of the frozen cohort inherited a 10% parental baseline instead of a ~21% one; the plateau reflects that missing growth. Recovery from 2011 onward corresponds to the post-disruption cohort's children finally dominating the school-age population. The regime fell in 1979, but education remained frozen through 1985 (9.5%). A generation (~28 years) from the end of the educational disruption (1985), recovery follows from 2011 — the generational lag the mechanism predicts (Prediction 1).

Countries that started at Cambodia's level in 1960 reached a median of 46% by 2015. Cambodia reached 36%. The buildings came back in 1991. Progress came back in 2011. That twenty-year gap is the home-niche shadow.

The shadow is not over. The grandparent channel (Section 9.7) predicts a second, deeper lag. Children reaching adulthood in 2020–2025 have parents from the post-reconstruction cohort (35–36% completion) but grandparents from the frozen cohort (9–10%). At low baselines, grandparent education predicts child outcomes independently of parent education, comparable in magnitude to but never larger than the parent's own ($\beta_{gp} = –0.048$ vs $\beta_p = –0.068$ for fertility). The Khmer Rouge fell in 1979; the parental shadow lifted in 2011; the grandparent shadow persists through at least 2035, when the first post-disruption cohort's children become grandparents themselves.

7.3 Non-Destructive Shocks (Prediction 3)

Where Cambodia tested what happens when the carrier is destroyed, the shock test asks the converse: what happens when institutions, income, or infrastructure are shattered while the educated population itself survives? Exogenous shocks that kill people depress life expectancy. They do not reverse the fertility transition. The pattern holds across every shock type in the dataset.

Table 1: Exogenous shocks and development outcomes.

Shock type	Country	LE effect	TFR effect	Edu effect
Civil war	Sri Lanka	Delayed 12 yr	On path	None
State collapse	Russia	−5 yr (back)	On path	None
Pandemic (HIV)	South Africa	−9 yr (back)	On path	None
Income destruction	AFC (5 countries)	None	On path	Continued rising
Population destruction	Cambodia	Stalled	Lags	**Reset**

Notes: LE effect: deviation from pre-shock trajectory. TFR effect: deviation from education-predicted path. Sources: World Bank WDI, WCDE v3. Sri Lanka in Section 8.3; Cambodia in Section 7.2.

TFR is a household decision running on the home niche — internal to the household, unreachable by war, virus, or collapsing state. LE depends partly on whether external conditions allow survival. When the two diverge — TFR continuing its education-predicted decline while LE collapses — the divergence identifies the shock as external rather than educational.

Russia (99% lower secondary completion, 1990). The Soviet dissolution crashed life expectancy from 69.5 to 64.5 (1988–1994) — 5 years lost in six. TFR, already at 1.89, continued below replacement throughout — 1.20 by 2000, never reversing. The state dissolved; the educated population did not. LE returned to its pre-collapse level by 2009 (68.7) and reached 73.1 by 2019 — surpassing the Soviet-era peak by 3.6 years. Education predicts the recovery trajectory, not immunity to state collapse.

South Africa and the HIV epidemic. South Africa tests Prediction 3 under a health shock whose biology is exogenous to education. Between 1990 and 2005 LE fell from 62.9 to 53.9 (the largest peacetime LE reversal in modern history) while TFR fell from 3.72 to 2.51 and lower-secondary completion rose from 64.5% to 82.6%. TFR registers contemporaneously to the mother's own depth and tracked the lower-secondary expansion on schedule. LE's behavioural response to a novel pathogen requires schooling depth at acquisition age, set in the prior generation's window — the cohort entering reproductive age 1990–2005 grew up under apartheid education too shallow to deliver it. Post-ARV, LE recovered to 66.1 by 2019. Full mechanism — apartheid stratification of the 1990 aggregate, circumcision and mining-migration acquisition biology, the Population Education Transition curve — is in Section 8.5.

Cambodia (Section 7.2) is the exception that proves the rule: the only shock that breaks the household mechanism is the one that destroys the household.

7.4 The Colonial Test (Prediction 5)

The famous settler mortality instrument does not measure institutions. It measures whether Protestant colonisers — who built schools — survived.

Acemoglu, Johnson & Robinson (2001) use settler mortality as an instrument for institutional quality: where European settlers survived, they built inclusive institutions; where they did not, they built extractive ones. The instrument cannot distinguish this from an alternative channel: where *Protestant* settlers survived, they built *schools*. The Reformation made mass literacy a theological obligation (Section 5.1). The Counter-Reformation removed it. Protestant colonisers (Britain, the Netherlands) transplanted mass education systems. Catholic colonisers (Spain, Portugal, France, Belgium) did not. AJR's instrument captures this difference and attributes it to institutions.

The data separates the two channels. On AJR's 64-country base sample, education at independence (1950 lower secondary completion) explains 52% of the variance in current log GDP per capita; AJR's own institutional measure (*avexpr*, average expropriation risk 1985–95) explains 53%. The two are highly collinear ($r = 0.62$): settler mortality predicts both, because where Protestant settlers survived they built the schools, and the schools made the inclusive institutions possible. The collinearity is the structural finding — the IV identification AJR claim cannot decompose the two channels because the same colonial-era variation drives both. Coloniser religion alone explains 6% of GDP variance, but adding religion to a model that already contains education raises R^2 from 0.518 to 0.521. Religion predicts GDP only because it predicts education. Once education is in the model, religion adds nothing. The channel is religion $\rightarrow$ schools $\rightarrow$ education $\rightarrow$ development.

The 2SLS first-stage diagnostics confirm the structural problem on AJR's actual variable. Coloniser religion is a strong instrument for education at independence (first-stage $F = 10.71$, above the Stock & Yogo threshold) and a borderline one for AJR's avexpr measure ($F = 9.61$, just below). The decisive test is the symmetric one: *avexpr is itself a strong instrument for education* ($F = 37.13$, three times the threshold). AJR's own variable, on AJR's own sample, predicts education with the same first-stage strength required to call any instrument valid. The exclusion restriction is empirically violated: settler mortality and avexpr both reach the outcome through education *and* through institutions, and IV identification cannot decompose them. Full first-stage and second-stage coefficients ($\hat{\beta} = +0.084$, $t = 6.74$ when avexpr instruments education) are in Table 14 and the IV diagnostics footnote below.[3]

[3] IV diagnostics: the bivariate horse race uses coloniser religion as the excluded instrument; education first-stage $F = 10.71$ (Stock & Yogo strong-instrument threshold cleared), *avexpr* first-stage $F = 9.61$ (marginal). The symmetric diagnostic — treating *avexpr* as the instrument for education on AJR's 64-country base sample — gives first-stage $F = 37.13$, three times

Latin America is the critical case. Spanish and Portuguese colonies were settler colonial — Europeans came and stayed in large numbers, built cities, established legal systems. AJR's framework predicts that settlement should produce inclusive institutions and therefore development. It did not. AJR require an ad hoc distinction between "inclusive" and "extractive" settler colonialism to explain why Latin America underperforms. The education account needs no such distinction: Catholic settlers brought no mass education tradition. Spain had 0.6% primary completion for its own 1875 birth cohort; Portugal had 0.1%. They could not transplant what they did not possess. The educational base at independence was correspondingly low (mean 11% lower secondary, 1950), and development followed the education, not the institutions.

The settler mortality instrument is a Protestant education instrument.

Ireland 1845–49 is the within-polity confirmation (Section 7.1): one Crown, one Parliament, one legal system — institutions held constant by being one polity — with mortality tracking the literacy gradient the Penal Laws had carved over 150 years. Mokyr's (1983) county-level horse race on the Irish cross-section already returned the same answer the 64-country sample returns: literacy is the active variable; institutional proxies are not.

The Protestant channel connects to Easterlin's (1981) original insight: the divergence in economic growth between early and late industrializers tracked the divergence in mass schooling — driven originally by the Reformation's insistence on personal scripture-reading, which produced mass literacy across Northern Europe. The distribution of development in 1960 was substantially the distribution of schooling in 1860. Goldin & Katz (2008) document the same transmission within the United States: the "human capital century" of 1910–1940, driven by the high school movement, explains the majority of twentieth-century income growth — a within-country longitudinal result consistent with the cross-country generational mechanism estimated here.

---

the threshold. The same colonial-era variation drives both education and institutions; the exclusion restriction underlying institutionalist IV identification is violated by AJR's own data. Full replication in `scripts/robustness/iv_2sls_colonial_icrg.py`.

8. The Country Histories

The natural experiments and the country histories are both history. The natural experiments isolate the mechanism inside single shocks; the country histories show it running across full development trajectories — schools built before factories, literacy before fertility decline, mothers' education before household reorganisation, in observable historical sequence. The panel's forward lag is a coefficient; in the histories, it is a calendar. The history carries the identification; the panel shows the resulting pattern is universal across the species.

Every country that developed did so by loading its children's long childhoods with formal schooling — fast under singular priority, slower under competing priority, but always through the same eighteen-year window, under every political system, at every income level, on every continent. What is in each country's history is what happened to its children during that window: who built the schools, who reached them, what was loaded into the children, and what the resulting cohort then did as adults.

Table 2: The empirical sequence — each country's schooling commitment and the year it crossed the development threshold.

Country	Dev.	TFR crossed	LE crossed	Onset	Rate	Lag	Gen.
Taiwan	**~1970**	~1970	~1970	1950s	2.15	~20 yr	**1**
S. Korea	**1987**	1975	1987	1953	2.13	~34 yr	**1**
Cuba	**1974**	1972	1974	1961 (40%)	2.27	~13 yr	**1**
Bangladesh	**2014**	1995	2014	1990s	1.30	~24 yr	**1**
Sri Lanka	**1993**	1981	1993	1940s–50s	1.20	~42 yr	**2**
China	**1994**	1975	1994	1950s+CR	1.50	~42 yr	**2**
Kerala†	**~1982**	~1973	~1981	Early 20thC	—	~65 yr	**3**
India	**2017**	1996	2017	1950s	0.87	~67 yr	**3**
Uganda	**—**	TFR 4.39	LE 67.7	None	—	—	—

Notes: Developed = year both thresholds crossed (TFR < 3.65 *and* LE > 69.8, the 1960 United States values, World Bank WDI). Expansion onset = year or level at which sustained educational expansion began. Rate = average annual change in lower secondary completion from expansion onset to development crossing, in percentage points per year. Lag = years from onset to development crossing. Generations = number of ~28-year cycles. CR = Cultural Revolution (1966–76). Cuba's high starting base (40%) and the 1961 reconstruction of losses from the post-revolution exodus shortened the lag below the typical ~28-year generation. † Kerala figures estimated from India Sample Registration System and census records; all other figures from World Bank WDI direct measurement.

The earliest modern example is Japan (1872): the Meiji compulsory education ordinance, implemented through repurposed temples with community teachers, no substantial budget. The sequence was mandate before resources — the state declared compulsory education and built the infrastructure afterward.

Korea is the calibration case — the fastest sustained expansion in the WCDE dataset (Section 8.1). All other rows in Table 2 are independent tests. The expansion rate determines generations to crossing (Section 5.1): at Korea's pace it compresses to one generation; at India's 0.87 pp/yr it spreads across three.

Three parameters predict crossing time: starting base, expansion rate relative to the Korea benchmark, and structural disruption. Disruption delays LE crossing — Sri Lanka's 12-year TFR-to-LE gap (1981→1993) maps onto the civil war; China's 19-year gap (1975→1994) reflects the residual from the 1959–61 famine working through adult cohorts. Korea, Taiwan, Cuba, and Bangladesh had no comparable disruption. The predicted depths match: Korea-pace cases in 20–34 years; half-Korea-pace in 42–45 years; third-Korea-pace in 60–70 years. The crossing dates are the test, not the inputs.

Read together, the cases hold across regime type (authoritarian, democratic, socialist, post-colonial, military), ideology (Confucian, Islamic, Catholic, Buddhist, Marxist-Leninist, Hindu-majority), and income level — $1,038 per capita at Korea's expansion, $1,159 at Bangladesh's crossing. This is Prediction 5 in its most direct form. What the cases add beyond the prediction is the calendar: in each, schooling commitment precedes development by a generation, observably and in order.

8.1 Korea and the Philippines (Predictions 1 and 5)

In 1950 the Philippines was ahead of Korea on income per capita, and started from a comparable educational base — 22% lower-secondary completion against Korea's 25%. By 2000 Korea had crossed every development threshold; the Philippines had crossed none. The income advantage favoured the Philippines; the educational regime that followed did not. What separated them was what each country chose to load into its children's eighteen years.

Taiwan and Korea are the timing and regime-independence test. Both were authoritarian during expansion; both crossed one generation after sustained expansion began. Taiwan and Korea crossed earliest — both at benchmark pace, both state-driven from colonial bases built by Japan (Taiwan 18%, Korea 25% completion by 1950). Korea expanded at 2.13 pp/yr, the fastest sustained rate in the WCDE dataset; Taiwan followed a nearly identical trajectory. Both had functioning markets — what Sen calls "growth-mediated security" — but the growth was mediated by education, not the reverse. State-driven singular expansion compressed development to 20–34 years against the 40–70 years seen in competing-priority cases.

American colonial education in the Philippines built that 22% base over half a century. In 1960, the Philippines was one of the most prosperous countries in developing Asia: GDP per capita of $1,124 against Korea's $1,038 (constant 2015 USD), ahead of Thailand ($592), Indonesia ($598), and far ahead of India ($313) and China ($241). The Philippines had a comparable educational base and the income advantage. But no post-independence government sustained the educational investment. Korea made education the singular priority; the Philippines drifted at roughly half that pace. The result: comparable colonial base, comparable income, divergent post-colonial trajectory. The Philippines crossed TFR in 2003 and LE in 2017, then wobbled around the LE threshold through the pandemic; by 2022 it sat at TFR 1.9 and LE 69.5 — effectively converged, two generations after Korea did so, at roughly half the pace of educational expansion. Education policy is a choice, not an inheritance. Income does not make the choice for you.

Korea also fixes the ceiling. From the 1950 starting position of twenty-five per cent lower-secondary, biology alone — nine years from policy decision to a fully loaded cohort — would have delivered universal completion at roughly twice the pace Korea achieved. Even the WCDE record ran at about half of biological maximum. The ceiling sits well above, not at, what any country has ever delivered. SDG 4's fifteen-year window to 2030 for universal upper-secondary (UN 2015) is three years longer than the biological minimum — correctly ambitious, feasible under singular priority. Where the target has been missed it has been missed under competing priority, not biology.

8.2 Kerala (Prediction 5 at sub-national scale)

Kerala tests the mechanism at sub-national scale: the same transition, the same lag, ~35 years ahead of India as a whole (Kerala ~1982; India 2017). Kerala crossed at ~1982 — the paradigm case of competing-priority expansion. Education had been building since the early twentieth century through social reform movements — gradual accumulation through social reform and state investment — competing-priority, not singular-priority at Korea pace. TFR crossed by ~1973; LE followed ~8 years later. Kerala had extensive state provision — public health, food distribution, land reform — but fertility declined because educated women made the decision, not because provision was available (Section 5.3). The Emergency sterilisations (1975–77) confirm the polarity: forced on the least-educated northern states, they brought down the government (Gwatkin 1979; Vicziany 1982); Kerala needed no coercion. State coercion either arrives after education has done the work, or fails where education has not.

The sub-national test has been run systematically. Drèze & Murthi (2001), using district-level panel data for 1981 and 1991, found that women's education was the single most important predictor of fertility differences across Indian districts, while urbanisation, poverty reduction, and male literacy showed no significant association. The cross-country result replicates within a single country at district level.

8.3 Four Further Cases

Each row below identifies the specific prediction the case most directly confirms.

Case	Crossed	Prediction tested and result
Sri Lanka	1993	**Prediction 3** (non-destructive shock). TFR crossed 1981; LE climbed to 69.0 by 1988, fell to 67.3 during the civil war, recovered to 70.0 by 1993. War disrupted the economy and life expectancy without breaking the home niche.
Myanmar	—	**Prediction 5** (regime independence). 73 years of military rule, GDP $1,025 (2015), education at 0.6 pp/yr. TFR fell from 5.9 to 2.3; LE rose from 44.1 to 65.3 — consistent with 43.5% completion. The regime suppressed speed; it could not reverse what the home niche had already carried forward.
Cuba	1974	**Prediction 5** (socialist regime). 40.3% completion in 1960; post-revolution exodus removed the professional class. The 1961 campaign deployed 268,000 volunteers (Prieto 1981), replacing an elite that left with a mass baseline that stayed. Cuba and Taiwan — Soviet-aligned and US-allied — crossed within four years.
Bangladesh	2014	**Prediction 4** (income independence). 11.4% completion in 1960, GDP $1,159 at crossing; Nepal followed in 2022 at $1,114, even lower. Two crossings below $1,200 in the dataset. Sustained commitment to girls' education from the 1990s. TFR decline tracks female education expansion, not contraceptive distribution (Bora et al. 2022).

8.4 China (Prediction 4 — the support-led residual is education)

China is the direct test of income-independence against the support-led-security thesis. Drèze and Sen attributed China's outcome to direct health provision; the panel shows the residual was education. China crossed in 1994. Three corrections to the standard narrative.

First, what is called an educational catastrophe was the largest educational expansion in Chinese history. The standard framing conflates the disruption of university education — a small percentage of the population — with the overall trajectory. For rural China — 80% of the population (National Bureau of Statistics 1982) — the Cultural Revolution era produced the largest population-weighted lower-secondary expansion in the WCDE 1870–2015 record. Cohort-over-cohort gains run +10.6pp for the birth 1951–55 cohort (in lower-secondary at CR onset), +15.0pp for birth 1956–60 (peak rural 民办 era), and +10.7pp for birth 1961–65 — adding roughly thirteen million additional lower-secondary completers in the 1956–60 birth bracket alone, more than thirty times South Korea's peak single-cohort gain and over a hundred and fifty times Cuba's. The cohort that came of school age *after* 民办 collapsed and before the 1986 compulsory-education law (birth 1966–70) shows the smallest gain in modern Chinese history, +2.3pp. Community schools (民办学校) brought secondary education to villages that had none (Pepper 1996; Unger 1982; Gao 2008).

Second, China is the most direct test of Drèze and Sen's support-led-security thesis. Drèze & Sen (1989, Table 10.6) show China's under-5 mortality at 32% of what GNP predicted — the largest deviation in their sample — and attribute it to "support-led security": barefoot doctors (赤脚医生), direct provision. But Sen regressed on GNP and found a residual. That residual is education.

When life expectancy and under-5 mortality are matched on mean years of schooling (age 20–24, $\pm$0.5 years of schooling, any calendar year; the metric used in this section's outcome-matching, distinct from the lower-secondary completion of the panel chapters — the two move together for China across the period) rather than GNP, China's advantage reverses. China's LE was *below* education-matched peers from 1965 to 1991 — 6.6 years below at 1965, still 2.7 below at 1980, converging only in 1991 — and its U5MR was *above* education-matched peers until 2000.[4] The barefoot doctors were not producing exceptional outcomes; China was underperforming what its education level predicted.

The barefoot doctors and the community school teachers came from the same CR-era rural mobilisation (Gao 1999). The doctors were local, culturally embedded, drawn from the villages they served — optimally delivered provision. When Deng dismantled the system after 1980, they left, because rational people freed from revolutionary ideology redirected their labour toward self-interest. The LE convergence rate was statistically unchanged: $+0.31$ years/year before 1981, $+0.30$ after ($\beta_3 = -0.007$, $p = 0.82$). LE continued to rise — from 64 to 69.8 by 1994 — because the educational baseline built into rural households did not leave with the doctors. Drèze & Sen (1989, pp. 215–221) document the resulting health crises and rising rural medical costs (Gao 2008; Liu et al. 2003), but these crises left no detectable mark on the trajectory.

China crossed to *above* education-predicted LE in 1992 and to *below* education-predicted U5MR around 2000 — two decades after the barefoot doctors were gone. Mean years of schooling rose from 5.9 (1965) to 8.0 (1980) to 9.6 (2000), tracking both outcomes across both eras.

Third, the TFR threshold was crossed in 1975 — five years before the one-child policy. China's fertility trajectory mirrors South Korea and Thailand's, both of which achieved equivalent TFRs without compulsory policy (Cai 2010; Babiarz et al. 2018). The policy did not cause the decline; it was imposed on a transition already in progress, matching peers who reached the same endpoint without coercion.

8.5 South Africa: Apartheid Stratification and the PET Curve

South Africa's 64% lower-secondary completion in 1990 is an aggregate masking apartheid stratification: near-universal white completion and much lower

[4] Mean years of schooling computed from WCDE v3 proportions (age 20–24, both sexes); LE and U5MR from World Bank WDI. Peer pool is all country-years within ±0.5 mean years of schooling of China's value at each year. The result is robust to peer-pool bandwidth: LE and U5MR gaps retain the same signs and comparable magnitudes at ±0.25 and ±1.0 mean years.

Black completion produced a bimodal distribution invisible in national data. The population bearing the brunt of HIV had far lower education than the headline figure suggests. The fall in TFR from 3.72 (1990) to 2.51 (2005) is the lag-difference between TFR and LE operating as predicted. Lower-secondary education drives fertility decline at the contemporaneous lag (panel $R^2 = 0.71$ for lower-sec → TFR at `LAG_TFR`), and the post-apartheid expansion of lower-secondary (64.5% → 82.6% between 1990 and 2005) was running through the shock — the current mothers' own depth was rising fast enough to flip fertility on schedule. The fall is not a frozen home-niche baseline surviving external disruption — 1990 would be too early to carry that claim, with TFR still at 3.72, above the development threshold. It is the ordinary lower-secondary-to-fertility mechanism continuing to run while the prior generation's depth — the cohort at acquisition age when the epidemic hit — remained too shallow in the affected majority to deliver the behavioural response to a novel pathogen that deeper schooling, accumulated earlier, carries.

The LE collapse runs through that secondary deficit. The Southern African epidemic was driven at the acquisition stage by factors education could not have prevented: low traditional circumcision prevalence (8–11% in the HIV belt vs 85–99% in Muslim West Africa), which three randomised trials showed provides a 60% biological reduction in male acquisition risk (Auvert et al. 2005). Circular mining migration also created connected sexual networks across the subcontinent. These are genuinely exogenous to education.

Education was in the story on both sides of the behavioural response. Early in the epidemic, more educated Africans were *more* likely to be HIV-positive: greater mobility, more partners, more agency exercised without information about a novel pathogen (Fortson 2008). Once information existed, the gradient reversed: educated populations changed behaviour first. The Botswana 1996 school-reform natural experiment found each additional year of secondary schooling caused an 8.1-pp cumulative HIV reduction (De Neve et al. 2015). The same pattern appears in rural Uganda, where educated individuals responded more to prevention campaigns, and the education–HIV gradient reversed from positive to negative between 1990 and 2000 (De Walque 2007). Baker et al. (2008) formalise this as the Population Education Transition curve: when a novel health risk enters a population, the educated are early adopters (positive gradient); once information exists, the educated change behaviour first (gradient reverses). The pattern replicates across cigarettes, processed food, and HIV — education is the mechanism in both directions, at the secondary-completion depth the behavioural response requires.

South Africa's 1990 lower-secondary figure, once stratified, sat below that depth for the majority who received the shock.

9. The Panel

9.1 What the Panel Sharpens

I now turn to society-scale evidence: what 185 countries did from 1950 to 2015, read through within-country fixed effects on the 144-country expansion sub-panel. The biology (Chapters 2–5) and the country histories (Chapters 7–8) carry the argument. The panel sharpens it. It is uniquely positioned to do what the case work cannot: show that the population-scale means of education, fertility, longevity, and child survival moved together across the full universe of countries; show that log GDP per capita, once stripped of education's contribution, predicts none of the development outcomes; show that the pattern holds in every region, era, and income band; and show that none of this depends on the estimator — every model class, from regularised linear regressions to a country-blind transformer, returns the same education contribution (Section 9.9). The biology and the natural experiments are what identify the channel (Chapter 7, Chapter 8, and the falsification in Chapter 10); the panel supports that at the scale of the whole species. The objection that the mortality decline is exogenous medical technology is answered separately (Section 11.7).

9.2 The Specification

Every panel result in this chapter comes from one estimating equation. For country i, a development outcome Y, and a lag L set by the biology of the channel,

$$Y_{i,\,t+L} \;=\; \alpha_i \;+\; \beta\,E_{i,t} \;+\; \varepsilon_{i,t},$$

where $E_{i,t}$ is lower-secondary completion (both sexes, age 20–24) in year t, α_i is a country fixed effect, and L is the outcome's biological horizon — 0 for GDP, 5 for fertility, 12 for life expectancy and under-five mortality, and 28 for the child-education step where the outcome Y is the next cohort's own completion (Section 9.3). Outcomes enter in logs where the table so notes. I estimate β within countries and cluster the standard errors by country.

The equation rests on four assumptions, each stated plainly.

- **Within-country variation only.** The fixed effect α_i absorbs everything time-invariant about a country — geography, founding institutions, culture, colonial inheritance. β is read from how a country's own schooling and its own later outcomes move together, never from comparisons between countries.
- **The cause precedes the outcome.** A cohort's lower-secondary schooling is completed in childhood, years before that cohort is measured as young adults. The outcomes I read all fall on or after that measurement — from contemporaneous income out to the next cohort's own schooling a generation later — so the schooling was finished first, a later outcome cannot have produced it, and the forward structure closes the reverse reading.
- **I do not control for GDP, and this is deliberate.** Income sits on the causal path: schooling builds the educated population that produces the income. Holding contemporaneous GDP fixed would block the pathway this paper traces and fold education's effect into a downstream proxy. Section 9.5 makes the point directly — stripped of education's contribution, income predicts none of the outcomes.
- **The lag is not knife-edge.** Each outcome's signal rises and falls gradually around its mechanism's window (Section 9.3); the finding is a property of the mechanism, not an artifact of the exact year read.

This is the whole apparatus the argument needs on the page. The headline robustness — that no estimator, from regularised linear models to a country-blind transformer, moves the answer — is in Section 9.9. The full battery (Frisch-Waugh-Lovell, unique-R^2 decompositions, the 20-test diagnostic suite, robustness sweeps, and the marginal-effect machinery) is in `scripts/`, indexed by `scripts/ECONOMETRICS.md`.

9.3 The Generational Lag

The same regression edu$(T) \rightarrow$ outcome$(T+L)$ measures different mechanisms at different lags L. The panel uses four timescales, anchored on the biological generation cycle.

Because childbearing, childrearing, and the working life each spread across a range of ages rather than concentrate on a single year, and education's effect on every outcome is continuous, the lag profile is smooth: each outcome's signal rises and falls gradually around its mechanism's window, with no knife-edge at any particular year. The specific lag I read an outcome at, within its biologically plausible window, is therefore not load-bearing — the result is a property of the mechanism, not an artifact of lag selection, and the lag-sweeps in this section and the next confirm it. I report each outcome at the center of its window for interpretability, not because the finding turns on the choice.

The biological generation cycle (∼29 years). The cohort-weighted mean age at childbearing pooled across the expansion-phase panel (UN WPP 2024, $n = 1{,}609$ country-years, 177 countries) is $\overline{\text{MAC}} = 28.8$ years. The interval moves only ∼1.0 year across lower-secondary completion bins (29.3 at low education, 28.2 at high) — educated women delay first birth but bear fewer children, and the mean barely moves. This invariance is what makes one generational step a well-defined unit at population scale. The 28-year cross-generation step — the lag for edu-to-edu transmission and the joint forward-prediction headline — is a slight underestimate of the empirical cycle; the lag-sweep profile (Section 9.6) is smooth enough that the existing $T{+}28$ results carry through.

Timescale	What it measures	Mechanism
Contemporaneous (lag 0 for GDP; lag 5 for TFR)	The 20–24 cohort's education and its near-term outcome — log GDP at the same year; TFR at $T{+}5$, when the cohort reaches its reproductive peak.	The cohort is already in its reproductive and productive window; literacy acts in real time. Mothers read contraceptive instructions and act on them; educated workers contribute to current GDP.
Childrearing window (lag 10–15; U5MR and LE)	The 20–24 cohort's education and both under-five mortality and population life expectancy 10–15 years later.	The cohort's own children must be born and survive to age 5 before the survival outcome registers. Life expectancy at birth is mechanically dominated by child mortality, so it moves with the same childrearing-window schooling rather than on a separate adult-longevity clock a generation back.
Cross-generation step (∼25–30 years)	The 20–24 cohort at T and outcomes observed at $T{+}28$: their daughters' education stock, the next generation's TFR via their daughters' decisions, and the next generation's U5MR via the system-level health environment the now-grandparental generation built.	One full generational step. Parent → child education in the language of the lag-decay table; equivalently grandparent → grandchild TFR from the $T{+}28$ vantage.

Outcome-specific lag structure. I read each outcome forward at its own biological horizon — TFR at $T{+}5$, under-five mortality and life expectancy at $T{+}12$, GDP at T — and run a same-sample horse race across primary, lower-secondary and upper-secondary completion in two blocks: the cohort's own schooling at T and the prior generation's at $T{-}28$ (`scripts/residualization/horse_race_lags_levels.py`). Because each outcome sits a biological lag ahead of the T block, that block is the parental generation — it leads the outcome by exactly that biological lag — and the $T{-}28$ block the grandparental one, leading the outcome by that same lag plus the 28-year generational step. This reproduces the per-outcome parent/grandparent lags of the decomposition in Section 9.7 (TFR: outcome-5 / outcome-30; U5MR: outcome-12 / outcome-37). For all three demographic and health outcomes both blocks carry signal: the grandparental block adds 0.091 (TFR), 0.130 (U5MR) and 0.043 (LE) of within-country R^2 above the parental block alone, and the parental block adds 0.114 (TFR), 0.105 (U5MR) and 0.068 (LE) above the grandparental block alone. The cross-generation signal is strongest for child survival, where the grandparental block carries signal at least as large as the parental one — the health environment the grandparental cohort built still moves the survival floor a generation on, the result Section 9.7 states at the coefficient level. Log GDP is the exception: the grandparental block adds essentially nothing once the parental block is in ($+0.011$), because GDP is genuinely contemporaneous — the educated worker's current output. Life expectancy, read at the childrearing window rather than a generation back, behaves like under-five mortality, grandparental block and all: it is mechanically dominated by child mortality — within country, log under-five mortality alone tracks most of its movement and child plus adult mortality together nearly all of it — so the schooling that moves it is the same childrearing-window schooling that moves under-five mortality, and the per-outcome transformer, trained to predict each outcome at its own biological horizon, places life expectancy at that window alongside under-five mortality.

Reading each outcome forward of the education that predicts it — whatever its biological lag — also disciplines the panel against reverse causality: education at T precedes every outcome it predicts, so the arrow cannot run backward.

The leader's policy timeline. A leader who funds today's primary and lower-secondary kids invests in cohorts aged 6–15. Those kids reach age 21 in 15 years and start to dominate the 20–24 cohort; the minimum-visible policy effect is about 15 years. This is a forward-projection statement about a different cohort — the kids in school today — not a panel-testable lag. The panel measures the cohort that already finished; the leader's timeline is about the cohort that hasn't.

9.4 The Convergence in Aggregate

At population scale the six core variables move together (Table 3). Across the 185-country universe, 1975–2015, parental completion roughly doubles and child completion rises by about a third, life expectancy rises by 7.7 years, total fertility falls by about a third, under-five mortality falls by 60%, and GDP per capita rises by about half. The parental column lags the child column by one generation by construction — the early-period parental mean corresponds to the 1950–1964 cohort of the same countries. The chapter reads convergence as this co-movement of population-scale means.

Table 3: Period means, six core variables.

Variable	1975–1989	1990–2004	2005–2015
Parental lower-sec completion (%, $T-28$)	26.4	39.4	53.9
Child lower-sec completion (%, T)	54.5	62.2	71.1
Log GDP per capita (const. 2015 USD)	8.07	8.19	8.52
Life expectancy at birth (years)	62.4	65.5	70.1
Total fertility rate (births/woman)	4.40	3.63	2.95
Under-5 mortality (per 1,000 live births)	91.0	66.8	39.2

Sample and sources. Education from WCDE v3 (Lutz et al. 2021): lower-secondary completion for the 20–24 cohort, both sexes, completed (not enrolled). All percentages refer to this measure unless noted. GDP per capita (constant 2015 USD), life expectancy, TFR, and under-five mortality from World Bank WDI. The headline regressions run on the *expansion sub-panel* — the 144 countries / 945 country-years for which child-year lower-secondary completion sits in $[10\%, 90\%]$, the active-transition window. This is the mechanism's own selection rule: countries already at the ceiling have nothing more to teach the panel about transition, and countries yet to enter have not yet loaded the channel. Full summary statistics and source tables are in Appendix B; the residualisation extension to $T = 1960–1990$ is described in Section 9.5; the long-run 28-country subsample (1900–2015) backs the multi-generational-decay finding in Section 9.5.

The 15 USSR republics. Goskomstat's 1960–1990 lower-secondary reporting is approximately honest for the six European-core republics (Baltics, Belarus, Ukraine, Moldova) and inflated 2.6 to 4.0 standard deviations on log U5MR for the eight Caucasian and Central Asian republics east or south of Moscow's longitude (Chapter 10 develops the falsification). The expansion-window screen removes all fifteen by construction; the headline is unchanged by their exclusion (Section 10.6).

9.5 Education vs. GDP at Population Scale

Where it matters most — countries still building their education systems — education predicts the next generation's attainment three times better than log GDP per capita does. The amplification curve — $\beta_g > 1$ at low baselines, $\to 0$ at saturation — is the population-scale signature of niche-stacking arithmetic. Once log GDP is residualised against education and only the leftover is used to predict outcomes, the leftover predicts essentially nothing. The two findings are the same finding viewed from two sides.

The amplification headline. Table 4 estimates child lower-secondary completion at T on parental completion 28 years earlier, on the active-expansion subsample (parental completion below 30%), with and without contemporaneous log GDP per capita. Both columns use the identical 672 country-years / 106 countries.

Table 4: Education predicts child lower secondary completion one generation later.

	Child lower secondary completion (t)
Parent edu ($t-28$)	1.434***
	(0.092)
Country FE	Yes
R^2 (within)	0.693
Observations	672
Countries	106

Notes: Country FE; country-clustered SE in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Putting log GDP into the same regression as a structural check barely moves the parental coefficient ($1.434 \to 1.320$, SE 0.074) while GDP attaches only a weakly significant 5.58 pp per log-unit (SE 3.19) and lifts within-R^2 only to 0.724 — GDP is downstream of education at societal scale, and the residualisation block below strips it out cleanly. *Produced by:* `scripts/tables/table_1_stepwise.py`.

Among countries where the expansion is still active, a 1-pp rise in parental lower-secondary completion raises child completion one generation later by **1.43 pp** ($t = 15.5$). Adding contemporaneous log GDP barely moves the parental coefficient ($1.434 \to 1.320$); GDP attaches only a weakly significant 5.58 pp per log-unit on the same 672 country-years. The stability of education's coefficient and the fragility of GDP's significance is the signature. The active-expansion screen removes the 15 USSR republics by construction, leaving the headline unchanged (Section 10.6).

The same comparison at three parental cutoffs on the GDP-merged panel (Table 5) returns education leading GDP at every level of the expansion; the parental coefficient is significant at $p < 0.01$ in every row.

Table 5: Education versus log GDP per capita as predictors of child lower secondary completion.

Cutoff	Edu β (SE)	Edu R^2	GDP β (SE)	GDP R^2	N	Countries
<10%	2.380*** (0.329)	0.613	16.518*** (4.164)	0.292	313	62
<50%	1.044*** (0.059)	0.670	17.514*** (3.861)	0.249	863	119
Full	0.707*** (0.058)	0.561	17.381*** (3.742)	0.265	945	144

Notes: Each row is a single-regressor within-country fit of child lower secondary completion at T on its predictor: parental education one generation earlier ($T{-}28$) for the Edu columns, contemporaneous log GDP per capita (T) for the GDP columns. Country fixed effects; country-clustered standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. "Cutoff" restricts the sample to country-years with parental lower secondary completion below the threshold; "Full" is the expansion sample (945 country-years, 144 countries, child-year completion in $[10\%, 90\%]$) — the broadest specification the paper's claim covers, not the GDP-merged subset. The cutoff rows are the active-expansion subsamples; the Full row is what the within-country fixed-effects estimator returns on the broadest sample, included so the reader can compare. Edu/GDP R^2 ratios: $2.1\times$ (<10%), $2.7\times$ (<50%), $2.1\times$ (Full); the ordering holds across the full 10%–90% sweep, confirming the amplification regularity. *Produced by:* `scripts/residualization/by_gdp_cutoff.py`.

Each generation gains more than it inherits. The generational amplification coefficient β_g is the within-country slope of child completion on parental completion at the one-generation lag (Glossary). $\beta_g > 1$ means a 1-pp rise in the parental cohort predicts a more-than-1-pp rise in the child cohort — the state's school niche is extending reach above what the home niche alone carries. $\beta_g \to 0$ at high baselines is ceiling compression (distinct from Lutz & Kebede's cross-level finding that education's effect on LE shows no diminishing returns from primary through tertiary). On the long-run 1900–2015 panel, $\beta_g = 3.10$ below 20% parental completion, 1.91 below 50%, 1.26 below 90%, and 1.05 unrestricted; within countries it traces the same compression curve from low to high baseline (Figure 4).

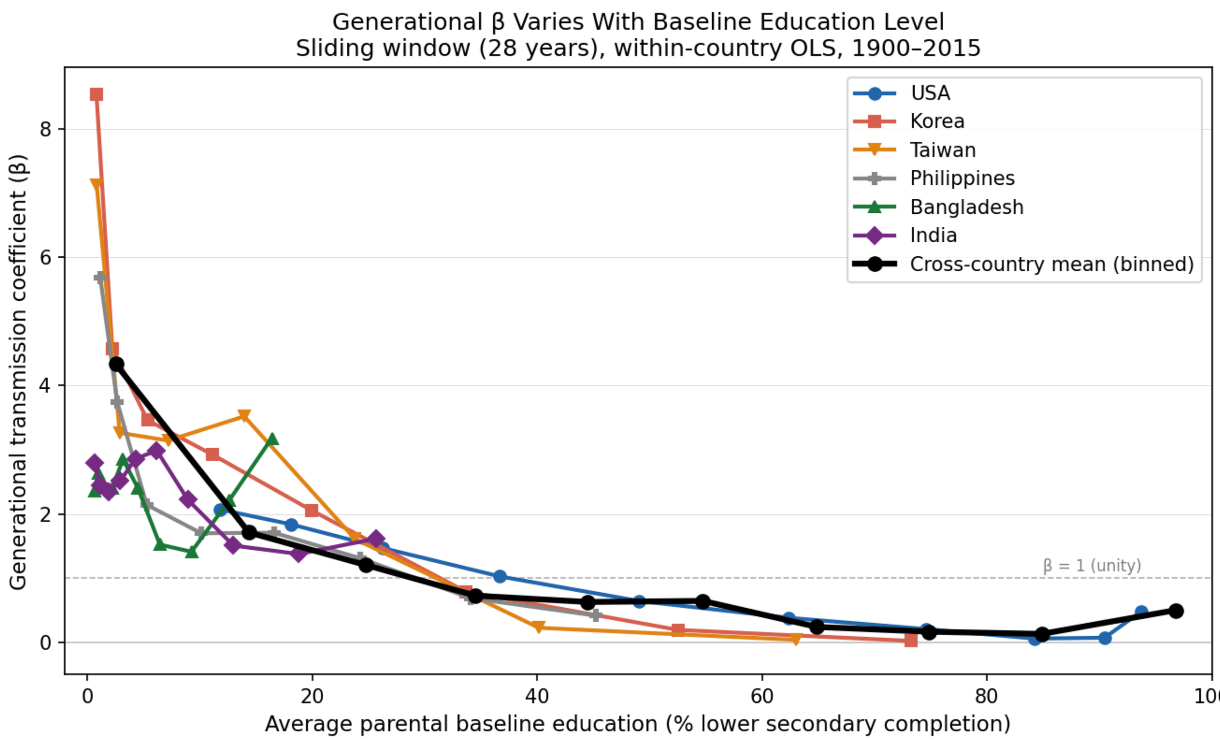

Figure 4: Generational β_g varies systematically with baseline education level. Each thin coloured line: 28-year sliding window within one country, against the window's parental baseline. Thick black line: cross-country mean across 185 countries ($n = 1{,}312$ country-windows), binned by parental baseline. Illustrative country trajectories (USA, Korea, Taiwan, Philippines, Bangladesh, India) are labelled; the universality claim sits with the black line. $\beta_g > 1$ at low baselines reflects state reach; $\beta_g \to 0$ reflects ceiling compression (WCDE v3 long-run cohort reconstruction 1875–2015; `scripts/wcde/long_run_generational.py`).

Quality vs. quantity in cross-section. The within-country panel (this section) shows lower-secondary completion capturing $R^2 = 0.67$ for TFR (lag 5), 0.43 for LE (lag 12), and 0.46 for U5MR (lag 12) — the load-bearing evidence is within countries as they expand schooling. A complementary 2015 cross-section ($n = 66$–67 countries with HLO test-score data) races lower-secondary completion against HLO cognitive quality on each outcome. In the cross-section, HLO dominates: for TFR, lower-sec $\beta_z = -0.19$ ($p = 0.14$), HLO $\beta_z = -0.51$ ($p < 0.01$); for U5MR and LE the HLO edge is larger still (Table 6). HLO is not absorbed by adding grandparent-generation lower-sec as a stock control: with lower-sec at T, lower-sec at $T-28$, and HLO jointly, HLO's coefficient holds at $\beta_z = -0.54$ ($p < 0.001$). Education quality, where measured, is genuinely informative across all three outcomes. The panel and the cross-section answer different questions — "what moves when a country expands schooling" versus "what differentiates countries today" — and lower-secondary coverage remains the operational policy lever in both.

Table 6: Quantity (lower-secondary completion) vs. quality (HLO) in 2015 cross-section.

Outcome (T=2015)	Quantity β_z (t)	Quality β_z (t)	Both-model R^2	n
TFR	-0.19 (-1.5)	-0.51 (-3.4)***	0.45	66
log(U5MR)	-0.16 (-1.7)	-0.70 (-7.6)***	0.78	67
LE	-0.12 (-1.1)	$+0.83$ ($+7.7$)***	0.62	67

Notes: Cross-sectional 2015 horse race. All variables standardised so β_z is the per-standard-deviation response. Quantity = lower-secondary completion at 2000 (the parent generation's own 20–24 cohort education, WCDE v3); Quality = harmonised learning outcome (HLO) test scores at secondary level from Angrist et al. (2021). Both specifications include log population and HC1 robust SEs. $^{***}p < 0.01$. The cross-section reflects between-country variance; the within-country panel result (lower-sec captures $R^2 = 0.67$ for TFR at lag 5) carries the load-bearing claim. HLO is not displaced by adding grandparent-generation lower-sec ($T-28$) as a cumulative-stock control — it captures something cross-country distinct from accumulated coverage. *Produced by:* `scripts/hanushek_horse_race_comprehensive.py` and `scripts/diagnostics/horse_race_with_earlier_cohort.py`.

The within-country panel and the 2015 cross-section make complementary claims. The panel says: as a country expands lower-secondary completion, all three outcomes move. The cross-section says: among countries today with HLO data, where you can measure cognitive quality, it is informative. Completion is the panel's operative variable because at population scale the question is whether the child sat inside the channel long enough for the categorical reorganisation to consolidate; the theory of why completion measures exposure rather than learning is in Chapter 4.2.

What rules out chance. The permutation null here defends the observed parent–child education slope; the residualisation block below runs the same null on education and residualised GDP together, where GDP's independent contribution falls inside it. A permutation null reshuffling the parent–child match across countries (200 iterations) places the observed slope more than 52 SDs above either null — the within-year null, which preserves global temporal trends, and the full shuffle, which breaks every systematic link (`scripts/robustness/permutation_null.py`). Common trends, panel autocorrelation, and serial correlation cannot generate the coefficient by chance. Within countries male and female lower-secondary completion expand together (within-country correlation > 0.9); the home niche moves as a unit, and the country panel cannot decompose by sex. The cross-sectional microdata literature locating the channel in mothers specifically (Caldwell 1979; Schultz 2002; Jensen 2010; Duflo 2012) is identified at the within-country individual level and is consistent with the panel result: the channel is mother-weighted in routing and household-located in measurement.

Non-depreciation of the channel under destructive disruption is the Cambodia test (Section 7.2).

GDP residualised: the null. Take each country's log GDP per capita and residualise out the part that education produced (Frisch-Waugh-Lovell; `scripts/residualization/education_vs_gdp.py`); use only what remains to predict development outcomes. The residual predicts nothing (Table 7).

Table 7: Residualised GDP analysis – education vs. log GDP per capita as predictors.

Outcome (T+lag)	Edu R^2	GDP R^2	Resid GDP β (SE)	Resid p	Resid R^2
Panel A: max-sample (per-outcome, ceiling ≤90%)					
Life expectancy	0.428	0.192	0.160 (1.418)	0.91	**0.000**
Total fertility	0.669	0.292	0.004 (0.324)	0.99	**0.000**
Child education	0.524	0.295	2.845 (3.472)	0.41	**0.005**
Under-5 mortality	0.457	0.170	2.933 (9.494)	0.76	**0.001**
Panel B: common sample ($n = 645$, 123 countries)					
Life expectancy	0.424	0.193	0.279 (1.428)	0.84	**0.000**
Total fertility	0.676	0.291	−0.005 (0.326)	0.99	**0.000**
Child education	0.562	0.297	2.899 (3.509)	0.38	**0.005**

Outcome (T+lag)	Edu R^2	GDP R^2	Resid GDP β (SE)	Resid p	Resid R^2
Under-5 mortality	0.434	0.170	2.933 (9.494)	0.76	**0.001**

Notes: Entry-cohort design (entry ≥10%, ceiling ≤90%); country FE; lower-secondary completion at T, outcome at T+lag where lag is biologically anchored per outcome (TFR: lag = LAG_TFR, reproductive peak; LE and U5MR: lag = childrearing window, LE being mortality-dominated; child-edu: lag = LAG_GENERATION, the cross-generation step); $T = 1960$–1990; country-clustered SE in parentheses. Residualised GDP is log GDP per capita after Frisch-Waugh-Lovell residualisation against education. Panel A reports each outcome on its own max-sample (sample sizes vary because LE/TFR/U5MR/child-edu have different country-year coverage). Panel B restricts to the intersection across all four outcomes ($n = 645$, 123 countries), so rows are directly comparable; the residual-GDP null survives the restriction. Unique-R^2, Maddison-backfilled, and ceiling-sweep variants in `scripts/residualization/`; all leave the headline intact. *Produced by:* `scripts/residualization/education_vs_gdp.py`, `scripts/tables/regression_tables.py`.

A 1-pp rise in lower-secondary completion at T predicts, at each outcome's biological lag, a gain of $+0.19$ years of life, -0.058 births per woman, and 1.33 fewer under-five deaths per thousand; the same-magnitude response to residualised log GDP is statistically indistinguishable from zero in every row. Education explains 0.43 to 0.67 of within-country variation; the leftover-GDP R^2 is at most 0.019, and it stays there across the range of biologically plausible lags, from the childrearing window out to a full generation (`scripts/robustness/lag_sensitivity.py`). Education is doing essentially all the predictive work. A 1,000-replication country-resampling bootstrap on the residualised-GDP coefficient ($\beta_{\mathrm{GDP}^{\perp}}$) confirms the null directly: the 95% percentile CI includes zero for log LE, log TFR, log U5MR, and child education at the headline ceiling (`scripts/robustness/robustness_tests.py`). A single permutation null on this panel makes the asymmetry direct: across 200 full-shuffle iterations the education coefficient sits 14 to 18 standard deviations above its null on every outcome — and 11 to 13 above it after I strip GDP out of education — while residualised GDP, education stripped out of it, sits inside its null, within two standard deviations, on all four (`scripts/robustness/permutation_null_symmetric.py`).

Why controlling for GDP is the wrong move. The prior is Prediction 4: GDP is the national-accounts face of the Level-2 (educated-CT) regime, so residualised log GDP per capita should have no independent predictive power once education's contribution is stripped. The biology gives the direction; the panel tests it. GDP is downstream of education at every scale that matters. The educated population produces the GDP. Controlling for GDP at the country panel level controls for education running through national accounts. The empirical claim is that once you stop doing that — strip GDP of education's contribution and use only the leftover — the leftover predicts nothing. The Asian Financial Crisis (Section 7.3) shows the same separation in observable time: GDP collapsed across five countries (Indonesia -14.5%, Thailand -8.8%, Malaysia -9.6%, Philippines -3.0%, Korea -6%) while lower-secondary completion continued uninterrupted in every one — Indonesia gained 5.4 pp through the crash. The shell can be destroyed; the channel survives. The reverse-direction reading sharpens it from the panel side: on the expansion sample ($n = 537$), parental-generation GDP alone predicts child education substantially ($\beta = 19.5$, $R^2 = 0.335$), but adding parental education collapses GDP's coefficient by 69% to $\beta = 6.1$ and adds only $R^2 = 0.021$ beyond education. Education's unique contribution beyond GDP is $R^2 = 0.280$ — 13.3× larger. GDP transmits through education, not vice versa.

Robustness to baseline stratification. The pooled level-U5MR residGDP × Post-2000 interaction is $+34.0$ ($p = 0.022$, $n = 867$; `scripts/robustness/u5mr_residual_by_year.py`); the pre-2000 log-U5MR within-country residGDP slope is -0.37 ($p = 0.018$, $n = 203$; `scripts/robustness/residgdp_by_baseline_fine.py`). The observational test is whether residGDP predicts mortality at a *given* parental-education baseline. Stratifying by parental lower-secondary completion at 5 pp resolution on the Maddison-backfilled panel, no individual bin shows a residGDP coefficient that clears bootstrap significance. The inverse-variance-weighted within-meta β on level U5MR post-2000 is $+5.28$, with a 2000-rep country-cluster bootstrap 95% CI of $[-104, +48]$ — indistinguishable from zero (`scripts/robustness/residgdp_strengthened_bin_test.py`). The pooled signal is between-bin composition: pre-2000 active-expansion cells sit at median parental completion of 28.6%; post-2000 cells at 38.2%, on a U5MR distribution that fell sharply over the intervening period. Within bins, given a country's parental education, residual GDP does not predict mortality — and the same null holds at each of the three timescales the chapter defines (lag 0 contemporaneous, lag 12 childrearing window, lag 28 cross-generation; bootstrap 95% CIs $[-77, +53]$, $[-69, +60]$, $[-83, +44]$ respectively; `scripts/robustness/residgdp_strengthened_bin_test_lags.py`).

The Lutz and Kebede (2018) 1990–2010 child-mortality acceleration is real. Vaccines, oral rehydration, vitamin A, IMCI, GAVI, the MDGs delivered absolute reductions in deaths per thousand. The acceleration runs through the schooled mothers of the 1970s and 1980s expansions reaching childbearing age and using those tools.

9.6 The Shape of the Response

Each outcome responds to lower-secondary completion at the lag where its mechanism runs. Log GDP is contemporaneous — the educated worker contributes to output in the same year. A woman's own fertility (TFR) follows about five years on, as her cohort reaches its reproductive peak; her children's survival (U5MR) registers at the childrearing window, about twelve years on. Life expectancy at birth moves on that same childrearing clock — lag 12 — because it is mechanically dominated by child mortality, not by a separate adult-longevity channel a generation back. None of these outcomes carries a clock of its own. They do not drift upward with calendar time; they move when schooling moves, and they have risen across the post-war decades only because schooling spread across those decades. Time is not the cause — education is — and the four outcomes co-move because the same lever drives each of them. The panel evidence below is the within-country signature: as a country expands lower-secondary, each outcome moves at its own lag.

Table 8: Lower-secondary completion and development outcomes at canonical lag per outcome.

Panel A: Lower-secondary completion at $T \to$ each outcome at its canonical lag (per `scripts/_shared.py`): GDP at lag 0; LE at lag 12 (childrearing window); TFR at lag 5 (cohort at reproductive peak); U5MR at lag 12. Each row is one outcome; the coefficient is education-only. Adding contemporaneous log GDP as a control barely changes any coefficient (Notes).

Outcome (lag)	Coefficient
log GDP (lag 0)	+0.0185***
	(0.0014)
log(LE) (lag 12)	+0.0042***
	(0.0003)

Outcome (lag)	Coefficient
log(TFR) (lag 5)	−0.0165***
	(0.0005)
log(U5MR) (lag 12)	−0.0327***
	(0.0014)
Country FE	Yes
N (GDP row)	1,812
N (LE row)	1,976
N (TFR row)	2,155
N (U5MR row)	1,575

Notes: Each outcome regressed on lower-secondary completion at T at its canonical lag per `scripts/_shared.py`: log GDP at lag 0 (`LAG_CONTEMPORANEOUS`); LE at lag 12 (`LAG_LE`, childrearing window); TFR at lag 5 (`LAG_TFR`, cohort at reproductive peak); U5MR at lag 12 (`LAG_CHILDREARING`). Country FE throughout; country-clustered SE in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Coefficients are semi-elasticities (proportional change in outcome per pp rise in lower-secondary completion). Adding contemporaneous log GDP at T as a control barely moves any coefficient: LE $+0.0039$ (0.0004), TFR -0.0153 (0.0008), U5MR -0.0290 (0.0021), on the GDP-merged samples (N = 1,633 / 1,812 / 1,253 respectively); the log GDP row is unchanged because the outcome is itself log GDP. The cross-generation (T+28) variant is retained as a forward-prediction identification robustness; see Section 9.5 for the reverse-direction check (GDP at T predicting child education at T+28). *Produced by:* `scripts/wcde/table_contemporaneous.py`.

Reading Panel A: each 1-pp gain in lower-secondary completion at T predicts, each outcome at its biological lag (Notes), 1.85% higher GDP per capita (lag 0), 0.42% higher life expectancy (≈0.29 years at LE $= 70$; lag 12), 1.65% lower TFR (≈0.06 fewer children per woman at TFR $= 3.5$; lag 5), and 3.27% lower under-five mortality (lag 12). All four are within-country effects; every Panel A coefficient is significant at $p < 0.01$. Across the four outcomes the predictor's coefficient barely moves once log GDP enters as a control (Notes) — the population-scale signature that GDP is downstream of education, not a parallel channel.

Four generations of reach. Figure 5 extends the lag-28 finding to four generations. The standardised within-country $|\beta|$ for lower-secondary completion predicting each outcome decays smoothly from lag 28 through lag 112; every outcome remains significant at $p < 0.001$ at four-generation depth (minimum $|t| = 5.9$, TFR). Under-5 mortality carries the deepest signature ($|\beta| = 0.331$ at lag 112, $|t| = 24.4$); LE, TFR and child education show the same smooth monotonic decay. No other input has comparable reach: log GDP per capita's R^2 falls below education's by lag 20–25 and cannot be tested past lag ∼45 on raw WDI data, by construction — the pre-1960 subsistence baseline (∼$400–600 constant 2015 USD) has no cross-country variation to test before literate populations existed. The construction-bounded backfill check is in `scripts/robustness/backfill_all_outcomes.py`.

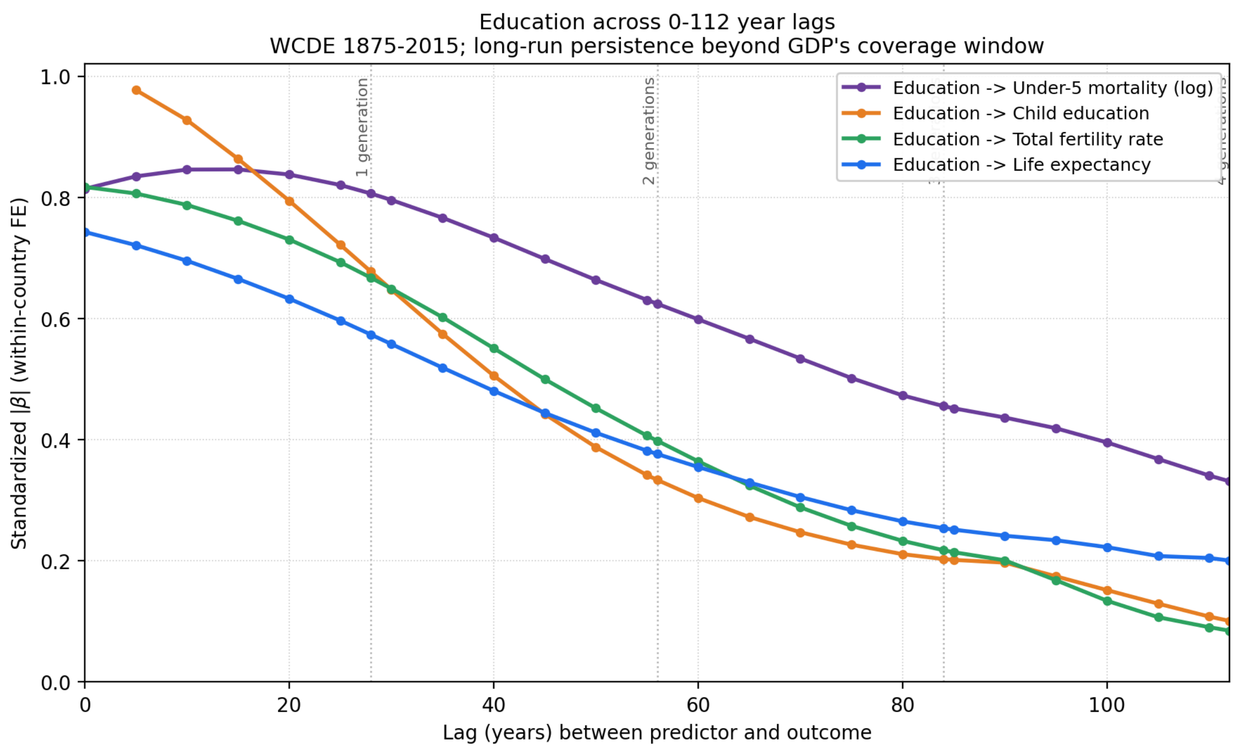

Figure 5: Education across lags 0–112 years. Standardised within-country $|\beta|$ for lower secondary completion predicting four development outcomes, extended to four-generation depth (WCDE v3 cohort reconstruction 1875–2015). Generational horizons 1–4 marked.

9.7 The Grandparent Channel

Does the grandparent's education predict outcomes *independently* of the parent's? Under biological anchoring — parent education at $T-5$ (the parental cohort at reproductive peak in the outcome year) and grandparent education at $T-30$ — both channels carry independent, highly significant signals on fertility everywhere in the panel. In the full panel $\beta_p = -0.050$ and $\beta_{gp} = -0.011$ (both $p < 0.01$; $n = 1216$, 152 countries); within-R^2 rises from 0.56 to 0.59 once grandparent education is added. Restrict to countries where parental completion is below 50% and both coefficients grow: $\beta_p = -0.068$, $\beta_{gp} = -0.048$ (both $p < 0.001$; $n = 464$, 80 countries); within-R^2 jumps from 0.69 to 0.74. The parent's coefficient is everywhere larger; the grandparent's coefficient never collapses to zero. The largest grandparent signal is on child survival (U5MR, at the childrearing lag — parent at $T{-}12$, grandparent at $T{-}37$): in the full panel $\beta_{gp} = -0.018$ on log U5MR, as large in absolute terms as the parent's ($\beta_p = -0.018$); within-R^2 rises from 0.407 to 0.653 once grandparent education is added ($n = 825$, 165 countries). Table 9 collects the cited coefficients.

Table 9: Grandparent channel — parent-only vs. parent + grandparent education.

Outcome (sample)	Parent-only β_p	+ GP: β_p	+ GP: β_{gp}	within-R^2	n (ctry)
TFR (level, parent <50%)	-0.096***	-0.068***	-0.048***	0.69→0.74	464 (80)
log U5MR (full panel)	-0.033***	-0.018***	-0.018***	0.45→0.65	825 (165)

Notes: Grandparent (GP) education predicts outcome via the biologically appropriate lag per outcome: for TFR, parent at $T - 5$ (parental cohort at reproductive peak in the outcome year) and grandparent at $T - 30$; for U5MR, parent at $T - 12$ (childrearing window) and grandparent at $T - 37$. Country FE; country-clustered SE. The TFR row uses lower-secondary completion in levels on the low-baseline subsample (parent completion $< 50\%$); the U5MR row uses log U5MR on the full panel. Parent and grandparent education are drawn from the same within-country series and are structurally correlated: in a country undergoing sustained expansion both regressors trend together, so the marginal grandparent signal collapses once parental baselines saturate. The low-baseline restriction in the TFR row isolates the regime where they vary enough to separate. $^{**}p < 0.05$, $^{***}p < 0.01$. *Produced by:* `scripts/robustness/grandparent_effect.py`, `scripts/robustness/grandparent_effect_all_outcomes.py`.

The pattern is what the generational gate predicts: at low baselines each educated person in the household adds independent weight; at high baselines the parent has internalised everything the grandparent could offer. This generates a falsifiable prediction the panel cannot test directly: any educated elder inside the child's kin/community radius (co-resident grandparent, aunt, village exemplar) should raise grandchild outcomes beyond what the parent's education predicts; equivalent elders present only on paper (deceased, non-resident) should not. DHS co-residence data can test it.

9.8 Universality Across Subsamples

The mechanism is universal — and universal in a stronger sense than a stable coefficient. A relationship that stays positive in every region shows it does not vanish anywhere; it does not, by itself, show the same mechanism runs everywhere, since unlike pathways can each yield a positive slope. The test is whether the *whole signature* repeats: not one coefficient but the coordinated response of all four outcomes, each at its own biological lag. It does. Two tests, panel and residual.

Panel, eleven subsamples. Re-estimating Column 1 of Table 4 on eleven subsamples (six regions, two child-cohort eras, three within-sample GDP terciles) produces eleven positive coefficients, every one significant at $p < 0.01$, with β ranging from 1.258 (Middle East & N. Africa, n=67) to 2.562 (South Asia, n=58). Sub-Saharan Africa, where most of the remaining 20% of humanity sits, returns $\beta = 1.270$ (SE 0.144, $R^2 = 0.692$, n=312, 40 countries) — tracking the full-sample headline. Re-estimating not that single coefficient but the full per-outcome specification (Table 8: lower-secondary completion at $T \rightarrow$ log GDP at lag 0, log TFR at lag 5, log LE and log U5MR at lag 12) on the same eleven subsamples reproduces the entire signature: in every subsample GDP rises, fertility falls, life expectancy rises, and child mortality falls, each at its canonical lag — all 44 coefficients (eleven subsamples $\times$ four outcomes) carrying the predicted sign and significant at $p < 0.01$. A different mechanism in each region could not produce one coordinated four-outcome pattern at the same four lags everywhere. Era and tercile splits rule out post-Cold-War and income-band artefacts. The universality claim is also tested on the institutional literature's own preferred sample: on AJR's (2001) 64-country former-colony base, colonial-era education explains a near-identical share of log GDP variance as AJR's *avexpr* institutional measure (Chapter 11, §11.9), and the settler-mortality instrument identifies the educational channel as well as the institutional one.

Residuals, eight over-performers. Some countries are educating children far beyond what their parents' generation would predict. Country fixed-effects residuals from a within-country regression of child on parent lower-secondary completion at the 28-year lag, using the full child-education panel (1875–2015, 185 countries), identify eight at $T{+}28 = 2015$: they are uniformly poor.

Table 10: Education over-performers (2015 FE residuals).

Country	FE residual above baseline	GDP per capita (2015)
Maldives	+37.2 pp	$9,645

Country	FE residual above baseline	GDP per capita (2015)
Bhutan	+25.5 pp	$2,954
Cape Verde	+24.9 pp	$3,415
Viet Nam	+22.6 pp	$2,578
Tunisia	+20.8 pp	$4,015
Nepal	+17.9 pp	$876
Bangladesh	+17.5 pp	$1,224
India	+14.9 pp	$1,584

Notes: Each within-country residual is child lower-secondary completion at $T+28$ minus the within-country slope $\hat{\beta}$ multiplied by parent completion at T, country-demeaned; $\hat{\beta}$ is estimated by country fixed-effects OLS across the *full* child-education panel (1875–2015, all 185 countries), not the active-expansion sample of Table 4. The figure reported is each country's residual at $T = 1990$ — child cohort observed in 2015. The residual measures deviation from the country's own historical trend predicted by the within-country slope. It is not a comparison against other countries at similar income, region, or other observables; a country can be a positive residual here simply because its child generation rose faster relative to its own parents than the panel slope predicts. *Produced by:* `scripts/tables/regression_tables.py`.

Nepal at $876 (2015), Bangladesh at $1,224, Vietnam at $2,578. Income is not the prerequisite for educational over-performance. The robustness battery (alternative cutoffs, lag lengths, PPML, balanced subpanel, Barro-Lee replication, the 20-test econometric battery) is in `scripts/` (Section 9.10). None reverses the ordering, and no check reduces the education/GDP R^2 ratio below 2.1$\times$.

9.9 Every Method Agrees

The within-country panel above used ordinary least squares with country fixed effects. A reader trained in econometrics will ask how much of the education-load-bearing result depends on that one functional form. To answer without leaning on a single estimator, I ran the same prediction problem — features at T predicting outcomes at the canonical lag — through a spread of model classes under country-clustered five-fold cross-validation: countries, not observations, are split across folds, so every scored country is fully unseen in training. Each method's education contribution is the drop in out-of-fold R^2 when the education feature block is zeroed at inference.

The within-fixed-effects estimator itself cannot be cross-validated this way — a held-out country's fixed effect is unidentified, so the linear class is represented by its regularised members, ridge and lasso. To those I add two tree ensembles and a universal transformer encoder trained with no country identity in its inputs, so it cannot memorise a single country.

Table 11: Education's out-of-fold R^2 drop when the education block is removed, country-clustered five-fold, each outcome predicted at its biological lag (LE@12, TFR@5, U5MR@12) on its own single-target panel.

Method	ΔR^2 (LE)	ΔR^2 (TFR)	ΔR^2 (U5MR)
Ridge (CV-tuned)	0.193	0.253	0.240
Lasso (CV-tuned)	0.256	0.355	0.275
Random forest	0.222	0.352	0.401
Gradient boosting	0.193	0.329	0.361
Universal transformer	0.234	0.262	0.394

Notes: Out-of-fold R^2 drop from zeroing the education block (WCDE current and cohort-historical attainment, Barro–Lee shares, derived gender-gap terms). Each outcome is predicted at its own biological lag (LE@12 via `LAG_LE`, TFR@5, U5MR@12) on a single-target panel. Transformer values are medians across 30 seeds; the linear and tree methods are single CV-tuned fits. Country-clustered folds throughout. *Produced by:* `scripts/ml/chapter9/spec_curve.py --parent` and `aggregate_parent_vantage.py`.

The five rows return the same answer. Education's out-of-fold contribution stays in the same band across every estimator, and the transformer — which imposes no functional form and cannot memorise a country — sits inside that band rather than above or below it. The result is not an artifact of the panel's estimator.

Education against income, head to head. The cleanest single cut pits education against its main rival. I trained the transformer on the full feature set, then ablated one block at a time and measured the loss in held-out accuracy. Removing education costs the model between a third and over half of its held-out accuracy — 0.31 for life expectancy, 0.31 for fertility, 0.52 for under-five mortality. Removing GDP costs almost nothing — 0.02, 0.002, and 0.007 for the same three. On unseen countries, with no functional form imposed, education does an order of magnitude more work than income.

A causal-style estimate. A double machine-learning estimator — education partialled against every other feature, cross-fitted — puts the per-level education effect at $+4.85$ years of life expectancy (95% CI $[4.63, 5.56]$), -0.94 children ($[-0.98, -0.89]$), and -36.5 deaths per thousand on under-five mortality ($[-39.1, -33.3]$). Every interval excludes zero.

Train on the past, predict the future. To rule out era-specific artifacts I trained the transformer only on schooling cohorts completed before a cutoff date and evaluated it on the era after. Across the four cutoffs and three outcomes, held-out R^2 never falls below 0.70 (under-five mortality, earliest cutoff) and reaches 0.93 by the 1990 cutoff; trained on pre-1980 cohorts alone, the model predicts post-1990 life expectancy at $R^2 = 0.84$. The trajectory is forward-projectable.

Falsification. Six placebo transformations break the education–outcome link in different ways. The three that should destroy it — shuffling outcomes across rows, replacing education with random noise, and scrambling country labels — collapse the education contribution to zero (median R^2 changes of -0.02, -0.005, and -0.06). Replacing education with absolute latitude leaves almost nothing (0.02). The two transformations that leave a real channel standing — within-year GDP rank and a year shuffle — retain a fraction (0.07 and 0.12), as they must, because income and time are themselves education-correlated. None approaches what the true features carry. (Full table: `scripts/ml/checkin/chapter9_placebos.json`.)

What the gap bought. A leader will ask what a country's schooling commitment actually delivered. For ten matched pairs I take country A's 1990 feature window, replace its education trajectory with country B's, and read the change in each outcome's prediction at its biological horizon (life expectancy and under-five mortality at $T{+}12$, fertility at $T{+}5$) — each country scored by the fold in which it was held out, so it is never read by a model that trained on it. The result is in Table 12.

Table 12: What each country's own education trajectory bought relative to its peer's (universal transformer, out-of-fold; $T = 1990$ features, each outcome read at its biological horizon: LE and U5MR at $T{+}12$, TFR at $T{+}5$).

Country A	vs B's education	ΔLE (yrs)	ΔTFR	ΔU5MR
Sri Lanka	Pakistan	+12.10	−2.19	−48.2
China	India	+7.06	−2.06	−16.9
Botswana	Zimbabwe	+4.80	+0.23	−11.8
Cuba	Dominican Republic	+4.61	−0.70	−8.4
South Korea	Philippines	+3.04	−0.30	−3.6
Vietnam	Myanmar	+1.84	−1.10	−11.8
Rwanda	Burundi	+0.47	+0.00	−14.0
Bangladesh	Pakistan	−0.76	+0.00	+4.0
Kenya	Tanzania	−1.39	−0.28	−4.8

Notes: For each pair the transformer is given country A's $T = 1990$ feature window with the education block replaced by country B's; the reported Δ is the change in that outcome's prediction at its biological horizon (LE and U5MR at $T{+}12$, TFR at $T{+}5$). Taiwan (no World Bank panel coverage) cannot be computed and is omitted. The effect scales with the size of the schooling gap: large where the two diverged hard (Sri Lanka, China, Cuba, Botswana), small where they did not (Kenya, Bangladesh). Bangladesh shows no gain because over the 1975–1990 window Pakistan was still *ahead* on lower-secondary completion — Bangladesh's documented overtaking is a post-1990, female-led catch-up outside this swap's horizon, and the model correctly declines to assign it backward. *Produced by:* `scripts/ml/chapter9/counterfactuals.py`.

What these methods establish. Country-clustered cross-validation, walk-forward projection, ablation, and counterfactual feature swaps establish that the panel's reading is not a methodological artifact — it survives every guard against overfit, country memorisation, time-period dependence, and functional form. The biology and the natural experiments establish the channel; this battery confirms that its population-scale signature is real. Leave-one-out fits and the full placebo and counterfactual detail are in `scripts/ECONOMETRICS.md`.

9.10 What the Panel Supports

The panel sharpens the population-scale pattern — the convergence in aggregate, the generational lag, the irrelevance of residualised log GDP per capita, the universality across subsamples, and the invariance of all of it to the choice of estimator (Section 9.9). The biology is fundamental, and the natural experiments, the country histories, and the USSR falsification carry the identification; the panel is the evidence that the same channel's signature holds across the whole species at once. The chapter reads those population-scale signatures rather than estimating a treatment effect — signatures of a structural channel that Chapters 2–5 deduce and Chapters 7–8 identify in observable time. The natural experiments in Chapter 7, the country histories in Chapter 8, and the USSR falsification in Chapter 10 are the rigorous work. Section 5.5 gives the structural reason no other variable could have carried the panel signature: the cognitive substrate is loadable only during the juvenile dependency window, and education is the only intervention that loads it; institutions, markets, and technology act on adults already formed. The standard toolkit — two-way fixed effects, instrumental variables, the marginal-control battery — is built for a different question; Section 11.8 sets out why it runs downstream of this channel rather than reaching it. The reader who wants to push on the panel will find every specification, robustness check, and variant in `scripts/` on GitHub, indexed by `scripts/ECONOMETRICS.md`: common-sample, Maddison-backfill, and ceiling-sweep variants of the residualisation; the country-shuffle permutation null; and the 20-test panel-diagnostic battery — including the staggered-adoption DiD, IV contests, and TWFE diagnostics run for completeness.

10. Hollow Education: What the Soviet Anomaly Tests

The mechanism predicts: if a population is reported to be educated but does not actually undergo categorical reorganization in the childhood window, then fertility and mortality will not shift to match the reported education. The Soviet 15 provide the natural experiment.

The paper's empirical engine is lower-secondary completion — the floor at which the cultural-transmission regime begins acting through household decisions (Section 5.6). I take the series from the WCDE v3 reconstruction (Lutz et al. 2021) because it reaches back to 1875 on a consistent definition and passes the phenotype tests laid out in Section 10.3. On the 185-country panel, WCDE reports near-universal lower-secondary completion for the fifteen Soviet republics by 1970.

But that near-universal reporting masks two different realities. The phenotype test reveals two patterns, not one. The six republics west of Moscow (Belarus, Ukraine, Lithuania, Latvia, Estonia, Moldova) sit within $\pm 1\sigma$ of the global education-phenotype fit on log U5MR — their reported completion is roughly real. The eight republics east or south of Moscow (Georgia, Armenia, Azerbaijan, Turkmenistan, Uzbekistan, Tajikistan, Kyrgyzstan, Kazakhstan) sit far above what their reported lower-secondary completion predicts — 2.6 to 4.0σ for the seven with multi-year under-five-mortality data (Uzbekistan's series begins only in 2020; Table 13). The split is categorical, not gradient: Belarus at 629 km west of Moscow shows a U5MR residual of $+0.43\sigma$; Georgia at 451 km east shows $+2.58\sigma$. Crossing the Moscow meridian moves a republic from "consistent with reported education" to "phenotype catastrophically inconsistent with reported education" in a single jump.

The chapter treats this as a diagnostic, not a footnote. It is a natural experiment run in reverse: the input side is corrupted, and the corruption pattern reveals the social geography of the corrupting institution. The Soviet metropole reported European-core educational standards uniformly across its empire. Where the underlying population was European-core — Slavic west, Lutheran-Catholic Baltics — the report was approximately true. Where the underlying population was Muslim or Caucasian periphery, the report was metropolitan fiction. And a state that inflates its schooling deflates its mortality by the same hand: in rural Central Asia Soviet child mortality was reported three to four times too low, and old-age survival inflated by age exaggeration (Anderson & Silver 1986). The under-five residuals that locate the anomaly (Section 10.4) run on those understated deaths — so they understate the true gap, not invent it. But the incorruptible witness is not child mortality, which the regime managed. It is fertility: a pro-natalist state that awarded medals for large families had no reason to fake its birth rates down. Fertility alone reports truthfully on the schooling, and it tracks the periphery's low-education neighbours, not the credential.

In the childhood frame, hollow education is exactly this: the eighteen-year window did not get loaded with the cognitive architecture the schooling was supposed to install — whether because the schooling was thin, the attendance partial, or the credential itself fictional over a population the state could not see. The credential measures only what was reported; the phenotype measures what entered the window.

The diagnostic yields three payoffs. First, a credibility case for the panel: the active-expansion screen and the hollowness flag independently exclude the same fifteen republics, so the headline results are robust to USSR exclusion by construction. Second, a reconciliation with the Hanushek–Woessmann "knowledge capital" line (2008, 2012): the HLO test-score measure, which the field has read as an alternative to completion-based education, turns out to be a multi-generation integrated output of completion itself. Third, a sharpened thesis on what schooling transmits. The cognitive architecture loaded in the dependency window is what produces the development response across all three outcomes (TFR, LE, U5MR); the credential without the phenotype does not. Hollow education is not education.

10.1 The Anomaly

Of the 154 countries that have crossed the developmental threshold by 2022, 24 were post-socialist — 13 USSR republics plus 11 Warsaw Pact and Yugoslav successors. The post-socialist crossers follow a pattern nothing else in the panel matches. Among market-economy crossers, the median country reached 90% lower-secondary completion 16 years *before* crossing the developmental threshold. Among post-socialist crossers, the median country reached 90% at the crossing year. Among the USSR core, Turkmenistan crossed 57 years after the reported date at which it reached 90%. Either these republics had reached 90% decades before their fertility and longevity trajectories responded — a 50-year mechanism gap with no precedent — or the 90% was not what it claimed to be.

The rest of the chapter tests the second hypothesis. Under the paper's mechanism, reported education at year T should produce a phenotypic response within a generation, each outcome at its own biological lag. Where it does not, either the mechanism is wrong or the reported education is wrong. The mechanism passes across the panel outside one bloc: 171 of the 185 countries behave as it predicts, and the 14 Soviet republics with a phenotype to test (Uzbekistan drops out of the panel on data coverage) share a single origin — one statistical office. Within that bloc the test does not break uniformly; it splits the republics at Moscow's longitude, as the chapter opening previewed. The European-core republics to the west show the phenotypic response their reported education predicts, while the republics to the east and south do not

(Section 10.4). The failure is the periphery the metropole could not see, not the credential itself.

10.2 What WCDE Reports vs What Was Observable

Three ground-truth checks against contemporaneous non-Soviet neighbors — Iran, Turkey, Afghanistan, Pakistan — with comparable pre-1917 educational baselines:

Gender parity in 1970. Central Asia reports a mean female-minus-male lower-secondary completion gap of -2.7 pp (near parity). Iran, Turkey, Afghanistan, and Pakistan report a mean gap of -15.6 pp, with several countries at -11 to -18 pp. Near-parity in 1970 rural Central Asia, with no comparable pre-Soviet mass-schooling infrastructure, diverges sharply from what comparable-baseline societies were producing the same decade. The Baltics' reported near-parity is plausible (pre-1940 European educational infrastructure); the Central Asian figure is not.

Primary minus lower-secondary. By construction, primary completion must exceed lower-secondary completion: one cannot complete the latter without the former. In 1980, non-Soviet neighbors show a mean primary-minus-lower-secondary gap of $+21.0$ pp — a fifth of the cohort completes primary without continuing to lower-secondary. Central Asia shows $+1.6$ pp: virtually every child who completes primary is reported as completing lower-secondary as well. This holds across individual Central Asian republics. Near-complete continuation is not what normal school systems produce; it is what uniform reporting produces.

Raw 1970 levels. Kazakhstan 94%, Turkmenistan 95%, against Iran 22%, Turkey 22%, Pakistan 16%, Afghanistan 6%. A 70-percentage-point gap from comparable pre-1917 baselines, built in half a century, is not what rural agricultural populations produced in this period anywhere else on the planet. Later cognitive evidence is consistent: Kyrgyzstan's 2009 PISA score of 350 sits at the low end of the global distribution; Turkmenistan, Uzbekistan, and Tajikistan do not participate in international assessments; Baltic PISA scores at 500–535 are consistent with their reported education and with separate pre-1940 educational infrastructure.

10.3 The Phenotype Test

The mechanism predicts that educational input at T produces phenotypic output within a generation, each outcome at its own biological lag. Where the predicted output does not appear, the input was not real. Figure 6 shows the test for under-5 mortality.

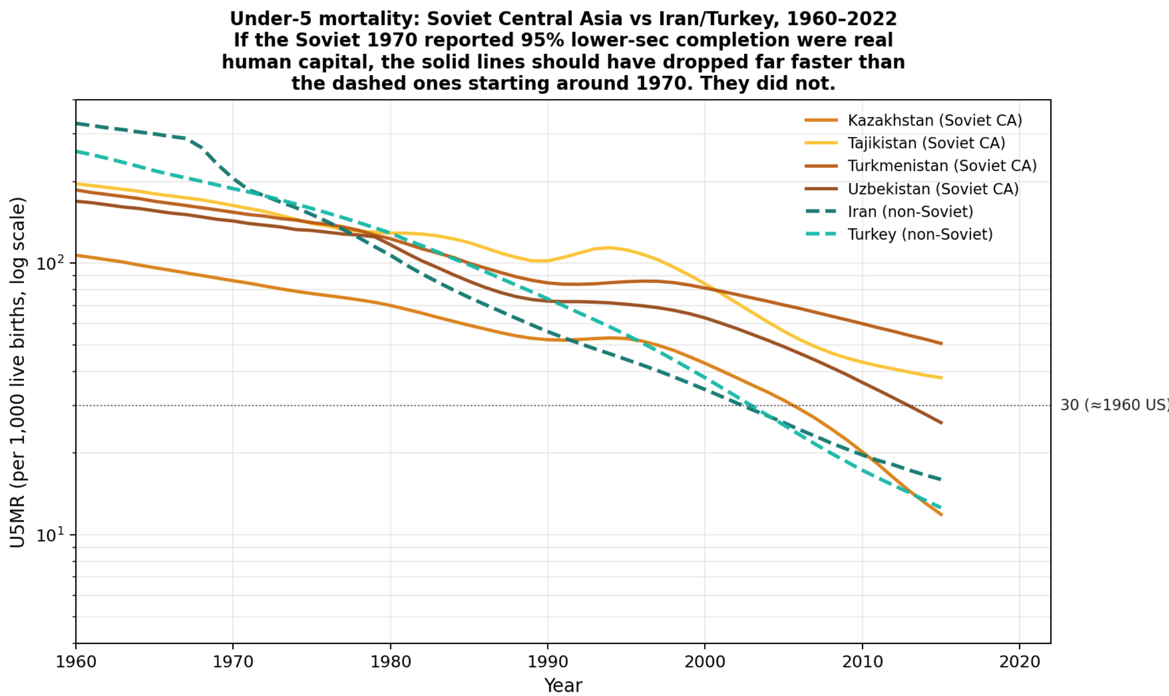

Figure 6: Under-5 mortality trajectories, 1960–2010. Central Asian and Slavic USSR republics (colored lines) start 1960 with large U5MR head starts over Iran (327 per 1,000) and Turkey (258). By 2010, Iran (20) has overtaken every Central Asian country except Kazakhstan. Under the mechanism, the country with four times the reported education in 1970 should dominate its neighbor permanently; no later catch-up should be able to reverse a head start built on five-times-larger secondary completion. The crossover says the head start was not real. *(Source: World Bank WDI.)*

Iran 1960 → 2010: 327 → 20 per 1,000 (94% decline). Kazakhstan 1960 → 2010: 107 → 20 (81% decline). Turkey 2010 at 17 beats every Central Asian country, Kazakhstan included. If the 1970 Soviet 95% had been real, Central Asia should have dominated Iran and Turkey permanently. But the head start was partly a reporting artifact: the Soviet live-birth definition excluded the most fragile newborns, understating infant mortality by 22–25% empire-wide and by a factor of three to four in rural Central Asia, where deaths went unregistered or were misattributed to the second year of life (Anderson & Silver

1986). The reported 107 was a floor; the true figure was far higher. What arrived after 1991 was not a Central Asian collapse but honest counting — Iran did not overtake the periphery so much as the periphery's numbers stopped being made. Iran and Turkey are running on school systems that delivered actual cognitive skill to the mothers whose children survived.

TFR tells the same story more quietly. Central Asian fertility trajectories from 1960 onward are indistinguishable from Iran's and Turkey's. There is no 1991 kink — the trajectories run smoothly through Soviet dissolution as if the state had had no independent effect on the household decision. What moved was the same thing that moved in Iran: mothers becoming literate in their own dialect, one generation ahead of the crossing.

The fertility test isolates the anomaly to one place, not to an ideology. Under WCDE the Central Asian and Caucasian republics carry a $+2.0$ SD fertility residual against the global fit, while the Slavic republics, the Baltics, the Warsaw Pact, and Yugoslavia all sit within a quarter of a standard deviation — their fertility matches their reported schooling. The anomaly is not socialism; it is the metropole reporting on populations it could not see.

10.4 The Moscow Meridian

The fifteen Soviet republics split categorically at Moscow's longitude. Table 13 shows the per-country phenotype residual against the global non-USSR fit (lower-secondary completion at T predicting LE and log U5MR at T, 1960–2020), ordered by signed east-of-Moscow distance.

Table 13: Phenotype residual per Soviet republic vs east-of-Moscow distance, WCDE lower-secondary completion age 20–24, global non-USSR fit.

Republic	East-of-Moscow (km)	LE residual (σ)	log U5MR residual (σ)
Latvia	-844	-1.06	$+0.75$
Estonia	-804	-0.89	$+0.57$
Lithuania	-771	-0.81	$+0.35$
Belarus	-629	-0.96	$+0.43$
Moldova	-547	-1.41	$+1.87$
Ukraine	-444	-0.95	$+1.09$

Republic	East-of-Moscow (km)	LE residual (σ)	log U5MR residual (σ)
Russia (metropole)	0	-1.40	$+1.13$
Armenia	$+431$	-1.42	$+2.77$
Georgia	$+451$	-1.51	$+2.58$
Azerbaijan	$+765$	-2.20	$+3.63$
Turkmenistan	$+1{,}291$	-2.31	$+4.00$
Tajikistan	$+1{,}934$	-2.20	$+3.73$
Uzbekistan	$+1{,}961$	-1.00	—[a]
Kyrgyzstan	$+2{,}285$	-1.80	$+3.23$
Kazakhstan	$+2{,}422$	-1.90	$+2.69$

Notes: Residuals are mean of country-year residuals against the global non-USSR education-phenotype fit at each year, scaled by the year-specific residual standard deviation. East-of-Moscow distance is the great-circle distance along the parallel through Moscow, signed positive eastward. [a] Uzbekistan's WCDE lower-secondary data covers only year 2020, no overlap with U5MR.

Three readings.

The west-east discontinuity. The six westward republics show U5MR residuals of $+0.35$ to $+1.87\sigma$ (mean $+0.84$). The seven eastward republics with multi-year data show $+2.58$ to $+4.00\sigma$ (mean $+3.23$). The categorical jump occurs at Moscow's longitude and is not bridged by closer eastward republics: Georgia at 451 km east shows $+2.58\sigma$, more than three times the residual of Latvia at 844 km west.

Distance is a proxy, not the mechanism. The U5MR residual scales with great-circle distance from Moscow at $r = +0.86$ ($n = 13$, Russia and Uzbekistan excluded). But within the eastward subgroup the slope is essentially flat ($r = +0.27$): once a republic is on the periphery side, additional distance buys little additional inflation. Within the westward subgroup the slope is positive but driven entirely by Moldova — the only westward republic the metropole did not consider European-core (Romanian-speaking, late annexation in 1940, predominantly rural).

The mechanism is bureaucratic, not deliberative. The Soviet system reported on the assumption that all republics had achieved European-core educational standards. For the European-core populations themselves, the assumption was approximately true and the reports were approximately honest. For the

Muslim Central Asian and Christian Caucasian populations the assumption was false and the report was a metropolitan fiction. The discontinuity at Moscow's longitude tracks the empire's identity boundary, not its administrative structure. Formal Politburo representation does not predict the residual: Kazakhstan had a full Politburo member for 17 years (Kunaev, 1971–1987) and shows the same inflation as Tajikistan which had none; Lithuania and Estonia had zero Politburo voice and the cleanest residuals on the panel.

The sharpened reading: the Soviet metropole could not lie about its own kind without immediate phenotypic embarrassment, so for the European core the reports stayed honest. For the periphery, the lie was containable — the populations were far, ethnically distant, illegible to the center, and produced no senior cadres who would correct the record. The phenotype test is what eventually broke the containment.

A blind confirmation. The gradient is not an artifact of the global linear fit that produced the residuals above. The universal transformer of Section 9.9, retrained on the 148-country expansion panel — which excludes all fifteen republics by construction (Section 10.6) — predicts each republic's child mortality from its reported features alone. Its prediction errors reproduce the meridian from the other side: the western republics land $+0.23\sigma$ above their predicted under-five mortality and the eastern republics $+1.43\sigma$ above, and the error scales with east-of-Moscow distance at $r = 0.78$. A method that shares no machinery with the regression — no global education–phenotype fit, no Soviet data in training — recovers the same boundary. (The correlation is read in level space; in log space it is dominated by the seed-unstable residuals of the very low-mortality western republics and is not the comparison.) This is corroboration, not identification: the identification is the discontinuity itself.

10.5 Hanushek's HLO Is Compounded Parental Transmission

The last forty years of development economics has increasingly treated test-score-based "learning outcomes" as the real human-capital variable, with completion treated as a noisy proxy. The Hanushek–Woessmann "knowledge capital" line makes this argument most explicitly: test scores predict growth, years do not, and where years and tests disagree the test is what matters.

The Soviet case is the first hard test of that line against the paper's generational mechanism. Over the contested cohort, HLO secondary scores and reported lower-secondary completion disagree catastrophically: Kyrgyzstan's 2009 HLO of 350 against 2010 reported completion of 99%; Albania's 412 against 96%. If HLO is the real human-capital signal and completion is noise, the paper's headline claims should survive stripping the contested cohort, which they do (Section 10.6). If completion is the real signal and HLO is downstream, HLO should be predicted by past completion — and crucially, by completion that runs further back than current schooling, since the stack of categorical capacities is built across generations: each generation installs rungs the next inherits as the home niche.

Figure 7 resolves it. On a 77-country panel excluding USSR, HLO secondary today is regressed on lower-secondary and primary completion at year $2010-L$, for $L = 0$ to 60 in 5-year steps. The curve is broad and approximately flat across the full sweep. Lower-secondary completion explains $R^2 = 0.504$ at lag 0 and $R^2 = 0.539$ at its peak (lag 10). Primary completion explains $R^2 = 0.469$ at lag 0, peaks at lag 25 ($R^2 = 0.549$), and is still at $R^2 = 0.489$ at lag 60. Primary completion 60 years before the test — the 1950 cohort, more than two generations back — explains nearly half the cross-country variance in today's teenagers' test scores, essentially as much as current schooling does (the bootstrap does not statistically separate the two).

A 2,000-draw country-resampling bootstrap confirms the shape but does not pin the specific peak. The 95% CI on the R^2 difference between peak and lag 0 includes zero in both panels, with the peak exceeding lag 0 in 92% of draws for primary completion and 87% of draws for lower-secondary. The point estimates place the peaks in the parental window; the bootstrap places them *somewhere* between current schooling and the 1950 cohort more than two generations back. What the cross-section rules out is the alternative that lag 60 collapses: the R^2 stays at 0.489 against 0.469 for the freshly-completed cohort's primary. The R^2 is flat where it should not be flat if HLO measured contemporary schools.

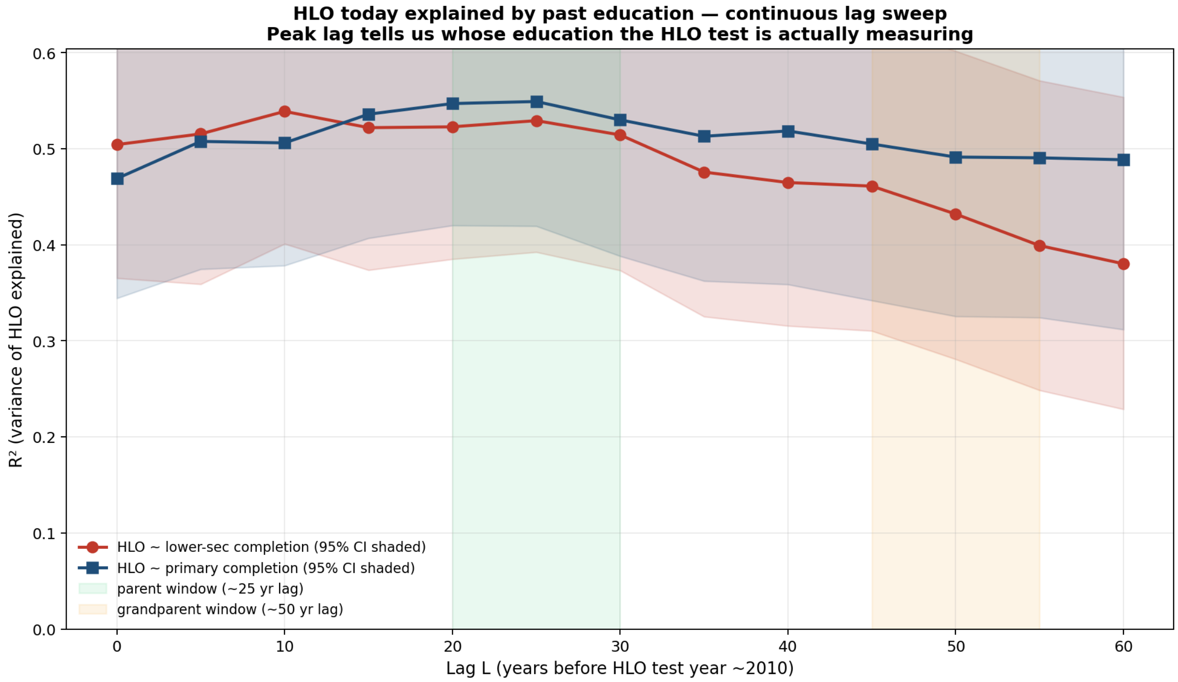

Figure 7: HLO secondary today regressed on WCDE completion at year $2010 - L$, for $L = 0, 5, \dots, 60$. Shaded bands are 95% country-resample bootstrap CIs (2,000 draws). The curve is broad and flat across the full lag sweep; primary completion at lag 60 (1950, more than two generations back) explains $R^2 = 0.489$ of today's variance, against $R^2 = 0.469$ for the freshly-completed cohort's primary. The freshly-completed cohort is a worse predictor than either parental or grandparental-generation completion. *(Source: Angrist et al. 2021 for HLO, WCDE v3 for completion.)*

The direct test on the parental window: HLO secondary today regressed on lower-secondary completion in 1990 (the parent cohort's 20–24 education) gives $R^2 = 0.523$, $\beta = +1.94$ HLO points per percentage point, $t = 9.1$, correlation 0.72, on 77 countries with USSR excluded. Fifty-two percent of cross-country variance in Hanushek's cognitive-skills measure is the paper's quantity measure, lagged one generation. The residual variance is what school systems add *above* the home-niche baseline: China ($+111$ points), Viet Nam ($+111$), Singapore ($+93$), Portugal ($+76$), Japan, Finland, Korea — East Asian education reforms and select rich systems do deliver real additional rungs above what the parental generation transmitted. The negative residuals — South Africa (-116), Ghana (-108), Albania (-86), Montenegro (-72) — mark school systems that undershoot the parental baseline, a credential-inflation diagnostic independent of the Soviet case.

HLO is not a snapshot of current school quality. It is a multi-generational education stock measured at the cohort that sat the test. The mechanism predicts this directly: each generation's schooling is encoded in the parents' inventory of installed rungs, which becomes the home niche the next generation absorbs the school niche through (Chapter 5). Three generations of that compounding produce the variance HLO measures, with contemporary schools adding the visible residual. The cross-country institutional-quality confound that the test-score literature has long worried about ("these countries are just better at everything") is addressed by the panel's country fixed effects (Chapter 9) and by the Korea/Philippines matched pair (Section 8.1), which holds institutional quality roughly constant and varies only the schooling regime, alongside the wider natural experiments and country histories (Chapters 7–8). The lag-sweep is itself cross-sectional — HLO is observed once per country, so the within-country fixed-effects machinery that anchors the rest of this chapter does not apply directly to it. It sits inside the surrounding identification structure as a consistency check: the cross-section behaves the way the panel's mechanism predicts. The causal weight is carried by the panel and the matched-pair experiments; the lag sweep is the test of whether HLO — read across the right time horizon — still tells the same story.

The horse race Hanushek called for partitions as the mechanism predicts: every outcome loads on the same lower-secondary completion, each read forward at its own biological lag rather than on a separate depth of schooling. Section 9.3 reports the full quantitative partition.

10.6 The Panel Excludes the Republics By Construction

The expansion-sample headline design already excludes the 15 USSR republics by construction: their reported lower-secondary completion sat above the 90% ceiling by 1975, placing them outside the active-transition window the regression covers. Re-running the panel with the USSR republics flagged as excluded changes nothing because they were never inside the sample: the parental-education coefficient stays at $\beta = 0.707$ ($n = 945$, 144 countries, $R^2 = 0.561$); LE, TFR, and U5MR forward-prediction R^2s also hold identically (0.384, 0.630, 0.650). The Goskomstat reporting concern affects countries the expansion-window design already screens out.

This is not survivorship under contamination. It is two independent design choices pointing at the same set of countries. The expansion-window screen is mechanism-driven: the panel covers countries while the parental-to-child completion ratchet is still operating (10%–90% lower-secondary in the cohort). The Soviet republics fall outside that window for the same reason any saturated system does. The reporting concern and the regression scope converge.

One sentence summarizes the chapter. My mechanism identifies which Soviet numbers were not real; the panel's expansion-window screen, designed on entirely separate grounds, excludes those same republics from the headline regression. It is the kind of agreement between two independent constructions that one does not get by accident.

11. Difficulties of the Theory

The chapter does one move eleven times. Each framework below identified something real about development — specialisation, wage returns, the Solow residual, geographic inheritance, capabilities, market democracy, institutional quality — and each stopped one level short of the substrate that makes the thing it identified possible. The structure of each section is the same: take what the framework saw, then name the layer beneath it took as given. Two of the eleven (Sen and Drèze, §11.5; Easterlin, §11.11) close the gap rather than rebut it — frameworks I complete rather than displace — but the architectural move is the same in all eleven.

11.1 Smith and the Loaded Labour

Smith (1776) founded the modern account of wealth on specialisation, taking as given the loaded Scottish labour that Knox's parish schools had produced over the preceding century (Section 3.4). The founding error propagated through two distinct descendant lines. The wage-frame line — Smith to Schultz to Becker (Section 11.2) — prices the loaded labour rather than naming what loaded it. The capabilities frame — Sen (Section 11.5) — enumerates what educated agency produces without specifying the categorical reorganisation that makes the agency possible. Both lines rest on the substrate; both omit it from their causal accounts.

Smith found the engine of wealth in self-interest twice: we owe our dinner to the baker's self-love, not his benevolence, and the merchant seeking only his own gain is led by an invisible hand to serve an end no part of his intention. But self-love is the one human constant — it bought no dinner and served no public for the three hundred thousand years before schooling. What changed in his Scotland was the population, not the motive: eighty years of parish schooling had made strangers legible enough to weigh, price, contract, and credit one another (Section 3.4). His hand is invisible because no one moves it; the schooling was invisible because everyone did. He found two marvels in self-love and missed the one beneath both — the loaded labour that was the only reason self-love had anything to coordinate.

The cost is not conceptual alone. Policy built on Smith tells countries to specialise into export markets to get rich. Korea did the opposite; Taiwan did the opposite; Bangladesh did the opposite. Each built the substrate first and the specialisation followed. The mechanism predicts this ordering; Smith's frame, taken as primary, reverses it.

11.2 Human Capital and the Wage Frame

Schultz (1961) and Becker (1964) founded the human capital school on the observation that education raises individual wage productivity and can be treated, economically, as an investment yielding private returns. The framework is the operating frame of labour economics and much of education policy. Measurement of the private return to individual schooling is well developed within it.

The mechanism is that, in the first-generation transition from illiterate to literate, the private labour-market return is the secondary channel. Education's primary effect during that transition is not on the newly-educated person's wages but on their children's lives — through the cognitive and decision-making traits transmitted parent-to-child (Chapter 5). A mother's education shows up in her children's mortality, nutrition, and school completion at effect sizes that dwarf its effect on her own wages. The Gakidou et al. (2010) decomposition (Section 11.7) captures this at population scale during the developing-country transition window: a half-of-mortality-decline intergenerational effect, not a labour-market effect.

Human capital theory treats this downstream effect as an externality and sometimes adds it to private returns. The mechanism reverses the ordering for the transition itself. During the illiterate-to-literate crossing the intergenerational population-level effect is primary and the private wage return is the smaller, more variable side channel. Every country whose population crossed the educational threshold received the developmental compounding, whether or not the policy frame named the mechanism correctly. Korea, Taiwan, and Cuba did it in one generation under singular priority; the others did it slower under competing priorities, but the channel ran the same way. Once a population is past the threshold and home niches are fully literate, the intergenerational effect compresses toward the ceiling and the wage return becomes the dominant marginal signal — but by then the developmental work is done.

The framework's narrower consequence is a misdirection of measurement. Decades of Mincerian wage-return estimation now anchor the policy question of whether schooling is "worth it" in developing countries. The answer from wage data is contingent and often uncertain. The answer from intergenerational outcome data is uniform and large. The disagreement is a frame choice, not a measurement dispute.

Galor's Unified Growth Theory (2011) extends the human-capital frame to the aggregate transition: rising returns to skill trigger a quantity–quality switch in fertility, which produces sustained growth. The empirical signature is the same one I report — TFR falls as education rises — and the cross-country panel is consistent with both engines. The mechanisms are not interchangeable. UGT's engine is the wage-return signal; mine is the categorical reorganisation that the dependency window installs whether or not the wage signal is present. The country histories identify the channel. Korea under Park, Cuba under Castro, Bangladesh through the 1980s and 1990s, and Mao-era China all expanded mass schooling under regimes where market returns to skill were absent, distorted, or politically suppressed. State commitment, not wage signal, drove the loading in each case. The shared scatter plot is what the two engines agree on; the country histories are where they part.

Heckman (2006; Heckman & Cunha 2007; Heckman & Mosso 2014) is the human-capital strand that has reached furthest toward the window. The empirical claim — that returns to investment in cognitive and non-cognitive skills are highest in early childhood and compound across the life cycle — runs in the same direction as the dependency-window argument: what is loaded in the juvenile window has higher leverage than anything added later. The Perry Preschool and Abecedarian follow-up data and the technology-of-skill-formation framework are the rigorous individual-level version of part of the population-scale claim I make. The distance is still a layer. Heckman prices the skills installed in the window and pushes policy toward earlier investment within a still-monetised return framework; the mechanism here is that the window itself is a species-level fact, that the categorical reorganisation it permits is not reducible to a stock of priceable skills (Section 4.2), and that the developmental compounding runs through population-scale home-niche routing rather than individual labour-market wages. The frames agree on the policy implication for the early years; they disagree on what the early years are doing, and on whose returns matter. The closer alignment is with Easterlin (Section 11.11) and Lutz & Kebede (2018), which this paper extends rather than displaces.

11.3 Growth Theory and the Residual

Solow (1956) founded modern growth accounting on a production function with capital, labour, and exogenous technology. The Solow residual — the share of cross-country output growth unexplained by measured capital and labour — is what the substrate looks like when you refuse to name it. Treating labour as an undifferentiated input assigns the unexplained share to “technology”, but technology is not a free input falling on the economy from outside it. Educated populations produce technology, and educated populations are what absorb it. The residual hides the substrate at the creation end and at the adoption end at once. Without categorical literacy at scale, the residual is small or zero — the pre-1800 record is the test. With it, the residual is the population’s own output measured as if it came from nowhere.

Romer (1986, 1990) closed Solow’s gap by treating technology as endogenous — produced by R&D spending inside the economy and non-rival in use. The move from exogenous to endogenous technology is a real advance, and the post-Romer growth literature is built on it. The engine still has no name. Ideas are produced by educated populations and spread among educated populations; the R&D-spending variable Romer measures is the budget line through which a country routes a small fraction of an already-educated workforce into formal research. The substrate is what makes the routing possible at both ends — researchers who can produce the ideas, and a downstream population that can absorb and apply them. Solow says technology is exogenous and unmeasured; Romer says technology is endogenous and measurable as R&D; neither names the educated population that produces ideas at the source and adopts them at the destination.

Mankiw–Romer–Weil (1992)[5] augmented Solow's production function with human capital as a stock variable — secondary enrolment of the working-age population — and reported that the augmented model explains most of cross-country income variation. This is the closest the growth literature comes to naming the substrate, and it is the most-cited move in the discipline. It still does not name it. The stock is priced as a labour-input multiplier, schoolyears monetised through wages, which is the same engine UGT runs at the aggregate scale (Section 11.2). What MRW measures is the cumulative output of past schooling cohorts entering the labour market; what it does not measure is the dependency-window loading that put them there, the home niche that gates absorption, or the household decisions through which the stock's children produce the next cohort's outcomes. MRW prices what was loaded without naming what loaded it.

The arc from Solow through Romer to Mankiw–Romer–Weil is the discipline working at the wrong layer. The residual was never a measurement gap to be closed by adding R&D or human capital as a stock; it was the trace of the substrate — the educated population, registered in the data as a number with no obvious owner. The growth literature has been pricing what loaded the labour without naming what loaded it — the same error Smith made (Section 11.1), with mathematics added.

11.4 Diamond and the Geographic Objection

Diamond (1997) argues that geography — the distribution of domesticable plants and animals, the east–west axis of Eurasia, the biogeographic inheritance of different continents — explains the differential trajectory of human populations. The argument is historical and long-run; it is not primarily about the post-1950 development window the panel here covers. On its own hori-

[5] David Romer; not the Paul Romer of the preceding paragraph.

zon it is largely right. Populations that inherited rich substrates of domesticable species developed settled agriculture earlier, accumulated food surpluses, and built the first stratified societies on those surpluses.

What the geographic frame cannot do is substitute for the mechanism that translates the geographic inheritance into developmental outcomes at any given point in time. The substrate for which environments could be rebuilt — the long juvenile dependency (Chapter 2) and the cultural-transmission channel (Chapter 3) — has geographic prerequisites in its deep evolutionary history; that is not in dispute. What geography does not do is fix the moment at which a particular country decides whether to put children in classrooms, and at what pace.

The Korea/Philippines pairing carries the point at country level (§8.1): the Philippines ahead in 1950 on income per capita, on lower-secondary completion, and on colonial educational inheritance, and behind by 2000 on every development outcome. Geography did not flip the rank order; the educational regime did. The panel is full of such pairings.

Geography is the substrate on which the mechanism runs. It is not the mechanism. The confusion of the two levels is the same shape as Smith's confusion of specialisation with the loaded labour specialisation operates on.

11.5 Sen and Drèze: Two Routings, One Mechanism

Sen (1999) and Drèze and Sen (1989) are the closest frameworks to the one I make. The distance is a single step: what generates the capabilities.

Sen argues that development is what people are able to do and to be, not what they earn. The list of capabilities — to live a long life, to be literate, to participate politically — is the right list, and measuring development in those terms is a correction of the GDP-primary frame. Drèze and Sen's claim that democracy with a free press prevents famine is a genuinely novel observation largely right on its own evidence.

What the frame underdetermines is the mechanism that produces the capabilities in the first place. Sen treats them as parallel expressions of human freedom — coequal axes of wellbeing. The mechanism is that they are serial consequences of a single substrate. Educated cognition makes categorical choice possible (Chapter 4); categorical choice, applied at the household scale, produces the fertility transition, the mortality transition, and the political capacity that the capabilities list comprises. The list is the output of one mechanism, not a set of independent dimensions.

The distinction matters most for the growth-mediated / support-led dichotomy. Drèze and Sen describe two routings of fiscal resources — through market growth or through state service provision — and identify democratic accountability as the feature that determines which populations receive which routing. Both routings, in their own data, produce development only where they reach an educated population. The China case (Section 8.4), matched on mean years of schooling to the fastest movers in the support-led list, shows the residual attributed to "support-led" routing is education itself. Costa Rica was treated as a similar exception in the 1980s — high life expectancy at low income, attributed to its universal-coverage health system — but South Korea, which began that period at much lower income and lower life expectancy than Costa Rica, ran the substrate path harder and now has substantially deeper schooling; by the 2010s Korea overtook Costa Rica on life expectancy and continues to climb past it. The residual that distinguished Costa Rica was real but bounded; substrate growth eventually overtook it. Bihar and Kerala (Section 7.1), under one constitution and one national press, isolate the mechanism: where most households are educated, the community-level shield forms and famine does not arrive; where most households are not, the shield does not form and press coverage alone does not substitute.

The frame is the same frame, read one layer deeper. The fiscal routing is how resources arrive; the educated household is where they become the capability.

11.6 Schooling and Learning

Pritchett (2013) argues that schooling in many developing countries fails to produce learning: children sit in classrooms for years and finish unable to read, do arithmetic, or reason at grade level. The measurement is accurate; cross-national assessments confirm wide gaps between years of schooling and measured proficiency. The policy consequence Pritchett draws is that schooling quality, not schooling coverage, is the development variable.

The biology disagrees on the mechanism. Duration of exposure to the categorical, symbolically organised environment of formal schooling is what restructures cognition (Chapter 4). Test scores measure one narrow output of that restructuring at a particular moment in a child's life; they do not capture the restructuring that has already taken place. The restructuring itself is observable at the neural level, not inferred from behaviour (Section 4.2). These structural changes persist whether or not the curricular content on which they were installed is retained for a particular test. The panel is the population-scale evidence: within countries, as lower-secondary completion expands, fertility, under-five mortality, and life expectancy all move (Chapter 9). The cross-country horse race in which cognitive test scores out-predict completion (Table 6) is not the counterexample it appears to be — I show in Section 10.5 that those same test scores are themselves predicted by the previous generation's years in school, so the cognitive quality Pritchett would substitute for coverage is in large part coverage already compounded one generation back. A child with nine years of low-quality schooling in rural Bangladesh reaches adulthood with a different cognitive architecture than a child with three years of high-quality schooling, and the household-scale behaviour that follows reflects the first child's architecture, not the second's score.

The schooling-without-learning frame is a level-of-measurement mistake. Learning is measured in content absorbed; the mechanism is that schooling's primary effect is categorical — the brain's reorganisation toward symbolic and propositional thought — and that content is the disposable vehicle on which duration rides (Section 4.5). This is a strong claim, and it implies that the curriculum wars that dominate education policy are second-order. The historical record is the test.

Four successful educational expansions ran through ostensibly different content regimes and converged on the same outcome. The convergence itself is the point: the wrapper at the entry of each expansion shed toward secular categorical instruction, because the categorical reorganisation schooling installs (Section 4.2) does not require doctrine and the doctrine drops away.

Scotland after 1560 taught scripture. Knox's parish schools were designed to make every child capable of reading the Bible, because Protestant salvation required direct textual access and Knox did not trust intermediaries. The content was religious at the entry. Over the following two centuries the curriculum secularised — arithmetic, geography, classical letters, and eventually the natural sciences entered the parish schoolhouse, and the scripture share fell. The categorical instruction that produced the highest parish-level literacy in Europe by Smith's lifetime — the loaded labour Smith observed in the pin factory (Section 11.1) — was secular by the time it scaled. Scripture was the starter; secular categorical schooling was what it became.

France under the Ferry Laws (1881–1882) taught secular classical literature, republican history, and a standardised national curriculum designed specifically to displace religious instruction. France entered where Scotland arrived. The effect was mass literacy across the French population, independent of whether the pupil retained the Latin declensions or the list of French kings.

Cuba in 1961 taught literacy through revolutionary primers — Castro's speeches, political slogans, revolutionary narrative. The doctrine was the surface. The teaching underneath was secular: phonemes, letters, reading sentences, doing arithmetic. By the time Cuba's literacy brigades had completed the campaign, the population had been categorically reorganised through secular instruction wrapped in revolutionary content, and the subsequent schooling system that built on it taught the same secular categorical material every other modernising schooling system teaches. Cuba's development trajectory — fertility transition, mortality transition, life expectancy tracking high-income countries despite low GDP (Section 8.3) — followed.

Bangladesh expanded schooling through the 1980s and 1990s with a secular curriculum that, by Pritchett's own measurement, produced substantial learning deficits. The content was weak in exactly the sense that motivates the schooling-without-learning critique. The effect was a fertility transition and a mortality transition that place Bangladesh among the paper's policy over-performers (Section 9.8).

Scripture-then-secular, secular-by-design, revolutionary-wrapper-on-secular, weak-modern-secular. Four entry conditions, one trajectory: secular categorical instruction at scale. The wrappers selected who built the schools and how the expansion was sold; they did not select what the schools did to the children inside them. Duration of exposure to secular categorical instruction selected whether the cognitive reorganisation happened at all. The curriculum wars mistake the wrapper for the payload.

Coverage is the primary policy variable. Quality is a second-order adjustment on top of coverage, not a substitute for it.

11.7 The Deaton Objection

Deaton (2013) attributes the post-1950 mortality decline to global health technology diffusion largely independent of education. The argument is: vaccines, antibiotics, oral rehydration therapy, and safe-water knowledge are cheap, globally available, and easily applied; once the technology is invented, populations benefit without needing to understand the science behind it; the curve rises because the technology spreads.

The distribution of outcomes does not agree. Vaccines have been globally available since the 1960s; under-five mortality fell sharply in some populations and barely moved in others over the same decades. The technology diffused; the outcomes did not, uniformly. Between the existence of a vaccine and a vaccinated child sits a household decision — to locate the clinic, to trust the schedule, to complete the series, to act on subsequent infection warning signs. Each of those is a small act of categorical reasoning on a medical claim. Educated mothers make those decisions reliably; uneducated mothers make them variably. Gakidou et al. (2010), decomposing the global mortality decline, attributed roughly half of the post-1970 reduction in under-five mortality to gains in women's education.

The panel confirms the channel. Residualized GDP — log GDP per capita stripped of education's contribution — has never exceeded $R^2 = 0.019$ as a predictor of under-five mortality or life expectancy in the paper's sample (Section 9.5). If technology were operating independently of education, residualized GDP should retain predictive power; it does not. What remains, when education is partialled out, is noise.

The objection also has a date, and the date refuses it. I run the same within-country regression of cohort education on later life expectancy and child mortality on the only populations whose vital registration reaches the nineteenth century: the early-industrialising states whose mortality the Human Mortality Database records and whose cohort education the WCDE series carries back to 1870, restricted to outcomes before 1945 — before antibiotics or mass childhood vaccination existed anywhere in the sample. The education coefficient is already present and already signed the way the modern panel signs it: life expectancy $+0.404$ (cluster SE 0.095), under-five mortality -2.977 (cluster SE 0.636). Fit that pre-antibiotic era alone and ask it to predict mortality from 1960 forward — the decades Deaton attributes to technology — and the same coefficients carry across the regime break ($+0.566$ and -3.464). The relationship that lowered child mortality before the antibiotic is the relationship that lowers it after.

That window is not pre-medicine; it is the sanitation and germ-theory revolution itself — clean water, sewers, pasteurised milk, the science of Snow, Koch, and Pasteur. But that revolution is the ratchet running (Section 3.2), not a wave arriving from outside it. Its science was the work of the thin elite layer Mokyr names the Republic of Letters, a layer that grows out of the loaded mass population beneath it (Section 11.10); its waterworks were built by educated states; and its hygiene was adopted household by household by populations already literate enough to act on a pathogen they could not see. Those populations had been schooled for generations before the science reached them, and the regression rides the within-country variation in that schooling. Education produced the technology at one end and absorbed it at the other, in the sanitation era exactly as in the antibiotic era. Deaton's diffusion has no exogenous end: both ends are the educated population.

The technology is necessary. It is not sufficient. The mechanism that converts the technology into outcomes is the educated household, one child at a time, and the community-level threshold at which that conversion becomes routine.

11.8 The Methodological Frontier and Its Limits

Two methodological waves have organised the discipline's empirical frontier for the past three decades. The credibility revolution (Angrist, Card, and Imbens — instrumental variables, regression discontinuity, and difference-in-differences) and the randomised-trial movement in development economics (Banerjee, Duflo, and Kremer — field experiments on poverty alleviation). Both produce rigorous local effects of marginal interventions inside operating systems. Neither reaches the channel I identify, and the reasons are methodological in the strict sense.

The credibility-revolution toolkit estimates causal effects of small treatments by exploiting natural variation that approximates random assignment. The country histories the paper rests on (Chapter 8) and the matched-pair natural experiments (Chapter 7) are clean comparisons of exactly the kind the credibility revolution methodology aspires to. What the discipline's headline applications of these methods have not done is converge on the channel as the developmental cause, because the methods have been pointed at marginal questions inside already-loaded systems — the wage return to one extra year of schooling, the test-score effect of class size, the labour-supply response to a minimum-wage change. Each finding is rigorous; collectively they accumulate without converging on the channel, because the methods are pointed at marginal questions and the channel is not a marginal question.

The randomised-trial methodology runs small experiments inside operating educational systems: deworming, conditional cash transfers, teacher incentives, textbook delivery. These produce defensible local effects on the margin of an already-loaded system. They cannot run a population-scale channel-loading experiment because no IRB will randomise 50 million children to twenty years of no schooling. History ran those experiments.

Korea and the Philippines began 1950 with comparable colonial education bases — the Philippines if anything ahead, on lower-secondary completion and per-capita income alike (Section 8.1); political accident assigned one to singular educational priority and the other to competing priorities; one generation later the developmental trajectories had diverged completely. North Korea and South Korea began 1953 as the same population on either side of a frontline; the channel diverged; the outcomes diverged. Cuba ran a literacy campaign in 1961 that loaded the population in eighteen months and crossed the development threshold by 1974. Cambodia's educated cohort was destroyed in four years and the developmental trajectory tracked the cohort's recovery rather than income's. China after 1949 and Niger after 1960 began as comparably impoverished, comparably illiterate states; one founded itself on mass-education commitment and the other did not; seventy years of compounding policy is the experiment.

These are RCTs at the right scale. The treatment is political commitment to mass education, assigned by historical accident with respect to outcome. The control is the country that did not get the treatment. The outcome is observable a generation later.

The hollow-education case (Chapter 10) completes the design. The Soviet republics carried reported lower-secondary completion among the highest in the world over a channel that was hollow; in the language of randomised trials, this is a manipulation-check failure — the treatment administratively imposed without actually loading the variable the theory claims is causal, and the outcome does not follow. The channel is what matters, not the schooling reported on paper.

The credibility-revolution and field-experiment methodologies are appropriate for the questions their applications address: marginal effects of marginal interventions inside operating systems. Applied to the question of why populations develop, the methodology cannot reach the mechanism because the mechanism does not operate at the margin. It is the channel itself. Borrowing that toolkit's vocabulary to defend the finding would borrow the yardstick those tools were calibrated against, and the channel is not on that yardstick. Evolutionary biology, in the 200 years since Wallace, has established the mechanism — drawing on comparative primatology, ethnography, developmental psychology, archaeology, life-history theory, and the evolutionary anthropology that synthesises them. Those methods have produced a settled picture of what the human species is. Development economics has produced a sequence of frameworks — Big Push, structural adjustment, institutions, randomised intervention — that replace rather than extend each other. One field has named the underlying mechanism; the other has measured its outcomes without naming what produces them. My natural experiments deliver the population-scale identification the field's methodology aspires to and cannot run. Where the discipline has read those experiments as "case studies" to be supplemented by something more rigorous, I read them as the rigorous part. The regressions are the supplement.

11.9 The Institutional Challenge

Acemoglu and Robinson (2012) argue that institutions — inclusive versus extractive — are the primary determinant of long-run development. The question is which way the causal arrow runs.

I claim the arrow is asymmetric: education builds institutions, but institutional quality does not measurably accelerate the next generation's schooling. The forward arrow is real; the back arrow is near-zero in the panel.

Institutions act on the adults a country has. Schools act on the children it will have. The window for arranging humans is open continuously; the window for making them is fixed by biology and closes around eighteen. A generation spent reforming institutions optimises the arrangement of an unchanged input. A generation spent schooling the children changes the input itself, and the next generation's institutions are built by different humans. The asymmetry is general: education runs at the speed of demographic metabolism (Lutz 2013); every other policy lever runs at policy speed (Section 5.5).

Institutions are built by educated populations. The civil servants, judges, journalists, health workers, and administrators who constitute institutional capacity are the educated cohort. Without the underlying educated population, institutional reform produces forms without function. Controlling for institutions is controlling for a product of education.

An institution is not a thing a country has; it is an equilibrium it holds against decay. The people who run any institution act in their own interest, and so do the people shut out of it — that is the baseline, not a flaw: incumbents capture it for rent, challengers tear it down to raise their own on the same terms. Every empire that rose had institutions adequate to its rising, and every one of them rotted, because the resting direction of an institution left to its own incumbents is toward extraction. So the presence of good institutions cannot be what explains development — good institutions are read off success after the fact, and they do not stay good. What holds the equilibrium against the rot is continuous pressure from below, and only an educated population can sustain it: a population that can read, coordinate as strangers at scale, monitor the institution, and make itself too costly to ignore. The check that matters sits in the population, which renews every generation, not in the elite, whose incentives are the same in every generation. Reform from the top rearranges the adults a country already has and reverts when they do; schooling changes who the next generation is. Institutions are reasonable now, where no empire's stayed reasonable before, not because the people running them became virtuous but because a literate population forces them to work.

This is where I part from the institutional account at the root. Acquiring institutions was never the hard part. We are a shared-intentionality species (Tomasello 2014): cooperation, and the building of institutions out of it, is our default — not a late achievement, but the baseline. The band had institutions; every human group builds them. So the question is never how a people comes to *have* an institution. It is whether, at the scale of millions, the institution can be *held* — and holding is a matter of force, not virtue. The costliness that forces an institution to work is not a special property of institutions; it is the whole of what any recognition rests on. A rival honours your claim only while crossing you costs more than honouring it. Property, the vote, the right not to be starved so that revenue can be collected — each is a recognition extended exactly as far as its denial can be made expensive, and no further. An institution is the durable shape that balance of force settles into, mistaken after the fact for the source of the recognition it only records. What the evolved inheritance cannot supply is the holding at population scale: its enforcement runs on sight and reputation, which reach the band and stop. Cooperation is destroyed by the free-rider — the member who takes the collective good and returns nothing — and the only thing that suppresses him is being seen: the band caught him by sight, the cultural group by reputation, and both detectors go blind at the scale of millions of strangers (Tomasello 2016). The incumbent who captures the institution for rent is that free-rider at the radius where sight and reputation no longer reach, and what restores the cost on him is a population literate enough to see the capture at a distance. A crowd is not strength. A population that can read, coordinate as strangers, monitor, and withhold in concert is — and mass literacy is the one renewable source of it.

Seen this way, the engine the institutional account proposes — the virtuous circle, in which inclusive institutions widen the franchise and spread secure property rights of their own momentum — is a balance of force misread as a property of the institution. Post-revolution England is held up as the birthplace of inclusive institutions — but inclusive for whom? About 400,000 men could vote before the Reform Act of 1832; after the celebrated extensions of 1832, 1867, and 1884, roughly 40% of adult men still could not, and no woman voted on equal terms until 1928. The same Parliament, in those same decades, ran the largest extractive machine in history: as many as ten million dead in the Bengal famine of 1770 under Company rule, one million dead inside the United Kingdom itself in the Irish famine of the 1840s, three million in Bengal again in 1943. The franchise extensions are read as the circle improving itself; they are nothing of the kind. Each is a bloc — workers, tenants, women — grown literate and coordinated enough to make its exclusion cost more than its admission. Nothing in the institution reached out to include them; they forced the door. Property tells the same story in its hardest form, because property is zero-sum where the vote is not: securing the landlord's title *was* extinguishing the commoner's customary right, and the Company's secure claim on Bengal's revenue *was* the famine. The right and the expropriation are one act seen from two sides of an edge, and the edge is who could make a claim expensive to deny — the propertied man who could organise and vote on the inside, the Irish tenant and the Bengali peasant and the unenfranchised on the outside, at the same time, under the same rules. Where a population was kept illiterate by design — the colonies — there was no force and there was pure extraction, through the same inclusive Parliament. "Inclusive" is not a property of an institution; it names who currently has the educated strength to stand inside the radius. The boundary of inclusion is the boundary of educated pressure (Section 5.2); the same scaled coordination that pulls the metropolitan worker inside extracts the colony (Chapter 13).

Institutions are therefore downstream of education by construction, not a parallel cause. A court, a ministry, a bank, a registry is a structure for coordinating non-kin at a scale the band never reached, held by written rule and symbolic role rather than by kinship and line of sight. It is not that stranger-cooperation is absent from our inheritance — it is that the inheritance enforces it only at band scale, and the apparatus that carries it to the scale of millions, the writing and the registry and the law and the literate population that can wield and read them, is what literate CT supplies (Section 4.3). Agriculture built large polities before that flip — on a literate elite over an illiterate-CT population — and that complexity rose and then collapsed (Section 4.3); only the flip to a whole literate population makes the scaled coordination hold. The scale an institution can reach is bottlenecked by the population's literacy, not the reverse. The capacity is the point, not the virtue: the same scaled coordination that builds a public-distribution system builds a war machine

(Chapter 13).

Country fixed effects absorb time-invariant institutional characteristics. Time-varying institutional change is itself driven by education. Qatar ($69,000 per capita, stable institutions by standard measures) delivered 3.1 percentage points below its generational education baseline in 2015. The institutions were present; the educational commitment was not; the development did not compound. The Philippines (Section 8.1) confirms the same point from the opposite direction: equivalent initial institutions, slower education, slower convergence.

India has had democratic institutions, an independent judiciary, and a free press since 1947. China had none. China expanded lower secondary completion from 10% to 75% in 40 years (1950–1990) at 1.6 pp/yr — under Mao, the Cultural Revolution, and the Great Leap Forward. India expanded from 10% to 37% in the same 40 years at 0.7 pp/yr — with every institutional advantage the literature claims to matter. India's institutions did not accelerate education; China's absence of them did not prevent it. Among countries in the active-transition window (lower-secondary completion 10–90% at decade start), the mean expansion rate peaked at 1.52 pp/yr in the 1960s — before most postcolonial institution-building — and fell to a trough of 0.74 in the 1990s, when the inclusive-institutions consensus was at its strongest. The series rebounded to 1.29 in the 2010s. Speed and institutional maturation do not co-move; in this panel they run counter.

Autocracy does not predict education speed. Merging Polity5 time-varying regime scores with the WCDE completion data (160 countries, 1950–2020, lagged 20 years to match the regime that made the schooling decision), polity2 explains less than 0.7% of the variance in education gain rates — the same null holds at 0-year and 15-year lags. Democracies are slightly faster on average: 10.3 pp/decade against 8.1 for autocracies at the 20-year lag.

What autocracy produces is not speed but *variance*: 74% of autocratic countries fall below the democratic median gain rate, but the fat right tail includes Korea and Cuba. Autocracies populate both the fastest and the slowest education trajectories in the dataset. Among 50 countries that transitioned, 30 were faster under democracy (paired mean $= +0.6$ pp/decade, $p = 0.57$).

Regime type is a variance amplifier, not a cause. The decision variable is education policy (Section 16), and it is orthogonal to regime type.

The colonial instrumental-variable contest (Section 7.4) makes the same point from the institutional literature's own preferred instrument; the structural test follows in Table 14.

The pattern is general. Low settler mortality marks where Protestant settlers survived to build schools. Coloniser religion picks up Reformation-era literacy. Each instrument's identification rests on its correlation with educational loading; remove the educational channel from the picture and the instruments' identification of any institutional channel collapses into a horse race the IVs were not designed to run. The instruments identify education; the literature has been calling it institutions.

The colonial test (Section 7.4) puts education and AJR's own institutional measure into a bivariate horse race on log GDP per capita across AJR's 64-country base sample. The two are highly collinear ($r = 0.62$), and adding coloniser religion on top of education does not raise the fit.

Table 14: Colonial education vs. AJR's avexpr as predictors of current GDP.

Predictor	R^2 (log GDP 2020)
Education at independence (1950 lower sec)	0.518
AJR *avexpr* (institutions, 1985–95)	0.525
AJR settler mortality (*logem4*)	0.544

Notes: AJR's base sample (n=61 with all data after WCDE/WB merge). Education and AJR's avexpr explain near-identical shares of log GDP variance because they share the same colonial-era upstream cause. Education at WCDE v3, GDP at World Bank WDI (constant 2015 USD), avexpr and logem4 from AJR (2001) replication archive.

11.10 Innovation and the Cultural Substrate

Aghion and Howitt (1992) is the Schumpeterian extension of Romer — innovation as serial creative destruction, with the same R&D-driven engine and the same substrate gap (Section 11.3). Mokyr is the harder target, because he reaches further than any major economist toward a cognitive-cultural substrate.

Mokyr (1990, 2002, 2009, 2016) argues that the Industrial Revolution was driven by a cultural substrate: the Republic of Letters, the rise of contestable beliefs, the seventeenth-century turn toward useful knowledge, and the Enlightenment's institutionalisation of open scientific inquiry. The argument has the right shape. A cognitive-cultural state of a population produces sustained technological growth; the state is historically specific and politically contingent; without it, the same material conditions produce no industrial transition. Mokyr is the only major economist who has reached for a substrate frame, and the substrate he names is real.

Mokyr's framework distinguishes propositional knowledge (the elite epistemic shift) from prescriptive knowledge (artisan technique and useful know-how), and treats the connecting institutions — mechanics' institutes, learned societies — as the transmission belt between them. The causal weight nonetheless runs from the elite tier down: contestable beliefs in the Republic of Letters change the norms, and the norms propagate to artisans through institutional channels.

The frame names the wrong layer. The Republic of Letters played in salons, journals, and learned societies of perhaps a few thousand correspondents across Europe; the Industrial Revolution was assembled by hundreds of thousands of parish-school graduates working looms, lathes, and forges in workshops across Britain. Knox's parish schools (Section 3.4) produced the loaded labour that made Smith's pin factory possible a century before Mokyr's Enlightenment thinkers met. Mokyr describes the elite layer; the population layer is in the data underneath his account, not central to it.

The two layers are not parallel; the elite rises from the population substrate. The engineers, philosophers, and experimentalists Mokyr describes are a small fraction — a hundredth, a thousandth — of a loaded population. Indian and Chinese brains in 1900 were no less capable than the engineers from those countries now filling Silicon Valley; what changed is that both populations were loaded through schooling. Take the loading away and the same brains do the work their forefathers did — subsistence agriculture, the labour of the unloaded. Song China, Abbasid Baghdad, and Renaissance Italy produced elite cognitive cultures without industrial transitions because the elite was thin and the population beneath it was not loaded. The Chinese imperial examination system selected a few individuals from that population; it did not load it. The population substrate is the layer at which the post-1950 convergence runs (Chapters 6–9), and at which the Industrial Revolution itself was assembled. Mokyr names the elite; I name the layer that grows it.

11.11 Easterlin's Question

Easterlin (1981) asked, and answered, the question I extend: *Why isn't the whole world developed?* Lutz and Kebede (2018) extended that answer empirically; I extend both.

The question's persistence in the literature four decades after Easterlin asked it reveals what the discipline of development economics has done with the answer he gave. The frame that took hold instead was growth-first, with the claim that education would follow, adjusted to whichever framework was ascendant: human capital in the 1980s, institutions in the 2000s, schooling-versus-learning in the 2010s. None of those frames permits Easterlin's answer, because each subordinates education to something prior. The substrate-is-primary claim is what those frames cannot absorb.

The whole world is not yet developed because the substrate-level investment is generational, and because most countries, in most decades since 1950, did not make the singular-priority decision (Section 16). The countries that did — Korea, Taiwan, Cuba — compressed the transition to one generation. The rest proceed slower under competing-priority regimes. Easterlin asked the right question and gave the right answer; I add the biology beneath it, the empirics that confirm it, and the invisibility argument that explains why his answer did not take.

12. Why the Loading in Childhood Is Invisible

Childhood is visible. The loading happening inside it is not. The invisibility is the structural reason every framework that has reached for the mechanism has stopped where the educated frame stops. The capacity to rank education as fundamental requires a cognitive frame that only education produces; and for households without it, the return is unimaginable. I make the invisibility precise here and close by turning it on the present synthesis.

12.1 Education Leaves No Physical Signature

An educated person is physically indistinguishable from an uneducated one. Same body, same biological drives, same appearance. Wealth builds visible infrastructure; disease marks the body; education rewires cognition and leaves no trace on the surface. The only observable signal is behavioural change — and the disciplines misread it. They see smaller families and call it wealth flows, or contraceptive access, or patriarchy relaxing, or telenovelas on television. They see longer lives and call it health systems. They see rising productivity and call it economic growth. Every discipline theorises the downstream behaviour without tracing it back to the cognitive reorganisation that produced it. Education has no physical signature, so the signature it does leave — changed decisions, across every domain, for the rest of the person's life — is assigned to whatever domain notices it first.

The timescale compounds the invisibility. Education works through demographic metabolism (Section 5.5; Lutz 2013): its effects appear as new cohorts replace older ones, on horizons of twenty-eight years, not in the year-over-year cycles that policy attention, news, and electoral politics are tuned to. A signature too slow for a news cycle and too distributed for a policy ledger is a signature easy to miss, even when its accumulated mark is the largest in the developmental record.

12.2 Invisible From Inside

Educated people rank education as important — one of many important things, alongside health, nutrition, governance, and security. They cannot categorise it as *fundamental* because that would require seeing their own cognition from outside. The act of ranking priorities is itself an educated act; the cognitive architecture that assembles competing inputs into a weighted comparison is the product of the substrate I describe. The medium cannot be seen by those thinking through it.

Pinker's (2018) *Enlightenment Now* is a representative case. Pinker attributes the post-1800 rise in life expectancy, literacy, wealth, and peace to "reason, science, humanism, and progress." The rise itself was education. Literacy is education by definition. Life expectancy, wealth, and peace are what educated populations produce — longer lives because educated adults care differently for themselves and their children, material gain because educated labour can be loaded with specialised skill, reduced violence because educated populations have fewer children, releasing the demographic and Malthusian pressure that drives conflict. Reason, science, and humanism are what schooling installs in the mind, the cognitive residue of the same mechanism. Pinker has described education twice — once as outcome, once as mechanism — without recognising that formal schooling is the common cause of both sides. The book is a compendium of the downstream outcomes of the mechanism in Chapters 2–4 without a chapter tracing them to schooling. A reader inside the educated frame reads this as persuasive because the frame is what produces the reading. From outside the frame, the chain is one step, not two: mass schooling produced successive cohorts of educated adults, and those cohorts produced the measurable gains Pinker documents.

Subjective wellbeing confirms the blindspot: across 159 countries, self-reported happiness tracks income, not education — people feel richer, not more educated, so they attribute progress to wealth. The signal educated populations receive from their own lives is material improvement, not cognitive transformation; the transformation is the medium through which they perceive the improvement.

There is a structural reason the misattribution runs so reliably to income. The convergence is emergent and leaderless — the flock's turn, not a movement's campaign. No one directs it from above and no one organises it from below; each household decides privately and sees only its own improving life. With no leader and no movement for the eye to land on, the observer credits the single thing that is visible and counted: the money. It is invisible from inside not because it is hidden but because a shape no one made looks like no one's doing.

12.3 Unimaginable From Outside

For uneducated households, the cost of schooling is concrete and immediate: the child not working, school fees, uniforms, distance to the school, the opportunity cost in daily labour and childcare. The return is abstract, delayed by a generation, and has no precedent in the household's experience. The frame required to evaluate the return — to imagine the cognitive reorganisation of one's future grandchild, and the household decisions that reorganisation will enable, and the population-level outcomes that will cascade — is itself the product of the education the household lacks.

Mullainathan and Shafir's (2013) work on the cognitive bandwidth tax of poverty is load-bearing here. Scarcity consumes the bandwidth required to evaluate abstract, delayed returns. An uneducated household under daily scarcity cannot process a multi-decade educational investment with the same cognitive apparatus an educated, resource-secure household brings; the apparatus itself is partially produced by the return they are being asked to imagine. The transformation is not just invisible — it is unimaginable. This is why state reach matters (§5.4): the state has to deliver the first generation of schooling into households that cannot voluntarily choose it because the choice itself requires capacities they do not yet have.

The people who could identify education as fundamental cannot, because they are inside it; the people who would benefit most cannot, because they are outside it. This is the structural form of the invisibility.

12.4 The Dilution Mechanism

Education simultaneously widens who the educated person invests in (literacy connects her to populations she has never met) and creates surplus to invest (fewer children, healthier children, longer productive life). The combination disperses investment across every cause simultaneously: health, poverty, gender equality, climate, governance, animal welfare, the arts. This is not a failure of educated people; it is the structure of educated behaviour. The educated adult, by the time she can evaluate priorities, has absorbed the frame in which priorities are plural by default.

The result is a global development discipline that treats education as one input among many. This is what the competing-priority regime (§5.4) looks like at the intellectual level: the educated frame produces dispersion as its native output, and the institutions built by educated people inherit the dispersion. Education sits at number four in the SDG framework not because the framers mis-specified, but because the frame through which they specified is itself the output of education operating at its widest radius.

The clearest contemporary demonstration of the dilution mechanism is the United Nations Sustainable Development Goals (UN 2015). The SDG framework was designed by the most credentialed development institutions on earth, with full access to the evidence I cite — including Haq and Sen's own Human Development Index, Easterlin's identification of mass schooling as the binding constraint, and the income–health literature showing education proxying through income. The result: SDG 4, Quality Education, listed fourth among seventeen parallel goals — alongside SDG 1 (No Poverty), SDG 2 (Zero Hunger), SDG 3 (Good Health), SDG 5 (Gender Equality), SDG 8 (Decent Work), and SDG 10 (Reduced Inequalities). Every one of these is a downstream product of education.

Income is what educated populations generate — stripped of education's contribution it predicts nothing (Section 9.5). Hunger is what educated households escape — not through provision but through the fertility transition that releases resources per child, the planning horizon that converts subsistence into surplus, and the household-level cognitive and behavioural capacity that prevents food shocks from becoming famines. 19 of 21 post-1950 famines occurred where lower-secondary completion was below 50% (Section 7.1). Gender equality is what the fertility transition produces first — the shift from biological fate to conscious reproductive choice that occurs at primary education, before labour force participation, before political voice, before any of the dimensions SDG 5 measures independently.

The SDG framework does not reflect ignorance of the evidence. It reflects the mechanism: the educated people who designed it experienced income, food security, health literacy, and institutional capacity as the normal baseline from which you design frameworks — not as products of education that the uneducated world does not yet have. Education at number four, competing for budget allocations against hunger programs that feed children today, poverty interventions with visible short-term outcomes, and health delivery systems that save lives this year, is not an oversight. It is what the competing-priority regime looks like when institutionalised at the global scale by people who cannot see they are inside it.

12.5 Historical Exceptions Confirm the Rule

Inside the development frame, education competes with health, nutrition, governance, and security in a cost-benefit comparison — and it always loses, because its returns are generational while the others' are immediate and visible. A child dying of diarrhoea today is an emergency; a child not sent to school today is no one's emergency — until her child arrives as one in 2050.

Every historical breakthrough bypassed this comparison. Knox, Meiji Japan, Korea, and Cuba made education non-negotiable for reasons that preceded any cost-benefit calculation — salvation, existential military threat, Cold War survival, revolutionary ideology. Each of these bypasses was produced by external compulsion that overrode the educated default of treating education as one priority among many. Knox wrote under the conviction that every Christian must read the Bible; Meiji reformers under the threat of colonisation; Park Chung-hee under Cold War military necessity; Castro under revolutionary ideology. In every case, the singular-priority decision was made *against* the default educated frame of plural competing priorities, not from within it.

This pattern confirms the invisibility claim. Without external compulsion, the educated frame produces dispersion, and dispersion produces the competing-priority regime (§5.4). With external compulsion strong enough to override the educated default, the singular-priority regime becomes possible — and where it has been attempted, the development threshold has been crossed within one generation, as Chapter 8 documented above. The policy consequence — what it takes to replicate a bypass from evidence rather than contingency — is Chapter 16.

12.6 This Applies to This Paper

My argument applies to its own production, and to the production of every argument, discovery, and specification in the history of ideas. I received categorical literacy through formal schooling. The biological components of the mechanism — Konner on juvenile dependency, Boyd and Richerson on cultural transmission, Hawkes on post-reproductive lifespan, Dehaene on literacy's neural reorganisation — and the development components — Haq, Sen, Easterlin, Lutz and Kebede — arrived through the same transmission channel

I describe. The synthesis is that channel operating on itself: an instance of what educated minds do with accumulated payload, no different in kind from any other output of the mechanism.

Every foundational specification is produced this way. Smith wrote from within an educated Scotland whose parish schools had been building literate labour for a century before him; he took the literate workforce as given and theorised the division of labour it enabled. Wallace wrote from the Malay Archipelago, having left formal schooling at fourteen; he carried the accumulated tradition of Linnaean taxonomy, Lyellian geology, and Malthusian population dynamics — each diffused to him through books and lower-tier British schooling rather than the Edinburgh–Cambridge elite circuit (Wallace 1870). Darwin reached the same synthesis independently because the tradition was ready. Shannon wrote from within the mathematical-engineering tradition of Nyquist, Hartley, Wiener, and Turing. Mendeleev assembled Döbereiner, Newlands, and Meyer into the periodic table; Meyer reached a parallel arrangement the same year. In each case the specifier is a node in a transmission chain that has reached the threshold where the chain's own mechanism can be articulated from within.

Traditions specify their own mechanisms at readiness-moments. The specific author is contingent; the readiness is not. Parallel synthesis — Leibniz and Newton, Wallace and Darwin, Meyer and Mendeleev — is the usual signature of a tradition arriving at its own self-description. What matters is that the accumulated components reach the point where synthesis becomes possible, not who performs the synthesis.

This closes the invisibility argument. The capacity to see the mechanism from inside is produced by the mechanism itself, operating across enough generations of accumulated knowledge that the channel becomes describable in its own terms. Education is invisible at shallow depths of accumulated understanding and specifiable at sufficient ones. The present synthesis is the tradition Haq and Sen opened in 1990 reaching the depth at which its own operative channel can be named. The paper is the channel speaking about itself — late, contingent, and no more exterior to the phenomenon than any other educated output ever has been. What separates this synthesis from competing ones is empirical survival across every line of evidence — the deep-history mechanism, the natural experiments, the country histories, the panel, and the Soviet falsifier — not its provenance.

13. The Dark Parallel

The positive case has been stated, defended, and shown to be invisible even from inside: education is the substrate on which development runs. What the positive case does not say, on its own, is what the same substrate does when it is pointed at something other than a population's own flourishing. This chapter says it.

The country histories and the European window are both history. The country histories show the educational substrate pointed at a population's own flourishing; this chapter shows the same substrate pointed outward at other populations and inward at subjugated groups. Both rest on the attested record.

13.1 Capacity, Not Virtue

The mechanism produces capacity for coordinated collective action (Prediction 6). That capacity forms the community shield of Section 7.1 — food reorganised under drought, fertility decisions made at the household level, child mortality reduced without waiting for the state. The same capacity, pointed outward at a rival population, organises armies, plans blockades, and sustains a war economy long enough to break the other side. Pointed inward at a subjugated group within a population, it organises domination at whatever scale the population's educational base permits.

The shield and the amplification are the same mechanism, rotated. An educated population is not a peaceful population. It is a capable population. What the capability is directed at is a political matter, not a biological one. The preceding chapters described what the mechanism does when it is pointed at a population's own survival. This chapter describes what it does when it is not. The capacity is the population-scale residue of how loaded each generation's childhoods were; the direction is set by what the directing class chooses to do with the residue, not by anything in the loading itself.

13.2 Demographic Pressure

Malthus (1798) identified the structural pressure: human populations, absent a check, outgrow the resources that sustain them. When growth outruns carrying capacity, the options narrow. Populations expand into new territory, or contract through starvation, or are destroyed in conflict with neighbouring populations doing the same. Peaceful coexistence is a function of slack — resource abundance relative to population. When slack disappears, the kill-or-be-killed structure forms. There is no middle ground when the resource base will not support the bodies. Cohort size is set by the mothers' fertility decisions in the previous generation, and those decisions are set by how loaded the mothers' own childhoods were. Pre-transition fertility — six or seven children per woman — is what unloaded childhoods produce at population scale.

The educational substrate supplies the organised response to that pressure at scale. Uneducated populations under Malthusian pressure can riot, migrate in disordered waves, starve. Educated populations under the same pressure can organise armies, project navies, build and sustain settler economies on other continents, administer conquests across distance. The pressure is the same; the instruments are different. This is why the great engines of inter-group violence correlate with demographic peaks more reliably than they correlate with particular ideologies. The pattern is older than ideology: chimpanzee coalitions raid neighbouring territories under resource pressure (Goodall 1986; Wrangham & Peterson 1996), and the primate substrate we inherit treats boundary-drawing and pressure-triggered coalitional violence as a package, not as a product of beliefs about what the fight is for.

The ideological frames human participants use to describe what they are doing — religious, national, civilisational — are the organising symbolism on top of that substrate. The frames do load-bearing work. Hostility is not murder. The primate baseline, and the ethnographic baseline of small-scale intergroup conflict, is episodic, opportunistic, bounded by the logistics of small-group mobilisation. Cumulative mortality across generations can be high, but the cumulative total is the product of repeated bounded engagements under chronic demographic pressure, not of single campaigns at industrial scale. Ideology does not unlock hostility — the orientation is already there. What ideology does is scale episodic bounded violence into sustained mass killing.

The out-group is reframed as an existential threat, as sub-human, or as the beneficiary of its own destruction ("we are doing them a favour while taking their land"), and the cost of killing is distributed through institutional hierarchy until it becomes bureaucratically ordinary. The pressure underneath is demographic and material; the frames are what convert the baseline into industrial violence. When the pressure releases, the frames lose their conscript force — the stake disappears and the conflict reverts to its baseline range.

The ideological frames are not epiphenomenal, but they are arbitrary. Out-group identification is the species' baseline orientation at the group boundary. Hare & Woods (2020) argue that human self-domestication selected for in-group cooperation — the positive trait that makes large-scale human coordination possible at all — and heightened out-group hostility is what that trait produces when a boundary is drawn. The two are consequences of the same evolutionary history but operate differently: in-group bonding holds the population together, out-group orientation governs how it meets another. Human populations will find an out-group; which category gets elevated to that role — religion, ethnicity, language, nation, caste, class — depends on which markers are locally available. The biology supplies the orientation but not the target, so coordination on any particular marker falls to whichever directing class has both the reach to impose one and a stake in the choice. The primate baseline already shows the logic: out-group is defined by whichever boundary is salient (a neighbouring troop, a stranger), not by the content of any belief about the stranger (Goodall 1986). In humans the markers are richer and elite-directable, but the orientation itself is inherited. The content of the othering is arbitrary. What is not arbitrary is that demographic pressure combined with out-group identification produces violence at whatever scale the educational substrate permits.

13.3 The European Window

European expansion from approximately 1500 to approximately 1960 has been read as educated Europeans dominating uneducated others. The reading is too clean. Europe in that window was not a uniformly educated civilisation. It was a small literate elite — clergy, merchants, administrators, military officers, later industrialists and engineers — sitting atop a largely uneducated mass population whose fertility had not yet begun to decline. Mass education in Europe arrived late. Scotland after Knox was exceptional; most of the continent reached mass schooling only in the last century and a half. England reached an Education Act in 1870; France reached the Ferry Laws in 1881–82; Prussia earlier but unevenly. The European fertility transition arrived in the same window, and for the same reason (Section 5.3).

What that configuration produced was a machinery, not a virtue. The educated elite had enough organisational capacity to build states, navies, finance, navigation, and industrial infrastructure (Section 3 on the tool sequence). The uneducated masses supplied the demographic pressure — pre-transition fertility produced a population boom that outran European agricultural yields — and the bodies: the soldiers, sailors, settlers, and conscripts who carried the projection outward. The populations the projection reached had neither elite capacity matching the projecting state nor a mass-level community shield; they absorbed the costs.

Ordinary European migrants, soldiers, sailors, and settlers were human beings responding to demographic and material pressure with the instruments the mechanism supplied. They are kin to the Mughal peasant conscripted into imperial service, the Qing subject sent to settle frontier provinces, the Roman citizen shipped to a colonia, the Phoenician family boarding for Carthage, the Bantu lineage absorbing smaller neighbours on the way south, the Arab trader extending the reach of caliphates across the Indian Ocean. Ordinary people in every imperial age did what ordinary people do: they took the opportunities available, they accepted the costs they did not see any way to refuse, and they died at rates they did not choose. There is nothing in the European window that is not also in the long human record of organised expansion under pressure. The scale is larger because the capacity is larger; the humans are what they have always been.

The Chola maritime expansion across the Bay of Bengal in 1025, reaching Srivijaya and Kedah, instances the same structure from an Indian base centuries before the European maritime projection. The substrate-and-pressure combination produces outward projection in whatever geography it obtains, and the form it takes (continental, maritime, tributary, settler) reflects local conditions rather than population-specific tendencies.

The directing class at the top of the imperial machinery made decisions whose costs fell on populations unable to resist. They are not named here because I write about a mechanism, not about individuals. They are also not defended.

The postcolonial account that locates the current state of the Global South primarily in colonisation cannot carry the weight it is asked to carry. But the account that locates those conditions in post-independence failure is equally uncalibrated. Post-independence states have operated, almost without exception, under the competing-priority regime (Section 5.4): education in the national plan, the budget, the five-year strategy, the international compact. Europe, for most of its four-century imperial ascent, operated under no priority at all. On the regime dimension I use, the post-independence Global South has operated at a level Europe itself did not reach until late in its own imperial age. What varies across the post-independence countries is speed, not direction. A handful — Korea, Taiwan, Cuba (Section 8) — made education the singular priority and compressed the transition to one generation. Most operate under competing priority and are proceeding slower; Bangladesh and Sri Lanka are the competing-priority over-performers (Section 8.3). I neither celebrate the first group as exceptional nor indict the second as negligent. Competing priority is itself a regime above anything Europe ran under for most of the period during which Europe dominated the world.

Everyone — coloniser, colonised, postcolonial — has been human in the same way, operating under different regimes at different moments.

13.4 The Evolutionary Baseline

Hierarchies, empires, out-group violence, and organised domination are not unique to the educational window, and not unique to Europe. They are the baseline condition of settled human societies going back to the first states in Mesopotamia and the Yellow River, and of pre-literate human societies before that. The biological chapters establish the substrate: de Waal (2013) on primate hierarchy; Goodall's (1986) four-decade Gombe record, and Wrangham & Peterson (1996) building on it, on coalitionary male violence as a deep ancestral trait shared with chimpanzees; Hare & Woods (2020) on the in-group orientation that accompanies self-domestication. The human species is not a naturally peaceful population interrupted by occasional lapses. Out-group violence and hierarchical domination are the baseline; inter-group peace at large scale is the historical anomaly and is the phenomenon that requires explanation.

The educational substrate does not create hierarchical domination. It amplifies the reach of whatever hierarchical arrangement a population's capacity permits. Small educated elites built the first states, the first organised religions, the first slave systems, the first conquest empires. Larger educated populations built larger and more durable versions of each. When mass education finally arrived in one region, the reach of organised violence scaled proportionally. This is not a charge against education; it is a description of what capacity does.

The honest reading is not that education corrupts an otherwise peaceful species. It is that the species is what it has always been, and education raises the ceiling on what the species can accomplish — in both directions.

14. Convergence and Peace

Three panel findings carry this chapter forward from convergence: the demographic transition follows mass education (Section 5.3); the community shield forms when most households cross the threshold (Section 7.1); residualised log GDP per capita carries no independent predictive power for life expectancy or fertility, and only a small bounded effect for under-five mortality (Section 9.5). The claim: the European window's projection has closed, but the engine that powered it — demographic pressure inside an uneducated substrate — runs in the background of the wars Europe finished by the 1970s and most of the rest of the world has not, and completing the transition retires it.

14.1 The Asymmetry Reverses

The European window (Section 13.3) was a single configuration: educated population plus demographic pressure on one side, absence of capacity on the other. The configuration depended on an internal asymmetry of childhoods — a literate elite above an unloaded mass — which supplied the organisational reach at the top and the demographic pressure and bodies at the base. The directionality required both terms. Capacity without pressure would not have projected; pressure without capacity could not have. That the projection ran outward and not inward was a fact about which side carried which term, not a fact about virtue or hostility on either side (Section 13.1).

The current era reverses one term and not the other. The educated populations have completed the transition; demographic pressure has lifted from them. The populations still mid-transition carry the demographic pressure but not the capacity. The historical configuration cannot run in reverse: uneducated demographic pressure has no instrument with which to project organised force against an educated population. Riots, disordered migration, and starvation are within reach; armies, navies, settler economies, and sustained inter-continental projection are not.

The floor for projection has also risen. Europe-1700 ran its outward window with sail-era instruments above a largely uneducated mass, projecting against populations whose only substrate was their own home niche. Today's projection — military, administrative, or extractive — requires deeper loading inside the projecting population (logistics chains, capital flows, administrative systems that no small-elite-on-mass configuration sustains) and faces populations on the receiving end whose state capacity, vector control, and defensive infrastructure are themselves outputs of the channel (Section 5.4, Section 7.1). The deep uncovered substrate that absorbed the European window is mostly gone. Each completed transition raises the floor for the next attempted projection by adding a layer of defensive substrate to the world it would have to cross.

The empirical signature of the asymmetry is already visible. Demographic pressure from the populations still mid-transition manifests as immigration into the educated north's labour stack — at every layer, from working-class agricultural and care labour up through professional-class engineering, medicine, and finance, up to CEO and founder ranks. The flow integrates into existing institutional structures rather than displacing them. The receiving populations absorb the labour rather than being conquered by it. The mechanism is the four-radii reversal (Section 14.4) running on the inward radius: people moving toward where the substrate is already in place, rather than the substrate's bearers projecting outward to where it is not. The next sections trace what happens as the configuration that produced this asymmetry completes.

14.2 The Malthusian Engine Runs Out of Fuel

The Malthusian engine (Section 13.2) that powers organised inter-group violence has one off-switch the species has ever found, and the preceding chapters have described it. The off-switch is loaded childhoods. When most of a population's children pass through eighteen years of formal schooling, the community shield forms (Section 7.1), household-level fertility decisions arrive (Section 5.1), population growth slows or reverses, and the demographic pressure that makes populations expandable-by-force or forced-to-expand releases. The window that produced the past four centuries of asymmetric organised violence (Section 13.3) was the window during which one region had arrived at the substrate and the rest had not. Convergence closes that window.

14.3 The Resource Engine

Resource competition is the other engine of organised violence — coal and the colonies, oil and the twentieth century's wars, water and the Sahel today. The energy question that haunted earlier calls for universal development is being resolved through the same channel as development itself: solar–wind–battery is the latest layer of the tool sequence (Section 3.3), produced by educated populations and adopted by educated workforces. The mechanism that retires the demographic engine also retires the resource scarcity that made expansion structurally rational in the fossil-fuel era. The variable that still requires a country-level choice is education (Section 2).

14.4 The Four Radii Reverse

The reversal is already partly visible. The four radii of educational effect (Section 5.2) describe how an educated population's capacity extends outward — self and children, close relatives, polity, humanity — with decreasing durability as the arc widens. Under demographic pressure, the outer radii previously produced the populations that got killed: conquered, displaced, starved, shipped across oceans in chains, left to die in famines the imperial administration preferred not to relieve (Section 13.3). The same outer radii now produce

GAVI vaccine delivery, SDG commitments, global health funding, and educational aid flowing from already-converged populations toward populations that have not yet crossed. The mechanism is the same four-radii structure; the direction has reversed.

Not all Europeans were in-group to one another in 1700, and the same was true in every prior imperial age (§13.3): in each loading window the circle widened slowly across class, region, nation, and civilisation as literate CT (§4.3) extended coalitional reach. The post-1945 institutional architecture is the first window in which the outermost radius — humanity — became operationally instantiated, imperfectly, by every substrate-bearing population at once. The four-radii structure has not changed; the outermost radius has slowly become reachable.

This is not new altruism, and it is not a Western phenomenon. It is the outer radii of every population that has crossed the threshold extending toward populations that have not — under the same four-radii structure that previously pointed the other way, when the populations at the centre were under pressure the convergence has now released.

14.5 The Peace That Remains

The peace is mostly already here. Asia has not invaded Europe in the seventy years since decolonisation; China and India, the two largest populations on Earth, are completing the transition without forming the demographic-pressure-plus-capacity configuration that powered the European window 1500–1960. The asymmetry-reversal has been running since decolonisation; what I predict is its continuation. Not the peace of virtue (Section 13.1) — the species will remain what it has always been, with its hierarchies and its arbitrary out-group orientations (Section 13.2) intact — but the peace that follows when the Malthusian engine runs out of fuel and the arbitrary othering humans always do loses the demographic pressure that has scaled it into mass violence for most of recorded history. Motive-removal is primary; defensive capacity is the backstop. Peace follows not from the brittle equilibrium of mutual deterrence but from the dissolution of the demographic configuration that made outward projection structurally rational — the nuclear era, capacity without motive-removal on both sides, produced suspended near-conflict across four decades, not peace. What remains is the ordinary evolutionary baseline at a scale the species has not previously experienced: a planet on which most populations have the capacity to feed themselves, to defend themselves, to adapt to climate shocks (Lutz, Muttarak & Striessnig 2014), and to pressure their own governments.

The peace claim and the development claim are therefore the same claim, through the same mechanism: completing the convergence is what dissolves the pressure that produced the four-century window, and the decision that extends education is also the decision that extends peace.

This is the claim the mechanism supports.

15. The Human Cost

My argument specifies a counterfactual; the panel supplies the coefficients that translate it into outcomes. Here I compute the implication. The counterfactual is the singular-priority pace demonstrated in Chapter 7; the coefficients translate education into both under-five mortality and fertility; what I report is children at stake. The mortality channel counts children dying before agency transfer because the mothers raising them carried unloaded childhoods forward. The fertility channel counts daughters not born to mothers whose loaded childhoods gave them the choice.

Life expectancy at birth is not a separable third channel. Once under-five mortality is controlled, education's effect on LE-at-birth collapses to zero ($\beta = 0.013$ years per percentage point, $p = 0.50$, $n = 483$). The LE-at-birth gain a developed country reports is the U5MR gain mechanically lifting the mean.

But adult mortality, measured directly as ${}_{45}q_{15}$, tells a different story. A one-percentage-point rise in female lower-secondary completion at T predicts a 1.25% reduction in female adult mortality at $T{+}28$; controlling for under-five mortality at $T{+}28$, the effect attenuates but remains significant (-0.37% per pp, $p = 0.001$, $n = 483$). The mother whose own childhood was loaded does live longer at older ages, beyond what her country's current child-mortality level captures. The margin is smaller than the children-not-dying count by roughly an order of magnitude — I do not convert it into a separate counterfactual-lives number here — but the channel is real, and Our World in Data's observation that life expectancy is rising at every age (Roser et al. 2013) shows up inside the panel as the within-country signature of this same channel.

Method. The mortality channel is the cross-country relationship $\log(\text{U5MR}) = \alpha + \beta \cdot \text{lsec}_F$, fitted on pre-2000 cross-sections of developing countries (USSR-excluded), giving $\alpha = 5.56$ and $\beta = -0.028$ per percentage point of female lower-secondary completion, applied at the childrearing window (the schooled cohort's children are at under-five risk about twelve years on). The fertility channel is $\beta_{\text{TFR}} = -0.058$ births per percentage point, applied at the cohort's reproductive peak about five years on, estimated on the same panel. Both channels operate together: a counterfactual education trajectory implies a counterfactual U5MR (via the log curve) and a counterfactual TFR (via the linear panel coefficient), which jointly determine the counterfactual deaths. Lives saved decompose into a mortality channel (educated mothers' children survive at higher rates) and a fertility channel (educated mothers have fewer children, so fewer children are at risk of dying in the first place).

The conservative anchor is Korea's empirical 1955–1985 expansion rate of 2.13 percentage points per year on lower-secondary completion (Chapter 8). Two upper-bound scenarios bracket the result: a 15-year linear ramp from each country's baseline to 95% completion, and a 9-year ramp matching the biological floor implied by the dependency-window argument (Chapter 2).

Anchor. The primary counterfactual starts at $T = 1990$. By that year the demonstration of singular priority had been visible for three decades: Korea, Taiwan, Singapore, and Cuba had each crossed both the TFR and life-expectancy thresholds; Korea's lower-secondary completion stood at 95%; the policy template was published, the data were available, and no developing country could plead informational uncertainty. $T = 1990$ is the year from which I count laggards as laggards. I report $T = 1980$ and $T = 1970$ as alternates — earlier anchors with longer windows, but each less defensible as the date past which there was no excuse.

The counterfactual replaces each country's actual lower-secondary trajectory with the scenario trajectory from year T forward. Effects on under-five mortality materialize at the childrearing window, about twelve years later, as the schooled cohort raises children; fertility effects materialize sooner, at the reproductive peak; the window therefore runs from $T + 5$ (the earliest channel) through 2025, with no projection beyond the panel's observation horizon. The USSR republics are excluded (Chapter 10), and the sample is restricted to developing countries, giving 100 countries with complete data at each anchor.

What I report is *Korea pace versus actual* — the additional lives that would have been saved had every country adopted Korea's empirical pace from T forward, beyond what actual education in fact delivered. This is the strict cost-of-laggards counterfactual.

Result. At the primary anchor $T = 1990$, Korea pace yields 104 million lives at stake through 2025: 38 million under-five deaths averted via the mortality channel, plus 67 million fertility-channel lives. The fertility channel now dominates: preventing a birth removes a child from mortality risk entirely, so as the fertility effect strengthens it draws lives out of the mortality count and into its own — the two channels partition the same children, never double-counting them. Earlier anchors yield correspondingly larger totals: $T = 1980$ gives 188 million (66 million mortality, 122 million fertility); $T = 1970$ gives 284 million (101 million mortality, 183 million fertility). The sensitivity envelope at $T = 1990$ is shown in Table 15: each successive scenario relaxes the pace constraint, raising the lives at stake from 104 million (Korea pace) to 195 million (9-year biological floor).

Scenario at $T = 1990$	Mortality	Fertility	Kids not born	Lives at stake
Korea pace (2.13 pp/yr)	38 M	67 M	635 M	104 M
15-year linear ramp to 95%	45 M	127 M	1068 M	172 M
9-year biological floor	51 M	144 M	1199 M	195 M

Table 15: Sensitivity envelope at the primary anchor $T = 1990$, end-year 2025, Korea-pace-vs-actual baseline. Mortality and fertility channels in millions of lives; kids-not-born is the cumulative reduction in births over the window. Each scenario sets a different pace from each country's 1990 lower-secondary baseline forward.

Caveats. The cross-country log curve is fitted on pre-2000 data; the counterfactual assumes the empirical relationship between education and U5MR holds in the counterfactual world. Only the smaller mortality channel rests on this cross-country curve; the now-dominant fertility channel uses the within-country panel coefficient ($\beta_{\text{TFR}} = -0.058$) that carries Chapter 9, so the headline is a composite read off both, not a single cross-country extrapolation. Actual populations are used; counterfactual populations would be smaller because the fertility transition would have arrived earlier, biasing the totals upward. The end-year-2025 horizon still truncates the predicted effect: education accelerated from $T = 1990$ keeps lowering mortality and fertility past 2025 as later cohorts reach the childrearing window, so the headline remains a floor at the panel's observation horizon. The 9-year scenario reflects the biological floor on cohort traversal, not a policy expectation. GDP is not modelled — it is downstream of education at societal scale and adds nothing once education is in (Section 9.5).

What this number means for policy I take up in Chapter 16.

16. The Decision

Every country that developed made the same decision: sustained investment in education. The variable is speed.

Singular priority is the policy logic that demographic metabolism imposes (Section 5.5). Every other lever — institutions, markets, fiscal rules, regulation — is reversible at policy speed. Education runs at metabolic speed: a country shifts its educational mix only as fast as new cohorts displace older ones. Education is the only intervention that alters the metabolism's input; the priority regime sets only how fast that input changes — singular priority fastest, competing priority slower, no priority not at all.

Every government on earth now claims to prioritise education (UN 2015, SDG 4; UNESCO GEM 2024). Competing priority is the norm. The decision I call for is singular priority — for education through the full dependency window, not for stopping at nine years. Singapore is the demonstration: when the option is given, three- quarters of the cohort continues into tertiary education.

SDG 4 commits every signatory, by 2030, to "ensure that all girls and boys complete free, equitable and quality primary and secondary education leading to relevant and effective learning outcomes" (UN 2015). That commitment — free, universal, and measured by learning rather than enrolment — is what singular priority asks. SDG 4 asks for twelve years of schooling, primary through secondary, and set its 2030 horizon with that in view — a fifteen-year window, three years past the biological floor (Section 8.1), correctly ambitious when it was adopted. That floor is set by the child, not the ministry: a cohort entering school under a maximum-speed commitment is fully loaded only on the day it finishes, one school-span later. No decision can deliver the dose faster than a child grows through it. So a country that commits this year reaches the target in twelve years, and a country that waits reaches it twelve years after it stops waiting. The target is correct; the only variable is when a country decides — and whether it finishes the window or stops short of it.

The commitment is universal; the apparatus is what each polity has to hand (Chapter 8). Four disciplines hold across every successful case.

Universal enrolment over school quality. Bad schools, untrained teachers, and large class sizes are not valid reasons to keep a child out. Knox 1696, Meiji 1872, Cuba 1961, and Bangladesh in the 1990s loaded the channel under conditions current education ministries would consider unacceptable (§4.5).

Coverage measured, not credentials. Tracking, retention, and learning-outcome measurement are the operational instruments because hollow reporting fails the phenotype test (Chapter 10).

Education protected through contraction. When fiscal pressure forces a cut between education and health, security, or infrastructure, education is the line that does not move — not because the others are abandoned (Cuba kept rural health, Meiji kept the army, Knox kept the churches) but because they run on the substrate education builds.

Commitment, not funding. The successful cases used near- peer teaching at scale — Cuba sent roughly a hundred thousand teenage *brigadistas* to teach adults through 1961, taking literacy from seventy-six to ninety-six per cent. None of these were donor-funded. Aid was not the funding mechanism in any successful case, and my argument does not require it.

This is the state's decision, and it does the one thing a household cannot do for itself: it reaches the first generation — the children whose parents were never schooled, and who therefore cannot load the home niche on their own.

That first reach is the whole of what the state is for (Section 5.4).

After that, the channel runs without instruction. A schooled generation loads its own children's childhoods as a matter of course; the loaded cohort becomes the next generation's home niche, and the ratchet clicks forward on its own (Chapter 5). No one decides this and no one could stop it — it is the biology of the dependency window, running at the metabolism's own speed, one cohort displacing the last.

But the leader is not the only one who decides. The convergence is built from private decisions that no one coordinates (Section 5.4); what is unwilled can also be willed. Every person already schooled can hold the discipline the successful states ran — education as the one cause, ahead of all the others, singular and not dispersed. The parent who loads her own child does this without deciding. The decision that remains open is for the children who are not hers: the ones the state has not reached, whose home niche carries no reading. To teach one of them, to build the school that holds them, to make their schooling the cause that does not move — that decision is available to every educated person, not only the one who governs.

I have written this for leaders. I have written it equally for everyone the leaders' schools already reached. The leader sets the pace for a nation's children; you decide what becomes of the child the pace has not yet reached.

One species, one biological channel of cultural transmission, one developmental form — across every political system and every continent. The cohort whose adulthood will produce the convergence is in school this year — or it is not. The window admits exactly what is loaded into it, and it closes one cohort at a time.

No one has decided on evidence alone. I provide the evidence.

What remains is the decision.

Acknowledgments

Melvin Konner's *The Evolution of Childhood* is the biological foundation I build on. The eighteen years that separate the human juvenile period from every other species is the fact from which everything else follows — this paper is what I take from his book. *The Tangled Wing* has been a longer, quieter presence.

Richard Wrangham established what the human species was before education reached it — what cooking made, what violence shaped, what self-domestication permitted. The brain the long childhood loads is the brain his cooked diet funded.

Sarah Blaffer Hrdy established that the household which transmits is never just two parents. The cooperative breeding that made the eighteen-year dependency period viable is the precondition I inherit.

Wolfgang Lutz's datasets made this analysis possible and his work on education and human capital set the terms on which it stands. He will recognise the question I am answering as his teacher's.

The CARTA symposia at UCSD provided a year's immersion in the scientific conversations underlying the biological arguments in this paper — evolutionary biology heard as a living community rather than a finished literature.

Frans de Waal spent his life showing what primates reveal about what is possible — the bonobo's imperfect, conflicted cooperation alongside the chimpanzee's lethal hierarchy. Without the first the household does not hold; without understanding the second we cannot see what education overcomes. He is missed.

Hans Rosling spent his life explaining the world as it is and showing us its progress. Patiently, talk after talk, country by country, he pulled apart the "developing vs developed" boxes and showed that we are all modern people. I owe that piece of the framing to him. He is missed.

Hunger and Public Action was one of the most important books I have read. The ten years I spent working to secure its Creative Commons release were a privilege; Jean Drèze was consistently supportive throughout. The views in this paper are my own. Drèze and Sen saw the right anomalies and asked the right questions. This paper is one answer.

The theoretical framework, analytical design, test specifications, arguments, conclusions, and all errors are my own. Implementation used AI assistance: I directed Claude (Anthropic) to write code and statistical analysis against

my specifications and architectural direction, and used Claude to assist prose drafting under my editorial direction.

A. Reproducibility

Every number, table, and figure in this paper is produced by a named script in the GitHub repository:

https://github.com/rkpagadala/the-long-childhood

 The repository is forkable and citable without further permission. The script index is `SCRIPTS.md`; each paper number is registered against its producing script in `scripts/verify_the_long_childhood.py` (run `make verify` to confirm coverage).

The panel headline ($\beta = 0.707$, lower-secondary parent $\rightarrow$ child completion, country fixed effects, expansion sample, 945 country-years, 144 countries, 1975–2015) is accompanied by a 20-test econometric diagnostic battery in `scripts/econometric_battery/`; the full test inventory, caveats, and outputs are in `scripts/econometric_battery/REPORT.md` and indexed by `scripts/ECONOMETRICS.md`.

β stays positive in every cell of every test. Leave-one-out $\beta \in [0.731, 0.774]$; 16 of 16 multiple-comparison tests survive at 5% after every correction in the battery; placebo nulls sit 25+ standard deviations below the real β.

B. Data

Table B1: Summary statistics for analysis variables.

Variable	n	mean	sd	min	max
Parental lower-sec completion (%, $T-28$)	1,480	41.6	32.9	0.0	100.0
Child lower-sec completion (%, T)	1,480	63.6	31.0	1.5	100.0
Log GDP per capita (const. 2015 USD)	1,338	8.30	1.48	5.14	11.67
Life expectancy at birth (years)	1,430	66.4	10.3	28.6	84.3

Variable	n	mean	sd	min	max
Total fertility rate (births/woman)	1,430	3.56	1.90	0.91	8.86
Under-5 mortality (per 1,000 live births)	1,392	62.5	64.9	2.2	334.0

Notes: Analysis panel: 185 countries, 1975–2015 at 5-year intervals, with parental education lagged 28 years (observed at $T-28$). The education+GDP panel covers 178 of 185 countries; seven are dropped because the World Bank WDI has no GDP series for them: North Korea, Taiwan, Micronesia (FSM), and four French overseas departments (French Guiana, Guadeloupe, Martinique, Reunion).

References

Acemoglu, D., Johnson, S. & Robinson, J.A. (2001). The Colonial Origins of Comparative Development: An Empirical Investigation. *American Economic Review*, 91(5), 1369–1401.

Acemoglu, D. & Robinson, J.A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown Publishers.

Aghion, P. & Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2), 323–351.

Anderson, B.A. & Silver, B.D. (1986). Infant Mortality in the Soviet Union: Regional Differences and Measurement Issues. *Population and Development Review*, 12(4), 705–738.

Angrist, N., Djankov, S., Goldberg, P.K. & Patrinos, H.A. (2021). Measuring Human Capital Using Global Learning Data. *Nature*, 592, 403–408.

Auvert, B., Taljaard, D., Lagarde, E., Sobngwi-Tambekou, J., Sitta, R. & Puren, A. (2005). Randomized, Controlled Intervention Trial of Male Circumcision for Reduction of HIV Infection Risk: The ANRS 1265 Trial. *PLOS Medicine*, 2(11), e298.

Babiarz, K.S., Ma, P., Miller, G. & Song, S. (2018). The Limits and Consequences of Population Policy: Evidence from China's Wan Xi Shao Campaign. NBER Working Paper No. 25130.

Baker, D.P., Collins, J.M. & Leon, J. (2008). Risk Factor or Social Vaccine? The Historical Progression of the Role of Education in HIV and AIDS Infection in Sub-Saharan Africa. *Prospects*, 38(4), 467–486.

Becker, G.S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. National Bureau of Economic Research; Columbia University Press.

Bora, J.K., Saikia, N., Kebede, E.B. & Lutz, W. (2022). Revisiting the causes of fertility decline in Bangladesh: the relative importance of female education and family planning programs. *Asian Population Studies*, 19(1), 81–104.

Boyd, R. & Richerson, P.J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press.

Brass, P.R. (1986). The Political Uses of Crisis: The Bihar Famine of 1966–1967. *Journal of Asian Studies*, 45(2), 245–267.

Cai, Y. (2010). China's Below-Replacement Fertility: Government Policy or Socioeconomic Development? *Population and Development Review*, 36(3), 419–440.

Caldwell, J.C. (1979). Education as a Factor in Mortality Decline: An Examination of Nigerian Data. *Population Studies*, 33(3), 395–413.

Cleland, J.G. & Van Ginneken, J.K. (1988). Maternal education and child survival in developing countries: the search for pathways of influence. *Social Science & Medicine*, 27(12), 1357–1368.

Davis, M. (2001). *Late Victorian Holocausts: El Niño Famines and the Making of the Third World*. Verso.

De Neve, J.-W., Fink, G., Subramanian, S.V., Moyo, S. & Bor, J. (2015). Length of Secondary Schooling and Risk of HIV Infection in Botswana: Evidence from a Natural Experiment. *Lancet Global Health*, 3(8), e470–e477.

de Waal, F.B.M. (2013). *The Bonobo and the Atheist: In Search of Humanism Among the Primates*. W.W. Norton.

De Walque, D. (2007). How Does the Impact of an HIV/AIDS Information Campaign Vary with Educational Attainment? Evidence from Rural Uganda. *Journal of Development Economics*, 84(2), 686–714.

Deaton, A. (2013). *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton University Press.

Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press.

Dehaene, S. (2009). *Reading in the Brain: The New Science of How We Read*. Viking.

Dehaene, S., Pegado, F., Braga, L.W., Ventura, P., Nunes Filho, G., Jobert, A., Dehaene-Lambertz, G., Kolinsky, R., Morais, J. & Cohen, L. (2010). How Learning to Read Changes the Cortical Networks for Vision and Language. *Science*, 330(6009), 1359–1364.

Diamond, J. (1997). *Guns, Germs, and Steel: The Fates of Human Societies*. W.W. Norton.

Drèze, J. & Sen, A. (1989). *Hunger and Public Action*. Clarendon Press.

Drèze, J. & Murthi, M. (2001). Fertility, Education, and Development: Evidence from India. *Population and Development Review*, 27(1), 33–63.

Duflo, E. (2012). Women Empowerment and Economic Development. *Journal of Economic Literature*, 50(4), 1051–1079.

Dyson, T. & Maharatna, A. (1992). Bihar Famine, 1966–67, and Maharashtra Drought, 1970–73: The Demographic Consequences. *Economic and Political Weekly*, 27(26), 1325–1332.

Easterlin, R.A. (1981). Why Isn't the Whole World Developed? *Journal of Economic History*, 41(1), 1–19.

Eveleth, P.B. & Tanner, J.M. (1990). *Worldwide Variation in Human Growth* (2nd ed.). Cambridge University Press.

Fortson, J.G. (2008). The Gradient in Sub-Saharan Africa: Socioeconomic Status and HIV/AIDS. *Demography*, 45(2), 303–322.

Frank, M.C., Everett, D.L., Fedorenko, E. & Gibson, E. (2008). Number as a cognitive technology: Evidence from Pirahã language and cognition. *Cognition*, 108(3), 819–824.

Gakidou, E., Cowling, K., Lozano, R. & Murray, C.J.L. (2010). Increased Educational Attainment and Its Effect on Child Mortality in 175 Countries between 1970 and 2009: A Systematic Analysis. *The Lancet*, 376(9745), 959–974.

Galor, O. (2011). *Unified Growth Theory*. Princeton University Press.

Gao, M. (1999). *Gao Village: Rural Life in Modern China*. University of Hawai'i Press.

Gao, M. (2008). *The Battle for China's Past: Mao and the Cultural Revolution*. Pluto Press.

Goldin, C. & Katz, L.F. (2008). *The Race Between Education and Technology*. Harvard University Press.

Goodall, J. (1986). *The Chimpanzees of Gombe: Patterns of Behavior*. Belknap Press of Harvard University Press.

Goody, J. (1976). *Production and Reproduction: A Comparative Study of the Domestic Domain*. Cambridge University Press.

Gordon, P. (2004). Numerical Cognition Without Words: Evidence from Amazonia. *Science*, 306(5695), 496–499.

Grantham-McGregor, S., Cheung, Y.B., Cueto, S., Glewwe, P., Richter, L., Strupp, B. & the International Child Development Steering Group. (2007). Developmental Potential in the First 5 Years for Children in Developing Countries. *The Lancet*, 369(9555), 60–70.

Gwatkin, D.R. (1979). Political Will and Family Planning: The Implications of India's Emergency Experience. *Population and Development Review*, 5(1), 29–59.

Hajnal, J. (1965). European Marriage Patterns in Perspective. In D.V. Glass & D.E.C. Eversley (Eds.), *Population in History: Essays in Historical Demography* (pp. 101–143). Edward Arnold.

Hamilton, W.D. (1964). The Genetical Evolution of Social Behaviour, I & II. *Journal of Theoretical Biology*, 7(1), 1–52.

Hanushek, E.A. & Woessmann, L. (2008). The Role of Cognitive Skills in Economic Development. *Journal of Economic Literature*, 46(3), 607–668.

Hanushek, E.A. & Woessmann, L. (2012). Do Better Schools Lead to More Growth? Cognitive Skills, Economic Outcomes, and Causation. *Journal of Economic Growth*, 17(4), 267–321.

Hare, B. & Woods, V. (2020). *Survival of the Friendliest: Understanding Our Origins and Rediscovering Our Common Humanity*. Random House.

Hawkes, K. (2003). Grandmothers and the evolution of human longevity. *American Journal of Human Biology*, 15(3), 380–400.

Heckman, J.J. (2006). Skill Formation and the Economics of Investing in Disadvantaged Children. *Science*, 312(5782), 1900–1902.

Heckman, J.J. & Cunha, F. (2007). The Technology of Skill Formation. *American Economic Review*, 97(2), 31–47.

Heckman, J.J. & Mosso, S. (2014). The Economics of Human Development and Social Mobility. *Annual Review of Economics*, 6, 689–733.

Herculano-Houzel, S. (2016). *The Human Advantage: A New Understanding of How Our Brain Became Remarkable*. MIT Press.

Heyes, C. (2018). *Cognitive Gadgets: The Cultural Evolution of Thinking*. Harvard University Press.

Hill, K. & Hurtado, A.M. (1996). *Aché Life History: The Ecology and Demography of a Foraging People*. Aldine de Gruyter.

Hrdy, S.B. (2009). *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Harvard University Press.

Jensen, R. (2010). The (Perceived) Returns to Education and the Demand for Schooling. *Quarterly Journal of Economics*, 125(2), 515–548.

Kaplan, H., Hill, K., Lancaster, J. & Hurtado, A.M. (2000). A Theory of Human Life History Evolution: Diet, Intelligence, and Longevity. *Evolutionary Anthropology*, 9(4), 156–185.

Konner, M. (2010). *The Evolution of Childhood: Relationships, Emotion, Mind*. Harvard University Press.

Kramer, K.L. (2011). The Evolution of Human Parental Care and Recruitment of Juvenile Help. *Trends in Ecology & Evolution*, 26(10), 533–540.

Lee, J.W. & Lee, H. (2016). Human Capital in the Long Run. *Journal of Development Economics*, 122, 147–169.

Liu, Y., Rao, K. & Hsiao, W.C. (2003). Medical Expenditure and Rural Impoverishment in China. *Journal of Health, Population and Nutrition*, 21(3), 216–222.

Lutz, W. (2013). Demographic Metabolism: A Predictive Theory of Socioeconomic Change. *Population and Development Review*, 38(s1), 283–301.

Lutz, W., Muttarak, R. & Striessnig, E. (2014). Universal Education is Key to Enhanced Climate Adaptation. *Science*, 346(6213), 1061–1062.

Lutz, W. & Kebede, E. (2018). Education and Health: Redrawing the Preston Curve. *Population and Development Review*, 44(2), 343–361.

Lutz, W., Reiter, C., Özdemir, C., Yildiz, D., Guimaraes, R. & Goujon, A. (2021). Skills-adjusted Human Capital Shows Rising Global Gap. *Proceedings of the National Academy of Sciences*, 118(7), e2015826118.

Malthus, T.R. (1798). *An Essay on the Principle of Population*. London: J. Johnson.

Mankiw, N.G., D. Romer, & Weil, D.N. (1992). A Contribution to the Empirics of Economic Growth. *Quarterly Journal of Economics*, 107(2), 407–437.

Marlowe, F.W. (2010). *The Hadza: Hunter-Gatherers of Tanzania*. University of California Press.

Mokyr, J. (1983). *Why Ireland Starved: A Quantitative and Analytical History of the Irish Economy, 1800–1850*. George Allen & Unwin.

Mokyr, J. (1990). *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.

Mokyr, J. (2002). *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton University Press.

Mokyr, J. (2009). *The Enlightened Economy: An Economic History of Britain, 1700–1850*. Yale University Press.

Mokyr, J. (2016). *A Culture of Growth: The Origins of the Modern Economy*. Princeton University Press.

Morris, I. (2010). *Why the West Rules—For Now: The Patterns of History, and*

What They Reveal About the Future. Farrar, Straus and Giroux.

Mullainathan, S. & Shafir, E. (2013). *Scarcity: Why Having Too Little Means So Much*. Times Books.

Myhrvold-Hanssen, T. (2003). Democracy, News Media, and Famine Prevention: Amartya Sen and the Bihar Famine of 1966–67. Masters thesis, University of Oslo.

National Bureau of Statistics of China. (1982). *Third National Population Census of the People's Republic of China*. China Statistics Press. [Rural population share: 79.7% in 1980.]

Ó Gráda, C. (1995). *Ireland: A New Economic History 1780–1939*. Oxford University Press.

Ó Gráda, C. (1999). *Black '47 and Beyond: The Great Irish Famine in History, Economy, and Memory*. Princeton University Press.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.

Pepper, S. (1996). *Radicalism and Education Reform in 20th-Century China*. Cambridge University Press.

Pinker, S. (2018). *Enlightenment Now: The Case for Reason, Science, Humanism, and Progress*. Viking.

Prieto, A. (1981). Cuba's National Literacy Campaign. *Journal of Reading*, 25(3), 215–221. [268,000 brigadistas mobilised for the 1961 campaign.]

Pritchett, L. (2013). *The Rebirth of Education: Schooling Ain't Learning*. Center for Global Development.

Romer, P.M. (1986). Increasing Returns and Long-Run Growth. *Journal of Political Economy*, 94(5), 1002–1037.

Romer, P.M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5), S71–S102.

Roser, M., Ortiz-Ospina, E. & Ritchie, H. (2013, ongoing). Life Expectancy. *Our World in Data*. https://ourworldindata.org/life-expectancy.

Rosling, H., Rosling, O. & Rosling Rönnlund, A. (2018). *Factfulness: Ten Reasons We're Wrong About the World — and Why Things Are Better Than You Think*. Flatiron Books.

Schultz, T.W. (1961). Investment in Human Capital. *American Economic Review*, 51(1), 1–17.

Schultz, T.P. (2002). Why Governments Should Invest More to Educate Girls. *World Development*, 30(2), 207–225.

Sen, A. (1999). *Development as Freedom*. Oxford University Press.

Smil, V. (2017). *Energy and Civilization: A History*. MIT Press.

Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*.

Solow, R.M. (1956). A Contribution to the Theory of Economic Growth. *Quarterly Journal of Economics*, 70(1), 65–94.

Stearns, S.C. (1992). *The Evolution of Life Histories*. Oxford University Press.

Tainter, J.A. (1988). *The Collapse of Complex Societies*. Cambridge University Press.

Tomasello, M. (2014). *A Natural History of Human Thinking*. Harvard University Press.

Tomasello, M. (2016). *A Natural History of Human Morality*. Harvard University Press.

Tooby, J. & Cosmides, L. (2010). Groups in Mind: The Coalitional Roots of War and Morality. In H. Högh-Olesen (Ed.), *Human Morality and Sociality: Evolutionary and Comparative Perspectives* (pp. 191–234). Palgrave Macmillan.

UNDP. (1990). *Human Development Report 1990*. Oxford University Press.

UNESCO. (2024). *Global Education Monitoring Report 2024/5: Leadership in Education — Lead for Learning*. Paris: UNESCO.

Unger, J. (1982). *Education Under Mao: Class and Competition in Canton Schools, 1960–1980*. Columbia University Press.

United Nations. (2015). *Transforming Our World: The 2030 Agenda for Sustainable Development* (A/RES/70/1). New York: United Nations.

Vicziany, M. (1982). Coercion in a Soft State: The Family-Planning Program of India. Part I: The Myth of Voluntarism. *Pacific Affairs*, 55(3), 373–402.

Walker, R., Gurven, M., Hill, K., Migliano, A., Chagnon, N., De Souza, R., Djurovic, G., Hames, R., Hurtado, A.M., Kaplan, H., Kramer, K., Oliver, W.J., Valeggia, C. & Yamauchi, T. (2006). Growth Rates and Life Histories in Twenty-Two Small-Scale Societies. *American Journal of Human Biology*, 18(3), 295–311.

Wallace, A.R. (1870). *Contributions to the Theory of Natural Selection: A Series of Essays*. Macmillan.

World Bank. World Development Indicators. Various years.

Wrangham, R. & Peterson, D. (1996). *Demonic Males: Apes and the Origins of Human Violence*. Houghton Mifflin.

Wrangham, R. (2009). *Catching Fire: How Cooking Made Us Human*. Basic Books.

Wrigley, E.A. (2010). *Energy and the English Industrial Revolution*. Cambridge University Press.