The Long Childhood

On the Convergence of Humanity

Krishna Pagadala[1]

May 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 docu-

[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); not an economist.
https://educationfirst.world
Replication: https://github.com/rkpagadala/the-long-childhood

ment the inverse — what one species has done with its window, post-1960, 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 (~25 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 developed country is one in which enough households have crossed this parental threshold. The 185-country panel (1950–2015), spanning four generational depths — one century — bears the mechanism out. A great-great-grandparent's school completion predicts their great-great-grandchild's child-survival outcomes a century later. Seven empirical signatures are deduced from the biology — generational timing, multi-generational persistence, asymmetric disruption, income independence, universality, amplification at low baselines, collective action — and tested against the panel; all seven hold. The identification is biological and natural-experimental; the cross-country panel is the test of predicted signatures, not the source of the causal claim.

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; vertical cultural transmission; developmental niche construction; convergence; fitness components; mass schooling; home niche; cultural-transmission niche; human development; demographic transition

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 11–12 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 14 diagnoses why the human-development apparatus has under-weighted its own central variable. Sections 8 (the country trajectories) and 13 (engagement with Smith, Sen, Deaton, and the institutional alternative) are the conversation with existing scholarship. 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 seven 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. The Robustness appendix (Section A) is for the reader who must satisfy a technical audience before acting. Sections 11–12 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.

The sections not on a path — the Glossary at the front, the Robustness appendix at the back — are reference. They are read when a main path requires them, not in sequence.

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, but it is not the identification.

This paper is written for readers who want to evaluate the evidence directly. I write as a software engineer of twenty-five years' practice, self-taught in the substantive literatures over the same period — not as an economist, not from any discipline this paper crosses. No specialised training is required to read this. Every number is traceable to a script that produces it from source data; the verifier (`scripts/verify_humanity.py`) makes the empirical argument equally checkable for the reader and the author.

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 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-five-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, 185 countries (Table 11); a long-run 1900–2015 panel on a 28-country subsample with self-determined education policy (Table A10) shows the same pattern at deeper baselines, where the post-1975 panel is weighted toward countries already approaching the ceiling. Per-country 25-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

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 ascent is not gradual. 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 century, 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 at the population scale is set by demographic metabolism (Lutz 2013): a country shifts its educational mix only as fast as new, more-schooled cohorts displace older, less-schooled ones. 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 the 185-country 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 11–12 extend my argument forward. Chapter 13 takes up the frameworks the biological claim must displace; Chapter 14 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 child-rearing 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). No input can outrank the summary of itself. No consequence can 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 future GDP per capita as strongly as it predicts LE and TFR (Table 13). Crossing dates under three threshold specifications shift by 10–30 years but the ordering holds (Table A11). Countries with high expansion rates cross any threshold within 14–16 years. The definition chosen here is the conservative one — it strengthens the alternative's position by removing from the outcome the one variable where income might have independent standing.

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.

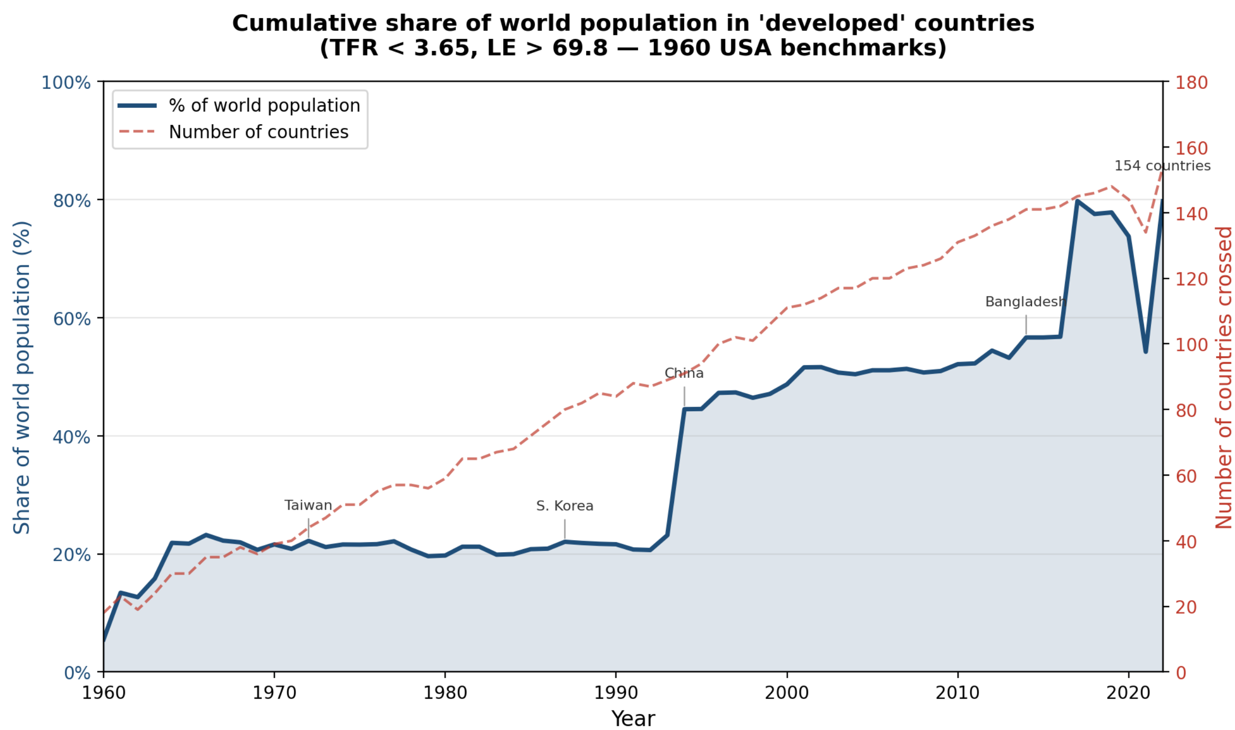

Figure 1: Cumulative share of world population crossing both development thresholds (TFR < 3.65, LE > 69.8), 1960–2022. Slow climb from 13% (1961) to ~20% (1993); vertical jump in 1994 when China crosses; 50% by 2001; 80% by 2017–2022. *(Source: World Bank WDI, WCDE v3.)*

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 the binding constraint and asked why it had not reached the whole world. Lutz and Kebede (2018) showed that the classic Preston curve — the apparent relationship between national income and life expectancy — 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 capture is not a separate precondition under cultural transmission; it is cultural transmission's oldest output, from fire forward. The 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 seven 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 fifty days. 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.

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. 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.

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 (~99%), 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.

2.5 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 (Ha-

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

jnal 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 the CT regime is locked at fifteen; later schooling and adult experience continue to shape it. 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 juvenile has around seven years of plastic learning embedded with an adult before independence. Seven years is long enough to absorb termite fishing, social hierarchy, and the foraging repertoire of the group. It is 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 — eighteen years, 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 water-hole 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 (Tomasello 2014): each generation of chimps rediscovers the termite stick; each generation of humans inherits whatever the prior generation learned. Fire is the seed case — caught once, never relost, and the condition for cooking, the expensive brain, and every subsequent layer of transmission. The next subsection traces the ratchet.

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 first. Cooking raised the net caloric yield of food enough to fund the metabolically expensive human brain (Wrangham 2009), which funded more transmission, which funded more tools. 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 why schooling specifically — among all the things that have loaded through the window — is the highest-bandwidth payload ever delivered.

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. Formal schooling is the latest, and by a considerable margin the highest-bandwidth, layer ever delivered. I ask why — what schooling does in the window that other transmission channels 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 10.7 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.

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.

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 primary completion at or above 66% and lower-secondary completion at or above 42% in the 20–24 cohort; the medians are 86% and 65% respectively ($n = 84$ after excluding the USSR, Europe, Cambodia, and countries already above threshold by 1960). For the fertility threshold (TFR < 3.65) the primary floors are looser (p10 $= 57\%$, median 79%; $n = 88$), consistent with TFR responding to primary completion alone while LE needs the deeper lower-secondary signal. The grandparent cohort, two generations earlier, sat at a median of 24% primary completion at LE crossings and 17% at TFR crossings, with a tenth percentile near 2% — so 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 — the outputs literate CT produces that no illiterate CT, however dense or long-running, produces — 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.

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.

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 makes the societal claim mechanistically coherent.

4.5 Duration Over Fidelity

What schooling installs is a cognitive reorganisation, not content. 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 software engineer 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 10–25 years earlier, the parental lag the mechanism demands. Tests are not an alternative to years; they are compounded years measured one generation downstream. Section 9.4 develops the partition; Chapter 10, §10.7 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 made possible by cooking, which Wrangham (2009) dates to roughly two million years ago. Once cooking crossed the threshold that made the energetically expensive human brain metabolically possible, the dependency window did the rest. 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 ~18-year juvenile 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 ~25 years. Section 5.4 develops the regimes; the cases (Chapter 8) trace them.

I take the ~25-year lag that runs through this paper to be biological turnover time: the interval for one full generation of children to pass through the school niche and become the near-adults of the next cohort's home niche. It is not the interval between policy and response. Policy is the click of the ratchet; the lag 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 9.8) operationalises it by outcome: primary completion moves TFR because basic literacy makes family size a choosable parameter; lower and upper secondary move life expectancy and under-five mortality (U5MR) because the deeper cognitive reorganisation lets a mother evaluate symptoms, weigh distant risks, and act on them.

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 9.4); 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 9.6). 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 built a singular-priority school niche alongside the home niche, accelerating a ratchet 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.

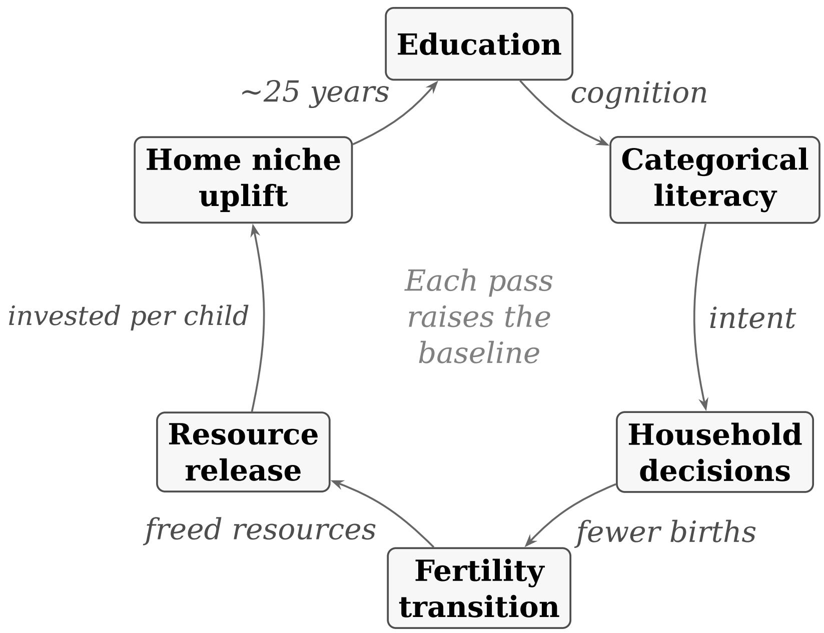

Figure 2: The generational transmission loop. Absorption from the school niche installs categorical literacy in the individual brain; at population scale, the same absorption sustains literate CT as the society's transmission regime. Categorical literacy shapes household decisions; household decisions 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 ~25-year cycle clicks the ratchet one notch.

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

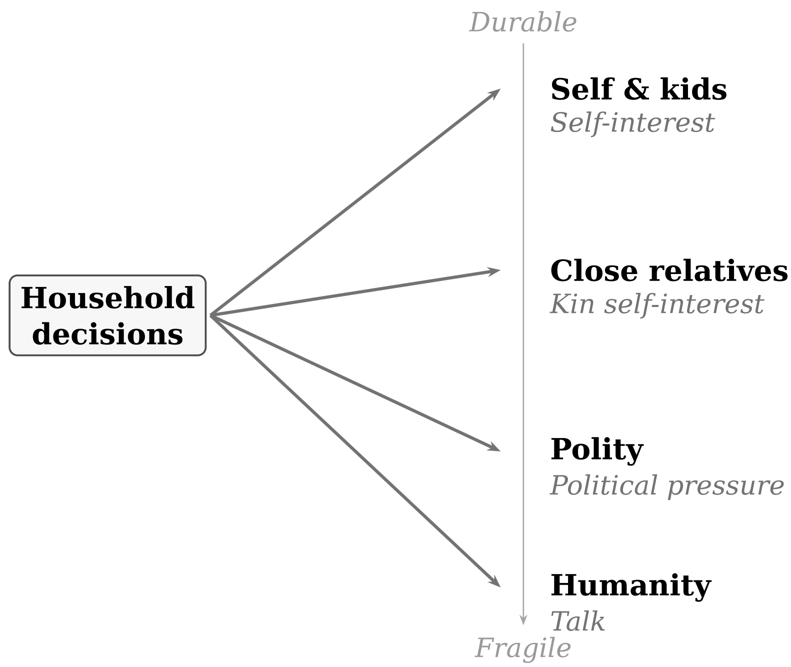

Figure 3: Four radii of educational effect — each a boundary draw of the same coalitional in-group mechanism, with durability decreasing as the boundary widens: action at the kin scale, political pressure at the polity scale, talk at the humanity scale.

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 — steepest at the bottom, 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 at primary education (Section 9.8), where women are barely literate. What primary education gives 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. 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.

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. The species evolved long post-reproductive lives because it evolved extended juvenile dependency because it evolved cultural transmission. 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. Their redrawing of the income–health relationship turns income-axis diagrams into education-axis diagrams, and the fit improves. 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.

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 (Mamun & Bongaarts 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.

Singular priority (Section 16) is the only intervention that alters the metabolism's input. 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 13.4 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 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 knee. Primary alone moves fertility, but not life expectancy or under-five mortality at the required scale; lower-secondary is the empirical floor for full convergence, with deeper exposure compounding the signal (Lutz & Kebede 2018). §4.3 reports the threshold-crossing record.

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 seven 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 juvenile dependency window 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 $T + {\sim}25$. The prediction is not 10 or 15 years; those are the intervals that direct income transfer or institutional change would operate through.

2. **Multi-generational persistence.** Because Parental Transmission (Chapter 5) runs through the dependency window in every generation, 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 — the loading is through the dependency window, not through the institution that delivered it (Sections 5.1 and 5.4) — destroying the educated population should reset generational progress, while destroying institutions, income, or infrastructure while leaving the educated population intact should not.

4. **GDP independence.** Because the causal path runs education $\rightarrow$ cognitive capacity (Chapter 4) $\rightarrow$ household decisions (Chapter 5) $\rightarrow$ outcomes, log GDP per capita stripped of education's contribution should have no independent predictive power for development outcomes: it sits downstream of the same mechanism that produced it. Operationalised

as residualization against education (Frisch-Waugh-Lovell), and tested in Section 9.8.

5. **Universality.** Because extended juvenile dependency and cumulative cultural transmission are species-level traits (Chapters 2–3), not institutional products, the pattern should hold across political systems, cultures, geographies, and colonial histories.

6. **Amplification at low baselines.** Because state reach brings first-generation learners into the school niche on top of what the home niche already carries (Section 5.4), each generation's attainment should exceed the parental baseline by the most where that baseline is lowest; at high baselines the gap should close as ceiling compression sets in. Formally: the generational amplification ratio β_g (defined Section 5.1) exceeds 1 at low baselines and approaches 0 at high ones.

7. **Collective action.** Household cognitive capacity does not stay in the household: it aggregates into collective-action capacity — the Level-2 categorical claim (Section 4.3). 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; 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 11). 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 14) — 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 12). 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 tests them universally across 185 countries.

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 country meeting both conditions will cross both demographic thresholds within thirty years of passing the 35% floor — Korea binds the rate-subset cluster at thirty; Indonesia's three-year overshoot reflects archipelago-scale sequencing. 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 7)

The famine test is the strongest empirical test of the collective-action prediction (Prediction 7). 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. Seventeen countries faced shocks comparable to or worse than those that produced famines elsewhere — Cuba's 35% GDP collapse when Soviet food imports were cut by 80% (1993), post-WWII Japan and Germany, Armenia under simultaneous blockade and war (1993), Ukraine under Russian invasion and grain blockade (2022) — and none experienced famine. Their median education was 71.6%, against 19.6% for famine countries. The difference is significant at $p < 0.00001$ (Mann–Whitney U). 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. The test assumes the energy substrate (Section 2): every shock in the dataset was a distribution failure, not absolute absence of food in the region. Educated populations respond by solving the distribution problem; uneducated populations cannot.

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 near 60%, Catholic Connacht and west Munster under 30%

(Ó Gráda 1995). Mokyr's (1983) county-level analysis found pre-famine literacy negatively associated with excess mortality net of potato dependence; 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%). Twenty-five years from the end of the educational disruption (1985) lands at 2010, matching the time-to-agency lag predicted by the mechanism (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 16) 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, with a coefficient nearly twice the parent's ($\beta_{gp} = –0.059$ vs $\beta_p = –0.033$ 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 2: 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; Barro-Lee v3.0 and WCDE v3 agree at 1990 but WCDE over-reports secondary completion for 1960–1980 cohorts, with phenotype implications developed in Chapter 10). 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 2.41 and primary completion rose from 78% to 91% — the post-apartheid expansion reaching the Black majority at scale. The 1990 lower-secondary headline of 65% was an apartheid aggregate; the population bearing the shock sat below the behavioural- response depth.

The composition matters because primary and lower-secondary depth carry different halves of the response. Primary drove the fertility decline: it was expanding through the shock, and TFR fell on schedule (R^2=0.65 for TFR-primary). Lower-secondary carries the behavioural response to novel pathogens; in the affected majority it remained too shallow. Acquisition-stage factors were genuinely biological — low traditional circumcision prevalence (60% biological reduction per three RCTs) and circular mining migration driving spread. Once information existed, the education–HIV gradient reversed: 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), and Baker et al. (2017) formalise the pattern as the Population Education Transition curve — educated populations change behaviour first once novel-risk information exists. Post-ARV, LE recovered to 66.1 by 2019. Full mechanism detail — apartheid stratification, circumcision prevalence in the HIV belt, and the Population Education Transition curve — is in Appendix A.3.8.

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 both schools and inclusive institutions. 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 Tables A7–A9; see also Appendix A.3.9.

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.

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.

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 T→T+25 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 one generation after it began loading its children's long childhoods with formal schooling — 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 the eighteen-year window: who built the schools, who reached them, what was loaded into the children, and what the resulting cohort then did as adults.

Table 4: 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	1953–65	2.13	~25 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 ∼25-year cycles. Cuba's high starting base (40%) and the 1961 reconstruction of losses from the post-revolution exodus shortened the lag below the typical ∼25-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 4 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 Deng's dismantling of the barefoot doctor network. 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, on lower secondary completion (22% against 18%), and on colonial educational inheritance. By 2000 Korea had crossed every development threshold; the Philippines had crossed none. The starting position 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 (25% and 18% 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 both the 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. Korea's 2.13 pp/yr is the WCDE record — the fastest sustained national expansion observed. From the 1950 starting position of thirty 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 (`scripts/korea_vs_biological_pace.py`). Korea 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, inside a larger country that has not crossed. 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

~1974; 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 cross 69.8 in 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.

Case	Crossed	Prediction tested and result
Cuba	1974	**Prediction 5** (socialist regime) and 6 (high starting base, short lag). 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 girls' secondary expansion, not contraceptive distribution (Mamun & Bongaarts 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 lower secondary gains in the WCDE 1870–2015 record: +10.6pp for the 1975 cohort, +15.0pp for the 1980 cohort. 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.[3] 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.

[3] 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 $\pm$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 $\pm$0.25 and $\pm$1.0 mean years.

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; Miller 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.

9. The Panel

A note to the reader from development economics. The identification of this paper's claim is not in this chapter. It is in Chapters 2–6 (the biological mechanism), Chapter 7 (the natural experiments), and Chapter 10 (the hollow-education falsification). What follows tests whether the population-scale signatures the biology predicts appear in the cross-country panel. The within-country fixed-effects design is insufficient as identification, and is not asked to do that work; it is asked whether the signatures are there. The regressions are the supplement; the natural experiments are the rigorous part. Readers who want to engage the identification should read those chapters first and return to this one as confirmation.

Education predicts every development outcome one generation forward; log GDP per capita, once education's contribution is removed, predicts none of them. These results follow from the biology — the seven predictions of the mechanism in Chapter 6, tested against 185 countries over 65 years. All seven hold.

9.1 Identification

Identification rests on three things: temporal ordering between cause and effect, natural experiments where the channel is severed or held fixed across other variation, and the falsification test in Chapter 10. The within-country regressions that follow show the mechanism's signature in the panel; they are not the identification.

Temporal ordering. Parental cohort precedes child cohort by biological necessity. Outcomes at $T + 25$ are predicted by education at T. Reverse causality from outcome to cause is ruled out by construction at the individual level. Within-country fixed effects absorb everything time-invariant about a country: institutions as of any baseline year, geography, ethno-linguistic composition, religion, colonial history, resource endowment, language. What remains is movement over time, and the 25-year lag is not a researcher choice — it is the time-to-agency interval derived from the dependency window (Prediction 1).

Why "standard controls" are the wrong question. Cross-country regressions in development economics conventionally include controls for institutional quality, geography, religion, colonial history, ethno-linguistic fractionalisation, and resource endowment, on the premise that these are confounders moving both education and outcomes. Each of these is either downstream of the channel or absorbed by the country fixed effect. Institutions are built by educated populations: civil servants, judges, journalists, and administrators are the educated cohort. AJR's own institutional measure (*avexpr*) is itself a strong instrument for education on AJR's 64-country base sample (first-stage $F = 37.13$, three times the Stock & Yogo threshold), so the exclusion restriction underlying institutionalist IV identification is empirically violated (Section 13.4, Tables A7–A9). Religion's developmental signature works through literacy: the Protestant differential is the post-Reformation push to read scripture in the vernacular, not a doctrinal effect on fertility or longevity. Colonial history is operative through where colonisers built schools and how widely; the historical record selects on this variable directly. Geography, ethno-linguistic baselines, and initial resource endowment are time-invariant within a country and are absorbed by the country fixed effect. The remaining time-varying institutional change is itself driven by education: Polity5 polity2 has effectively no power to predict education gain rates ($R^2 < 0.01$ across all measurement lags; Section 13.4, Table A5). Time-varying institutional quality enters as an appendix robustness check and does not move the headline.

The standard-controls toolkit was built to isolate the marginal effect of a marginal regressor. The educational transition is not marginal; it is the loading of the channel that those controls are downstream of. Adding them is measuring shadows of the channel and calling them confounders. Identification in this section is therefore not by control adjustment but by the design choices that follow.

Bad-control logic. Education predicts log GDP per capita one generation forward (Table 13). Because education produces GDP, GDP sits on the causal path (education → GDP → life expectancy, fertility). Controlling for GDP strips out part of education's effect and biases the result against the mechanism (Pearl 2009; Angrist & Pischke 2009, ch. 3). Regressing outcomes on GDP compares education's product against the education that produced it; the proximity of GDP to the outcome gives it an unearned advantage in the comparison.

Natural experiments and the falsification test. Population-scale natural experiments do the identification work that no within-country panel regression can do alone. Chapter 7 tests Prediction 3 on cases where the channel is severed, held fixed across institutional variation, or imposed administratively without real loading. The destructive-disruption case is Cambodia (population destruction → generational shadow). The non-destructive shocks are Sri Lanka (civil war, home niche unbroken, development on schedule), Myanmar (73 years of military rule, outcomes tracked education), China (barefoot doctors dismantled, LE continued rising), and the Asian Financial Crisis (GDP wiped, education untouched). The matched-pair tests are Korea vs Philippines (identical 1950 colonial base, divergent state commitment) and Bihar vs Kerala 1966 (same constitution, free press, central government; literacy 22% vs 55%; famine vs no famine). The treatment in each case — political commitment to mass education, administrative disruption of the educated cohort, or both — was assigned by historical accident with respect to outcome. The lag-decay shape (Figure 5; Section 9.7) is the panel's confirmation that the natural experiments' generational timing reproduces in the cross-country record: education and log GDP per capita produce different decay profiles, with education peaking at lag 25 while GDP's signal is concentrated at short lags.

The falsification test is the USSR case (Chapter 10). Goskomstat reported high lower-secondary completion across the 15 republics 1960–1990; the report was approximately real for the six European-core republics and metropolitan fiction for the eight Caucasian and Central Asian republics. The channel was hollow where the population was peripheral. If reported education sufficed for convergence, the USSR republics should sit on the same TFR/LE/U5MR trajectory as countries with comparable reported completion. The European-core six approximately do; the eight peripheral republics sit 2.6 to 4.0σ above the fit on log U5MR. Excluding the 15 republics strengthens every headline in the paper. Hollow loading does not produce convergence. This is the cleanest falsification test of the mechanism's sufficiency claim, and the channel passes it.

9.2 Data

The analysis covers 185 countries from 1950 to 2015, with education measured as the share of 20–24-year-olds who finished lower secondary school.

Education data are from WCDE v3 (Lutz et al. 2021): lower secondary completion rates for the 20–24 age cohort, both sexes, which reflects completed education rather than enrolment. All education percentages I report refer to this measure unless otherwise noted. Coverage: 185 countries, 1950–2015 at five-year intervals. The core generational panel uses 1975–2015 (1,665 country-years); the residualization analysis (Section 9.8) extends to T=1960–1990 with education interpolated to annual values for precise threshold identification. Long-run analysis uses a 28-country subsample extending back to 1900, selected for data availability; the generational coefficient varies systematically with baseline education level (Section 9.6). GDP data: World Bank, constant 2015 USD, per capita, log-transformed. Life expectancy, TFR, and under-5 mortality: World Bank WDI. Parental education: each country's lower secondary completion lagged 25 years (T−25).

The 15 USSR republics inherit Goskomstat's 1960–1990 reporting: 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 that sat east or south of Moscow's longitude. Barro-Lee v3.0 partially corrects the level but leaves residual bias on life expectancy and child mortality. Chapter 10 documents the convergence signatures of the anomaly and confirms that every headline result I report survives — and strengthens under — exclusion of those 15 republics.

9.3 Descriptive Statistics

The analysis panel is 185 countries observed at 5-year intervals 1975–2015 (1,665 country-years; 1,466 observations after merging the World Bank WDI GDP series, which covers 178 of 185 countries). Parental lower secondary completion is observed at $T-25$; child completion, log GDP per capita, life expectancy, TFR, and under-five mortality at T.

Table 7: Period means, six core variables.

Variable	1975–1989	1990–2004	2005–2015
Parental lower-sec completion (%, $T-25$)	27.0	42.7	56.2
Child lower-sec completion (%, T)	52.3	62.2	71.1
Log GDP per capita (const. 2015 USD)	8.07	8.19	8.52
Life expectancy at birth (years)	61.4	65.5	70.1
Total fertility rate (births/woman)	4.53	3.63	2.95
Under-5 mortality (per 1,000 live births)	98.8	66.8	39.2

Notes: Five-year observation bands. Each variable shifts monotonically across the sample window as the global educational expansion propagates into convergence. Parental completion rises 29 pp over the sample; child completion rises 19 pp; total fertility halves ($4.53 \to 2.95$); life expectancy rises 8.6 years; under-five mortality falls by 60%. The parental column lags child by one generation (25 years) by construction: the 1975–1989 parental mean of 27% corresponds to the 1950–1964 cohorts of those same countries. Full summary statistics (n, mean, sd, min, max) and the complete 185-country list are in Appendix A.3.10.

9.4 Completion as the Operative Variable

I measure whether people finished school, not how well they scored on tests. Completion is the right measure for a generational mechanism.

Section 4.2 catalogued the jumps formal education installs — germ theory, heliocentrism, the natural numbers extending without bound, atoms, evolution, the laws of motion — and the stack they compose into. Each is a kind-flip; there is no halfway position between germ theory and miasma. The empirical question for the panel is whether a given child crossed enough rungs of the stack to reorganise their adult decisions. Completion — the certificate that the child sat inside the schooling environment to its conclusion — is the proxy that scales across 185 countries. Test scores measure which rungs stuck on a single day; completion measures whether the exposure happened at all.

Why exposure, not learning, is the operative variable. Section 4.5 established the payload: cognitive reorganisation, not content retention. The panel uses completion — not years alone, not test scores — because completion is where the mechanism registers at the generational scale. What formal schooling provides is not a specific curriculum but sustained exposure to literate culture: structured time, abstract categorisation, written language, numeracy, adult authority organised around knowledge transmission. The longer the child remains inside this environment, the deeper the reorganisation. Primary gives capacity over reproduction (Section 5.3). Secondary gives capacity over health navigation, economic participation, civic engagement. Tertiary gives capacity over complex systems. Each level is more time inside the categorical change.

A child who completed nine years of poor-quality schooling and scored badly on every assessment has still spent nine years inside an environment that reorganised their cognition categorically relative to a parent who never entered the building. The test score measures how well the payload was delivered. Completion measures whether the child was inside the channel long enough for the reorganisation to consolidate. The life-course data confirm it: no generation in any country has completed school at a lower rate than its parents, and the life-course returns to education show no decay with age — exactly what a reorganisational mechanism predicts and a content-retention mechanism does not.

What the curriculum's breadth is for. Years of sustained exposure to biology, chemistry, history, mathematics, and literature produce adults who spent years inside five different formal knowledge systems — each with its own mode of abstraction, its own vocabulary, its own relationship between evidence and conclusion. The breadth of the curriculum is not an attempt to produce polymaths. It is sustained duration of exposure across multiple modes of formal reasoning, each contributing to the same cumulative reorganisation. Tests measure whether the student retained the content of each subject; the mechanism operates through the years the student spent inside the reorganising environment.

The generational gate. The distinction sharpens at the generational gate. The educated parent's role in the home niche is not to teach their child what they learned in school. It is to hold the expectation that their child belongs in school — to see completion as normal, to monitor attendance, to assume that the educated world is where their child will live. This expectation is the gate. An uneducated parent cannot open it, because they have no experience of the world on the other side. An educated parent opens it automatically — through self-interest, not ideology — because school is where they came from and where their child belongs. The gate does not require the parent to be a good teacher. It requires the parent to have been a student.

This is also why the grandparent effect is strongest at low baselines. Where education is scarce, the grandparent is a second educated person in the household — and, in a village where education is rare, a visible exemplar to the wider community. At high baselines, the entire society holds this expectation and the grandparent's independent contribution disappears — exactly the pattern the data show (grandparent β significant below 50% parental completion, zero above it).

From reorganisation to specialisation. The reorganised individual naturally specialises. Categorical literacy — the capacity for systematic learning, abstract reasoning, structured problem-solving — is not an end state. It is the precondition for further skill acquisition. The educated individual can absorb trade skills through apprenticeship, on-the-job training through employment, formal knowledge through college, deep expertise through a PhD. The specific pathway does not matter. What matters is that each pathway requires the cognitive infrastructure that formal schooling built. The uneducated individual is not excluded from specialisation by lack of opportunity alone. They are excluded because the capacity to learn a structured skill systematically — to read a manual, follow a protocol, internalise an abstract procedure — was never constructed.

The empirical confirmation. A direct test partitions the outcomes by the rungs of the stack each requires. On a 104-country cross-country panel with USSR republics excluded, using Angrist et al. (2021) HLO test scores and WCDE completion: for TFR, quantity (primary completion) dominates at $\beta_z = -0.64$ ($t = -7.4$); adding test-score quality contributes $\beta_z = -0.25$ ($t = -2.9$), taking R^2 from 0.70 to 0.72. Primary completion alone drives the fertility transition. For child mortality (log U5MR) the partition reverses: quantity goes insignificant ($\beta_z = -0.16$, $p = 0.09$); quality dominates ($\beta_z = -0.70$, $t = -7.6$, R^2 = 0.78). For life expectancy: quantity insignificant ($t = -1.1$); quality dominates ($\beta_z = +0.83$, $t = +7.7$, R^2 = 0.62). The three outcomes separate by stack height. *Fertility is a threshold flip that runs on the first rungs of the stack and registers even where schooling is shallow; child survival and adult longevity require composed stacks — germ theory plus dose-response plus the hygiene chain — and do not respond where credentials are hollow (Section 4.2).*

Read alone, the U5MR and LE results vindicate Hanushek — but the vindication is the mechanism reasserting itself, not an alternative to it. A continuous lag sweep (Chapter 10, §10.7, and Figure 11) shows that HLO secondary today is best predicted not by current schooling but by the *parental* generation's completion 10–25 years earlier; R^2 at the great-grandparent lag (60 years) remains near 0.49. Hanushek's test scores are not a rival framework. They are the paper's quantity measure integrated across three generations of parental transmission. The horse race partitions the way the biology predicts, and Chapter 10 develops the partition.

9.5 Empirical Strategy

Each country is compared only to its own past. The question: does parental education predict what happens to their children one generation later?

Primary specification — country fixed effects regression:

$$E_{it} = \alpha_i + \beta_1 E_{i,t-25} + \varepsilon_{it}$$

In plain language: a country's education level today is predicted by its own education level 25 years ago (the parents' generation). The α_i term gives each country its own baseline — so the test is not "which countries are more educated?" but "within the same country, does the previous generation's education predict the next generation's?" This absorbs everything permanent about a country: institutions, geography, culture, colonial history. What remains is change over time. The 25-year lag is not a researcher choice — it is the time-to-agency interval derived from the biological mechanism (Prediction 1).

Comparison specifications replace parental education with log GDP per capita, or include both:

$$E_{it} = \alpha_i + \beta_2 \ln Y_{it} + \varepsilon_{it}$$

$$E_{it} = \alpha_i + \beta_1 E_{i,t-25} + \beta_2 \ln Y_{it} + \varepsilon_{it}$$

For development outcomes (life expectancy, fertility, child mortality), I measure outcomes 25 years after the education measurement: education at time T predicts outcomes at T+25. Temporal ordering rules out reverse causation; time-varying confounding is handled by the natural experiments in Section 9.1:

$$O_{i,t+25} = \alpha_i + \gamma_1 E_{it} + \gamma_2 O_{it} + \varepsilon_{it}$$

Including the initial outcome level (O_{it}) is deliberately conservative: it absorbs the country's prior trajectory, so the test asks whether education predicts improvement beyond whatever trend was already underway.[4]

Residualization strategy. If education causes GDP, part of GDP's apparent predictive power is really education's effect in disguise. To test whether GDP contributes anything on its own, the design strips out the portion of log GDP per capita that education produced and uses only the remainder. Mechanics and results are in Section 9.8 (Table 17).

On year fixed effects. Adding country-and-year fixed effects on the full

[4] Both-sexes completion is used throughout; female-only specifications produce stronger effects (Section 9.8).

185-country panel collapses the parental coefficient from 0.483 to 0.083 — the mechanical signature of the staggered-adoption pathology, not evidence against the mechanism. Three-quarters of the 2WFE weight sits on comparisons between newly-expanding countries and already-high-completion controls (Goodman-Bacon decomposition). The Callaway-Sant'Anna (2021) estimator, limited to clean treatment-timing comparisons, recovers the effect: $+7.9$ pp at the generational lag and $+21.4$ pp by year 35, with pre-treatment coefficients indistinguishable from zero (Figure 7). Full decomposition weights, event-study pre-trend tests, and the continuous-treatment discretisation caveat are in Appendix A.1. I keep the one-way country-FE specification as the headline because it corresponds to the "how much do educated countries transmit to their children" test I set up.

9.6 Education vs. GDP as Predictors of Attainment (Prediction 2 and Prediction 6)

Where it matters most — countries still building their education systems — education predicts the next generation's attainment 3× better than log GDP per capita does. This is the test for Prediction 2 (education persists across generations, GDP does not) and Prediction 6 ($\beta_g > 1$ at low baselines).

Table 9 estimates child lower secondary completion at T on parental completion 25 years earlier, on the active-expansion subsample (parental completion below 30%), with three increasing controls: contemporaneous log GDP per capita, a parental-quadratic, and year fixed effects. All four columns use the identical 629 country-years / 105 countries.

Table 9: Education predicts child lower secondary completion one generation later.

	Child lower secondary completion (t)					
	(1)	(2)	(3)	(4)	(5)	(6)
Parent edu ($t-25$)	1.376***	—	1.270***	2.039***	0.830***	1.570***
	(0.083)		(0.059)	(0.273)	(0.133)	(0.409)
Parent edu^2	—	—	—	-0.026***	—	-0.019*
				(0.008)		(0.010)

	Child lower secondary completion (t)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log GDP (t)	—	13.66***	5.02*	4.57*	3.97	3.99
		(3.90)	(3.04)	(2.75)	(3.11)	(3.00)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	—	—	—	—	Yes	Yes
R^2 (within)	0.699	0.214	0.724	0.746	0.644	0.717
Observations	629	629	629	629	629	629
Countries	105	105	105	105	105	105

Notes: Dependent variable: child lower secondary completion (percent, 20–24 cohort, both sexes, year t). Sample restricted to country-years with parental completion below 30% at year $t - 25$ (active expansion); identical 629 country-years / 105 countries across all six columns. Country-clustered standard errors in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Columns: (1) parent only, country FE – the headline one-way FE; (2) log GDP only, country FE – the bivariate horse race against parental education; (3) parent + log GDP, country FE – bad-control test, the parental coefficient barely moves ($1.376 \rightarrow 1.270$) while GDP attaches only a weakly significant slope; (4) col 3 + parent^2 – the negative quadratic coefficient implies diminishing returns (β_g compression near the ceiling); (5) col 3 + year FE; (6) col 4 + year FE – the parental coefficient attenuates but stays significant at $p < 0.01$ across both two-way-FE columns, unlike the full-panel 2WFE collapse in Table A1, because restricting to the active-expansion subsample minimises the staggered-adoption pathology that drives it. GDP is excluded from the headline as a bad control (education $\rightarrow$ GDP $\rightarrow$ outcomes; Pearl 2009); Table 17 residualises it before the comparison.

Reading the columns directly: among countries where the expansion is still active, a one-percentage-point rise in parental lower secondary completion raises child completion one generation later by **1.38 pp** (column 1; $t = 16.4$). Column 2 runs log GDP alone on the same sample: the GDP coefficient is statistically significant (13.66, $p < 0.01$) but explains only $R^2 = 0.214$ of within-country variation, far less than education's $R^2 = 0.699$ in column 1. Column 3 puts both into the same regression: the parental coefficient barely moves ($1.376 \rightarrow 1.270$) and attaches only a weakly significant 5.02 pp per log-unit of GDP per capita — on a percent-for-log-point metric, education's effect is estimated nearly five times more precisely than GDP's on the same sample. Allowing curvature (column 4) reveals diminishing returns and raises the linear coefficient to 2.04: the linear specification in column 1 understates the effect at low baselines. The two-way fixed-effects columns (5) and (6) attenuate the parental coefficient (0.83 and 1.57 respectively) but both remain significant at $p < 0.01$, while GDP's coefficient loses significance entirely once year FE absorb the global trend. The stability of education's coefficient and the fragility of GDP's significance is the empirical signature.

The Table 9 headline is conservative: excluding the fifteen USSR republics whose Soviet-era lower-secondary reporting is documented anomalous in Chapter 10 raises the parental-education coefficient from 0.483 to 0.538 and the within-R^2 from 0.457 to 0.507.

The next table (11) presents the headline education specification (Table 9, column 1) at three representative cutoffs on the same GDP-merged panel, so differences across rows reflect predictive power within different stages of the expansion rather than sample composition. Education leads GDP at every cutoff; the parental coefficient is significant at $p < 0.01$ in every row.

Table 11: 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%	2.022*** (0.258)	0.590	15.535*** (3.291)	0.296	275	58
<50%	1.053*** (0.059)	0.697	17.311*** (3.942)	0.247	829	116
Full	0.483*** (0.034)	0.533	15.787*** (2.114)	0.245	1,665	185

Notes: Each row is a single-regressor within-country fit of child lower secondary completion at T on the regressor at T (parental education for the Edu columns; log GDP per capita 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 unrestricted panel (1,665 country-years, 185 countries) — the max-sample estimate on every observation available, 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: $2.0\times$ (<10%), $2.8\times$ (<50%), $2.2\times$ (Full); the ordering holds across the full 10%–90% sweep in Table A4, confirming Prediction 6.

Not driven by one region or one era. Fifteen subsample re-estimations of Column 1 (six regions, two eras, three within-sample GDP terciles) all yield positive coefficients significant at $p < 0.01$ (Section 9.9, Table 19); the headline is not an artefact of East Asia, the post-Cold-War period, or any income band.

Each generation gains more than it inherits (Prediction 6). The generational amplification coefficient β_g is the within-country slope of child completion on parental completion at the one-generation lag, with country fixed effects (Glossary). $\beta_g > 1$ means a one-pp rise in the parental cohort predicts a more-than-one-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 life expectancy shows no diminishing returns from primary through tertiary.

On the long-run 1900–2015 panel (one-way country FE), $\beta_g = 2.86$ below 20% parental completion (within-country R^2 explains 70% of child-completion variation), 1.83 below 50% (79%), 1.24 below 90% (77%), and 1.04 unrestricted. The post-1975 panel that backs Table 9 runs lower at every cutoff because it is weighted toward countries already approaching the ceiling: $\beta_g = 0.483$ unrestricted (Table 11), rising to 1.376 at the $<$30% active-expansion subsample where the channel is still moving. Adding year fixed effects on the active-expansion subsample (Table A1) attenuates further but still leaves β_g above 1: 1.032 at $<$20% ($t = 5.8$) and 1.019 at $<$10%. The full-panel 2WFE collapse to 0.083 is the staggered-adoption pathology developed in Appendix A.1, not a refutation of the prediction.

Ceiling compression within countries. Within individual countries, β_g varies systematically with baseline education level (Figure 2). The United States shows β_g=1.9 when its baseline was 13% (early 20th century), declining to β_g=0.08 when its baseline reached 92%. Korea shows β_g=6.5 at 1% baseline, declining through 3.6, 1.8, and 0.2 as its baseline rose to 59%. Taiwan follows a nearly identical trajectory (β_g=5.1 at 1.2%). The Philippines (Section 16) shows a flatter curve: β_g=4.4 at 1.5% declining more slowly to 0.4 at 49%. The ceiling-compression pattern is robust across all countries with sufficient historical depth. Bangladesh shows β_g≈2 throughout its observed trajectory (1–18% baseline), consistent with a country still well below the ceiling. India follows a similar pattern at β_g≈1.6 through 28% baseline.

Non-depreciation holds universally — the Cambodia test (Section 7.2).

9.7 Education Predicts Development Outcomes 25 Years Forward (Prediction 1)

Prediction 1 says outcomes at $T{+}25$ follow education at T with a one-generation lag. Knowing how educated a country's parents were predicts how long their children will live, how many children they will have, and how much

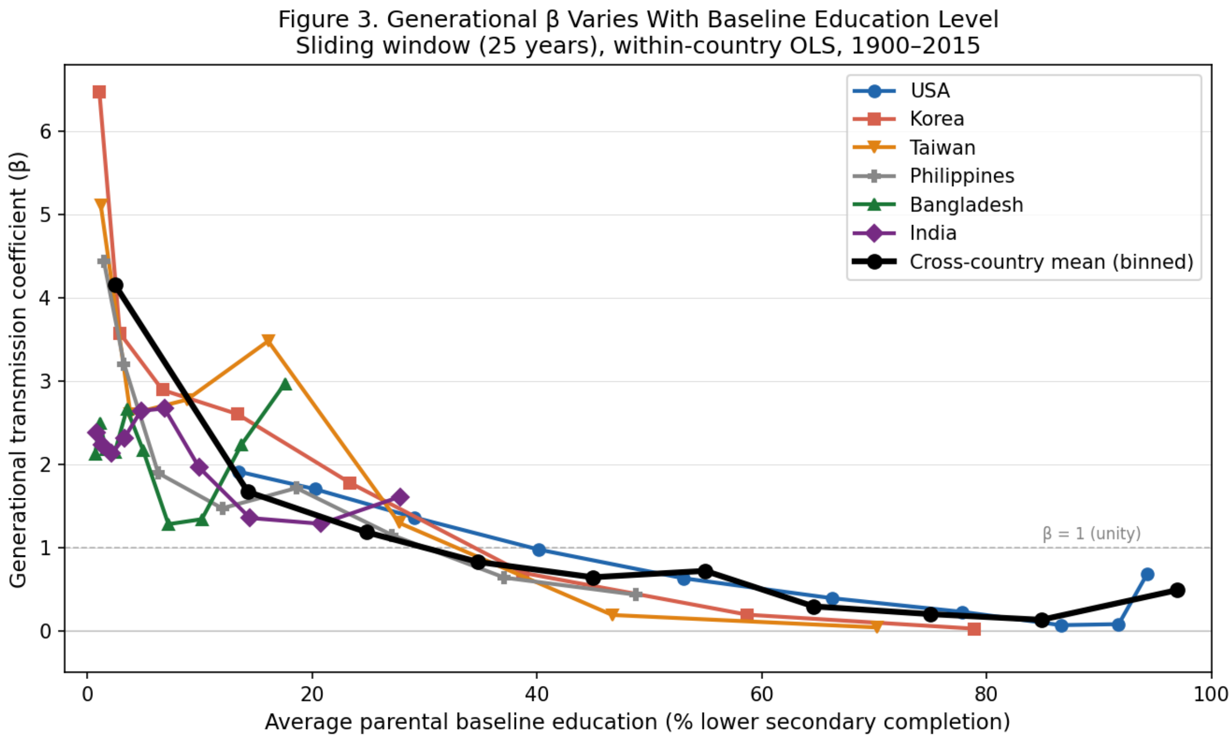

Figure 4: Generational β_g varies systematically with baseline education level. Each thin coloured line: 25-year sliding window (6 child cohorts) within one country, against the window's parental baseline. The x-axis is an implicit time dimension: each country's points move from left (low baseline, early periods) to right (high baseline, later periods) as schooling expands. The thick black line is the cross-country mean across *all* countries with sufficient data (185 countries, $n = 1{,}388$ country-windows), binned by parental baseline in 10-percentage-point intervals — not just the six illustrated trajectories. The six countries are illustrative selections: USA spans the long ceiling approach (1900–2015); Korea and Taiwan are the showcase rapid-state cases discussed in §8; Philippines is the matched-base, slower-state contrast in the same case section; Bangladesh and India represent the active-expansion countries currently below the ceiling. The appendix figure (Figure 12) shows every country-window with an OLS fit and 95% confidence band so the reader can assess the universality claim independently of the country selection. $\beta_g > 1$ at low baselines reflects state reach; $\beta_g \to 0$ at high baselines reflects ceiling compression. *(WCDE v3 long-run cohort reconstruction 1875–2015, broader than the 28-country regression panel in Table A10.)*

they will earn one generation later — the panel pattern is consistent with the one-generation transmission mechanism the theory specifies.

Table 13: Education and development outcomes one generation forward.

Panel A: Education(T) → development outcomes (T+25). Four-column stepwise on max-sample. Predictor: lower-secondary completion at T (TFR row uses primary, marked in label). Each row is one outcome; each column is one specification. Education coefficient with country-clustered SE in parentheses.

Outcome (T+25)	(1) edu only	(2) + initial	(3) + log GDP	(4) + year FE
log GDP per capita	+0.0149***	+0.0114***	+0.0114***	+0.0037
	(0.0013)	(0.0016)	(0.0016)	(0.0023)
log(LE)	+0.0029***	+0.0015***	+0.0021***	−0.0005
	(0.0002)	(0.0004)	(0.0004)	(0.0004)
log(TFR) [primary]	−0.0137***	−0.0124***	−0.0116***	−0.0082***
	(0.0008)	(0.0009)	(0.0010)	(0.0012)
log(U5MR)	−0.0276***	−0.0080***	−0.0096***	−0.0013
	(0.0014)	(0.0022)	(0.0025)	(0.0021)
Country FE	Yes	Yes	Yes	Yes
Year FE	—	—	—	Yes
N (log GDP row)	1,205	927	927	927
N (log LE row)	1,259	1,252	925	925
N (log TFR row)	1,259	1,254	927	927
N (log U5MR row)	1,225	1,225	911	911

Panel B: Log GDP per capita(T) → lower secondary completion (T+25). Reverse regression, four-column stepwise.

	(1) GDP	(2) + edu	(3) + year FE	(4) + edu+yr FE
Log GDP (T)	+14.42***	+2.87	+2.30	+1.20
	(1.79)	(2.02)	(2.06)	(2.18)
Initial education (T)	—	+0.485***	—	+0.145**

	(1) GDP	(2) + edu	(3) + year FE	(4) + edu+yr FE
		(0.048)		(0.067)
Country FE	Yes	Yes	Yes	Yes
Year FE	—	—	Yes	Yes
R^2 (within)	0.263	0.495	0.009	0.032
Observations	927	927	927	927
Countries	172	172	172	172

Notes: Panel A regresses each outcome at $T{+}25$ on the predictor at T in a four-column stepwise: (1) predictor only; (2) + initial outcome at T; (3) + log GDP per capita at T; (4) col 3 + year FE. Panel B regresses lower-secondary completion at $T{+}25$ on log GDP per capita at T in the same stepwise structure (b1 GDP only, b2 + initial edu, b3 GDP only + year FE, b4 + initial edu + year FE). Country fixed effects throughout; country-clustered standard errors in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All Panel A coefficients are semi-elasticities (proportional change in the outcome per percentage-point rise in education); units therefore comparable across outcome rows. "Max sample" means each spec drops NAs only on the regressors actually used — the max-sample estimate on every observation available for that spec. The common-927 sample (intersection across all four outcomes plus log GDP) appears as the appendix robustness in Table A6; coefficients move negligibly between the two samples. The TFR row uses primary completion because the steepest fertility decline occurs at primary (Section 5.3). The other three rows use lower-secondary, consistent with Table 9. Panel B's drop from $\beta = 14.42$ to $\beta = 2.87$ when initial education enters (col 1 → col 2) is the bad-control pattern (Pearl 2009): initial education sits on the causal path log GDP → initial education → future education, so controlling for it absorbs log GDP's contribution by design (see Section 9.1). Adding year FE alone (col 3) also kills GDP's predictive power ($\beta = 2.30$, $p = 0.26$) — once the global time trend is absorbed, GDP retains essentially no within-country signal. *Sample relative to Table 9.* Table 9 and this table use deliberately different filters because they answer different questions. Table 9 tests whether parental education predicts child education one generation forward, so its filter is on the *predictor at the time of measurement* — country-years where parental completion at $T{-}25$ is below 30% (active expansion in the parental cohort), 1975–2015, $n = 629$ / 105 countries. The table here tests whether being inside the educational expansion *window* predicts development outcomes, so its filter is on the *country's expansion phase* — country-years where lower-secondary completion has crossed 10% but remains below the 60%/90% ceiling, 1960–1990, $n = 927$ / 179 countries (GDP row) and $n = 1{,}295$ / 185 countries (LE/TFR rows). The two filters select different country-year sets by construction: the first conditions on the parental-cohort level, the second on the country's expansion stage. Coefficients in the two tables are therefore not directly comparable in magnitude; both are consistently estimated within the FE model on their respective sample.

Log GDP per capita alone explains 27% of child education variation across the full panel (Panel B col 1). But add the parents' education to the same model and GDP's apparent contribution collapses — the model's explanatory power jumps from 27% to 50%, almost entirely from education, not GDP.

Reading Panel A column 2 directly (the headline "edu + initial outcome" spec): each 1-percentage-point gain in lower secondary completion predicts, one generation forward, 1.14% higher GDP per capita, 0.15% higher life expectancy (≈0.10 years at LE $= 70$), and 0.80% lower under-five mortality.

For TFR the operative channel is primary completion: each 1-percentage-point gain in primary at T predicts 1.24% lower TFR (≈0.04 fewer children per woman at TFR $= 3.5$). All four results are within-country effects controlling for each country's starting point. Every coefficient in Panel A columns 1–3 is significant at $p < 0.01$; the year-FE column 4 attenuates U5MR and LE to non-significance, but the lower-secondary log(TFR) coefficient remains $-$0.82% per pp ($p < 0.001$) and the headline log GDP coefficient remains positive. Across all four outcomes the predictor's coefficient barely moves between columns 2 and 3 once log GDP enters as a control, confirming that GDP is downstream of education rather than an independent channel.

The education-vs-GDP predictive gap holds for child education at every cutoff tested.

For life expectancy the gap is even wider: among countries below 30% completion, education explains 33% of the variation vs 2% for GDP — a 14× gap. Among countries below 10%, education explains 51% while GDP explains 1.6%. Even that small GDP signal vanishes once education's contribution is stripped out (Section 9.8). The same pattern holds for the other outcomes: at completion below 50%, fertility shows a 32× edu-over-GDP ratio (edu R^2=0.359, GDP R^2=0.011); at completion below 10%, log under-5 mortality shows a 39× ratio (edu R^2=0.570, GDP R^2=0.015).

Lag-decay shape (Prediction 2)

Education's signal survives four generations across every outcome tested. Log GDP per capita's collapses within one.

Each lag corresponds to a generational relationship, and the relationship names what the number actually measures. The same pattern holds for all four development outcomes I track:

Table 16: Lag-decay of education on development outcomes.

Lag	Gens	Relationship	LE	TFR	log(U5MR)	Child edu
0	—	contemporary	0.743***	0.817***	0.814***	—
25	1	parent → child	0.597***	0.693***	0.820***	0.722***
50	2	grandparent → grandchild	0.411***	0.452***	0.664***	0.388***
75	3	great-grandparent → great-grandchild	0.284***	0.258***	0.502***	0.227***
100	4	g-great-gp → g-great-grandchild	0.223***	0.134***	0.395***	0.152***

Notes: Standardised within-country $|\beta|$: outcome on lower-secondary completion at T$-$lag. Each series demeaned by country mean and standardised by pooled within-country SD; coefficients are therefore not directly comparable to the unstandardised β in Tables 9 and 13. Signs: $+$ for LE and child education, $-$ for TFR and U-5; absolute values shown. Child education at lag 0 is omitted because predictor and outcome are the same series. Sample: 142 countries (intersection of edu / GDP / LE / TFR / U-5), outcome panel 1960–2015. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$; all reported coefficients are significant at $p < 0.001$ (minimum $|t|$ at lag 100 is 10.9, TFR).

Every outcome's coefficient survives four generations. Under-5 mortality carries the deepest signature: $|\beta| = 0.395$ at century depth with $|t| = 34.8$, roughly twice the strength of any other outcome and the only one where $|\beta|$ exceeds 0.3 after four intervening households. Life expectancy, fertility, and child education show the same smooth monotonic decay, all still significant at 100 years ($|t|$ of 18.4, 10.9, and 12.4 respectively). Roughly a century of generational transmission running through three intermediate households leaves a statistically detectable trace in every dimension of human development I define. (R^2 values show the same monotonic decay pattern. They are not directly comparable across outcomes because they conflate effect size with predictor variance and outcome noise; see `checkin/outcomes_r2_by_lag.json` for the R^2 decay curves.)

No other development input has comparable reach. Log GDP per capita's R^2 is already fading by lag 20–25 and falls below education across the 0–45 GDP-availability window (Figure 5); past lag 45, GDP cannot be tested at all on raw WDI data — not because the data is missing but by construction. Pre-literate subsistence societies had uniformly low GDP per capita (~$400–600, constant 2015 USD), so there is no cross-country GDP variation to test before literate populations existed. The uniformity is not offered as evidence for the mechanism; it is a reminder of what the mechanism had to produce before GDP per capita could become a variable at all. Institutions, quality scores, and provision fail the same generational-depth test: none survives at the depth education reaches. This is not a separate finding for each alternative — it is a single methodological principle derived from the mechanism itself. Backfilling pre-1960 GDP at $500 does not rescue GDP's predictive power for any of the four outcomes. The backfilled-GDP coefficient crosses zero and loses significance by lag 55–60 for life expectancy, fertility, log under-5 mortality, and child education (within-R^2 drops below 0.01 at the same points), while education's coefficient remains significantly signed through lag 100.

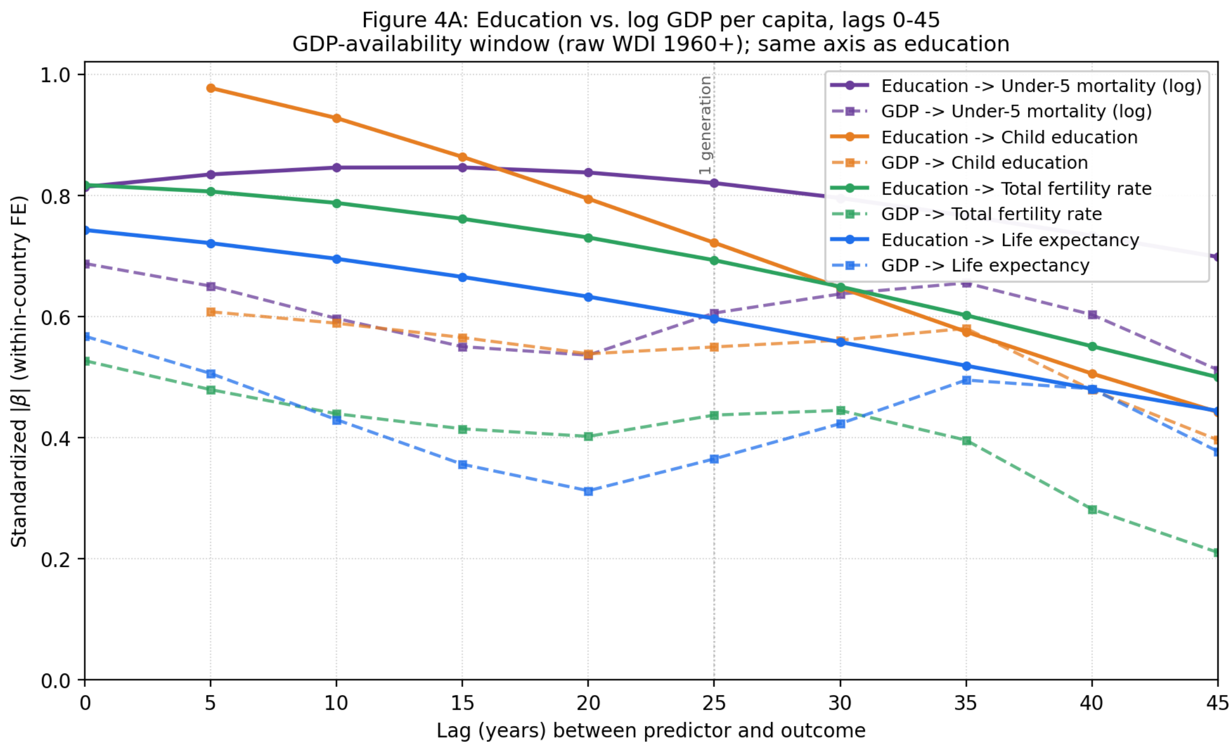

Figure 5: Figure 4A: Education vs. log GDP per capita, lags 0–45 years (the GDP-availability window). Solid: standardised within-country $|\beta|$ for lower secondary completion predicting four development outcomes. Dashed: log GDP per capita predicting the same outcomes; raw WDI coverage starts in 1960 so lag 45 is the last lag with usable GDP cells. Both predictors share the same axis; education leads GDP at every outcome and every lag inside this window. Each predictor and outcome is divided by its pooled within-country standard deviation before estimation, so coefficients are comparable across outcomes but *not* directly comparable to the unstandardised β in Tables 9 and 13. Sample: 142 countries (intersection of education, GDP, LE, TFR, U5MR data availability for outcome years 1960–2015), narrower than the headline samples in Tables 9 and 13. *(WB WDI and WCDE v3.)*

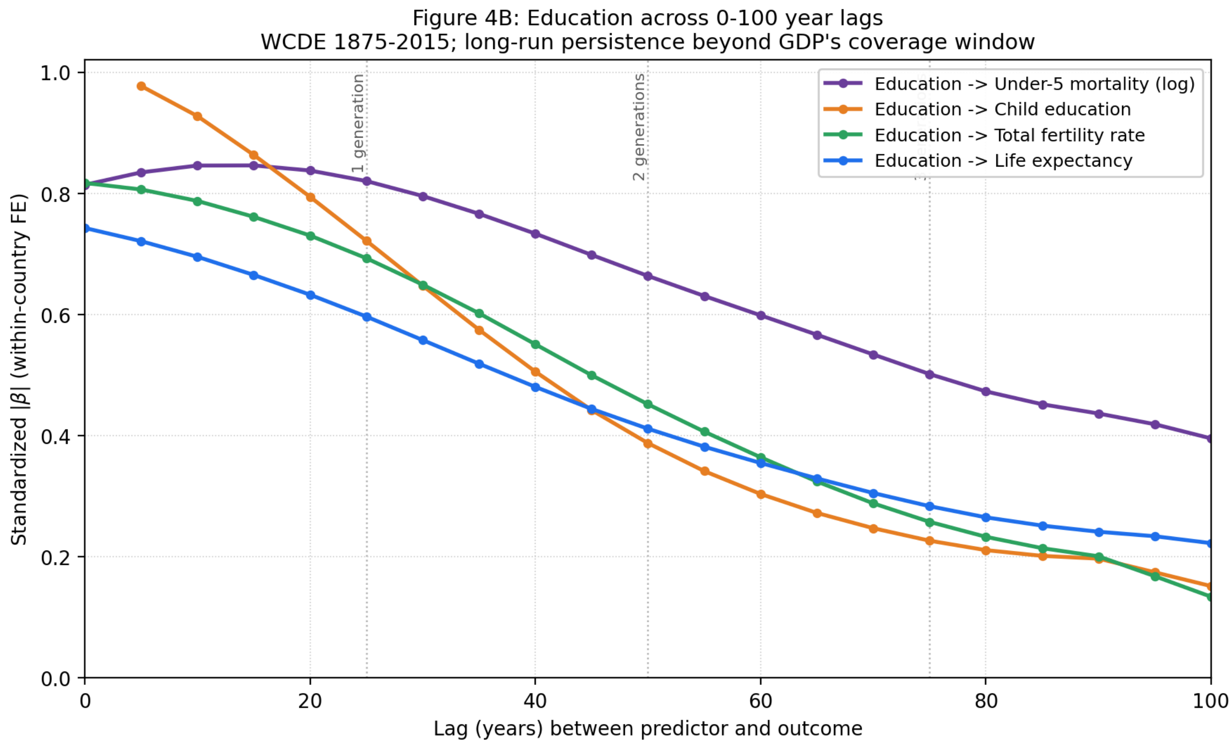

Figure 6: Figure 4B: Education across lags 0–100 years. Same standardised within-country $|\beta|$ as Figure 5 but extended to four-generation depth (WCDE v3 cohort reconstruction 1875–2015). GDP is omitted because raw WDI data start in 1960 and cannot be tested past lag ∼45 (Figure 5); the backfilled-GDP exercise in §A.3 confirms that GDP's coefficient crosses zero and loses significance by lag 55–60 for all four outcomes. Generational horizons 1–3 marked.

The compounding-generations signature (Prediction 2)

The mechanism in Section 5 predicts that where a country first crosses the 10% lower secondary completion threshold, the effect on child completion should grow with the generational window, not plateau at $T+25$. Parents who crossed raise children more likely to cross, who raise children still more likely to cross. The Callaway-Sant'Anna event study confirms this shape directly: +1.3 percentage points immediately, +6.6 at ten years, +14.9 at twenty-five years, +21.4 at thirty-five years. Pre-treatment coefficients are indistinguishable from zero — countries that will expand and countries that will not follow identical trajectories before the crossing. The signature is the monotonic compounding through generations after the crossing, not the cross-sectional level at any single lag.

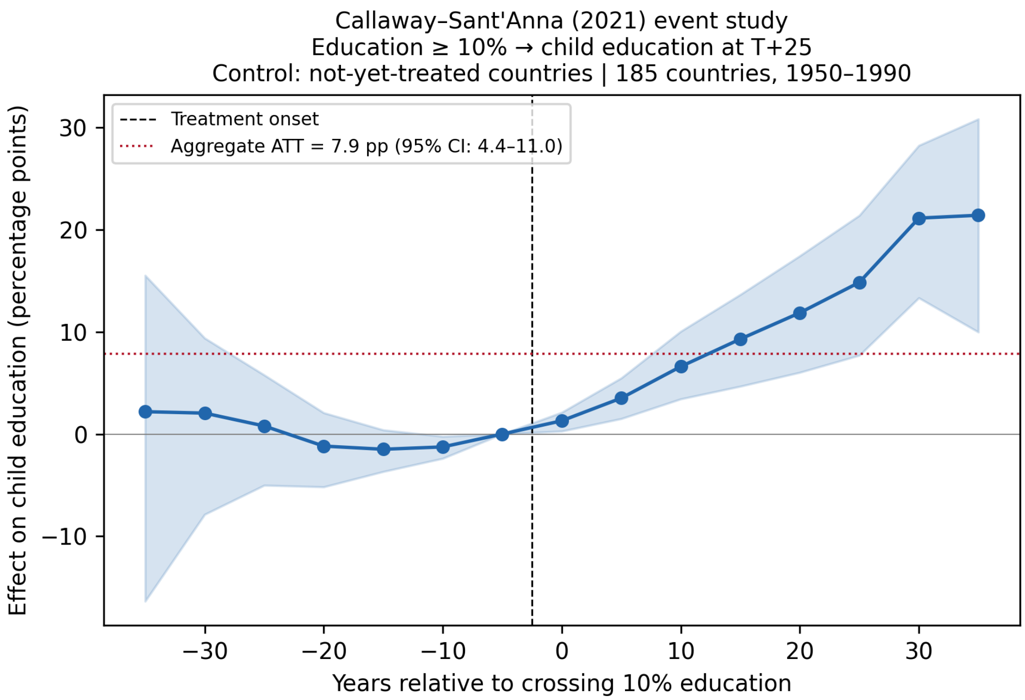

Figure 7: Generational compounding after the 10% threshold crossing. Treatment: country crosses 10% lower secondary completion. Outcome: child education at $T+25$. Control group: countries not yet treated at each comparison point. Shaded band: 95% confidence interval (500-replication clustered bootstrap). Pre-treatment coefficients near zero; post-treatment effects grow monotonically from +1.3 to +21.4 percentage points over 35 years — the compounding-generations signature the mechanism in Section 5 predicts. Callaway-Sant'Anna (2021) estimator. Full decomposition and staggered-adoption diagnostics in Appendix A.1.

The grandparent channel

The lag-decay results show that education predicts across multiple generations. A sharper test: does the grandparent's education predict outcomes *independently*, after controlling for the parent's own education?

For fertility in the full panel, the answer is no — grandparent education adds nothing once the parent's is controlled (β_{gp}=+0.001, p=0.80). The parent absorbs the grandparent's entire contribution to fertility. But restrict to countries where education is still scarce (parent completion below 50%) and the grandparent effect emerges: β_{gp}=−0.059 (p=0.0003), nearly twice the parent's own coefficient (β_p=−0.033). R^2 jumps from 0.39 to 0.46. Where education is scarce, having an educated grandparent independently reduces fertility beyond what the parent's own education predicts. The effect is not gendered: grandfather and grandmother education produce comparable coefficients in the low-education subsample (β=−0.054 and −0.050 respectively), consistent with oblique transmission — in a village where education is rare, any educated elder is a visible exemplar.

For life expectancy, the grandparent effect is significant across all education levels (β_{gp}=+0.070, p<0.0001), adding 3.6 percentage points of R^2 beyond the parent alone. Roughly half the grandparent's total effect on LE runs through the parent; half is direct. For grandchild education, the grandparent coefficient is β_{gp}=+0.271 (p<0.0001), with an R^2 gain of 5.2 points.

The largest grandparent effect is on child survival. Regressing log under-5 mortality on both generations' lower secondary completion in the full panel, the grandparent coefficient is larger in absolute terms than the parent's (β_{gp}=−0.018, β_p=−0.016; $|\beta_{gp}/\beta_p|$=1.16, both p<0.0001). Within-R^2 rises from 0.377 to 0.561 once grandparent education is added — a gain of 18.4 percentage points, roughly five times the life-expectancy gain and the largest grandparent contribution of any outcome I track. The grandmother's cohort carries an independent signal for child survival that, in magnitude, matches the mother's own education — the signature expected if the mechanism is kin transmission of child-care practice (sanitation, feeding, vaccination) within the household and kin radius.

The pattern is what the generational gate predicts (Section 9.4): at low baselines each educated person in the household adds independent weight; at high baselines the parent has internalised everything the grandparent could offer. The grandparent effect on fertility is real, but it operates precisely where it matters most — the steepest part of the transition.

This generates a falsifiable prediction distinguishing the kin/community-radius account from a pure household account. Any educated elder inside the child's daily kin/community radius — co-resident grandparent, village elder, aunt, older cousin — should raise grandchild outcomes beyond what the parent's education alone predicts. Equivalent elders present only on paper (deceased, non-resident, absent from daily life) should not. The channel runs through physical presence and daily interaction at Radii 1–2 (household and kin/identity; Section 5.2), not through inherited status. That the grandfather and grandmother coefficients come out nearly equal in the low-education subsample (β=-0.054 and -0.050) is consistent with the same reading: any educated elder works, not specifically the grandparent of the same sex. DHS co-residence data can test the presence-versus-absence prediction directly.

9.8 GDP Has No Independent Effect (Prediction 4)

Prediction 4 says log GDP per capita carries no independent predictive power once education's contribution is removed. Take each country's log GDP per capita and residualise out the part that education produced; use only what remains to predict development outcomes. The residual predicts nothing — with one bounded exception in post-2000 child mortality, examined below.

Table 17: Residualised GDP analysis – education vs. log GDP per capita as predictors.

Outcome (T+25)	Edu R^2	Raw GDP R^2	GDP R^2 (edu removed)	p-value
Life expectancy	0.472	0.179	**0.003**	0.56
Total fertility rate	0.478	0.175	**0.000**	0.87
Child education	0.524	0.303	**0.008**	0.31
Under-5 mortality	0.284	0.046	**0.023**	0.11

Notes: Entry-cohort design (entry ≥ 10%, ceiling ≤ 90%); country fixed effects; lower secondary completion at T, outcome at $T+25$; $T = 1960–1990$; country-clustered standard errors. "GDP R^2 (edu removed)" is the predictive power of log GDP per capita after stripping out education's contribution (Frisch-Waugh-Lovell). Sample sizes: $n = 822$ (152 countries) for LE and TFR; $n = 856$ (157 countries) for child education; $n = 787$ (147 countries) for under-5 mortality. $T = 1960–1990$ is the full WB WDI GDP window (raw coverage starts in 1960); the pre-1960 backfill exercise that pushes GDP to lag 100 (missing years fixed at the $500 subsistence floor) is in §A.3 and shows GDP's coefficient crosses zero and loses significance by lag 55–60 for all four outcomes.

Reading the table directly: a one-percentage-point rise in lower secondary completion at T predicts a one-generation-later gain of $+0.11$ years of life, -0.032 births per woman, and a -2.7 percentage-point drop in log under-five mortality; the same-magnitude response to residualised log GDP per capita (with education's contribution stripped out) is statistically indistinguishable from zero in every row. The R^2 columns summarise the same thing at the country-panel level: education explains a non-trivial share of within-country variation across all four outcomes (0.284 to 0.524), while the leftover-GDP R^2 is at most 0.023 in any row, and the coefficient carrying it fails conventional significance ($p \geq 0.11$ in all four). Education is doing essentially all the predictive work.

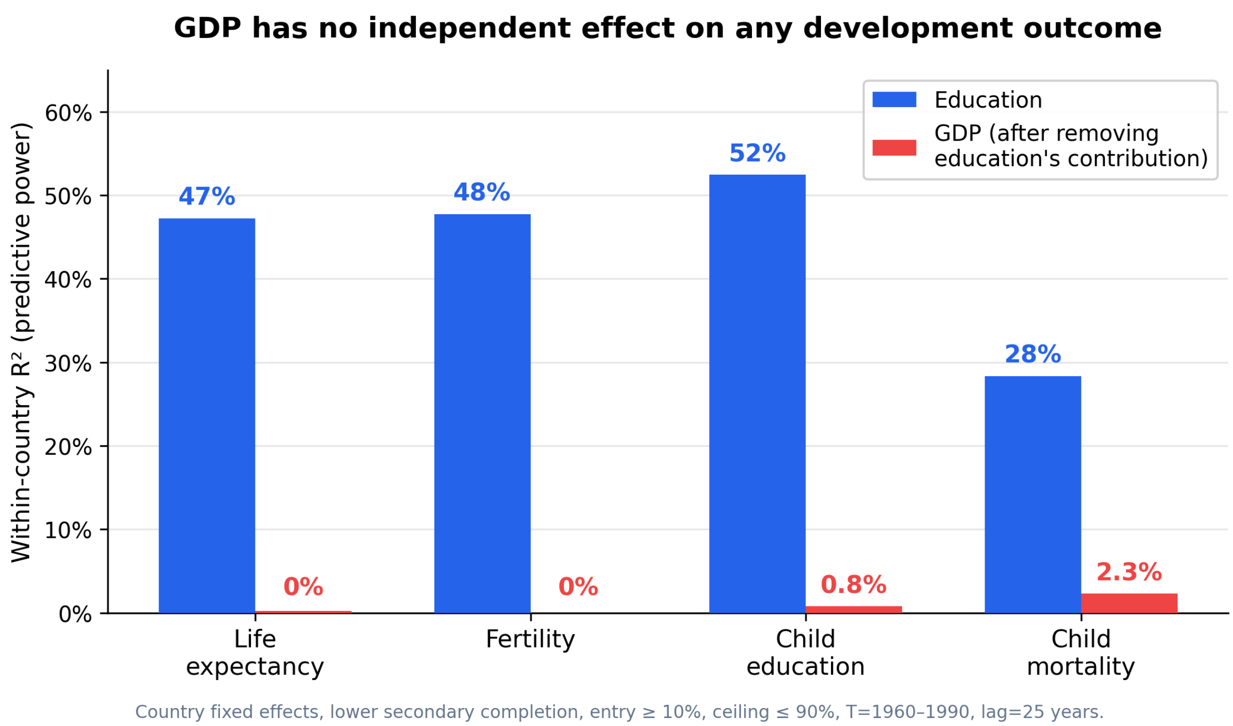

Figure 8: How much does education vs. log GDP per capita predict development one generation forward? Blue bars: education's predictive power. Red bars: GDP's predictive power *after stripping out what education produced*. Log GDP per capita predicts nothing on its own. Entry-cohort design (entry ≥ 10%, ceiling ≤ 90%), country fixed effects, lower secondary completion, T=1960–1990, lag=25 years.

Child mortality: one bounded exception. Child mortality is the only outcome where residualised log GDP per capita shows a marginal independent effect. Before 2000, leftover GDP explains 0.9% (not statistically significant); after 2000, 2.7% (p=0.03). The effect appears only at intermediate education levels and vanishes below 20% completion in any period. The paper's claim does not require every residual to be education: GDP, through external provision (GAVI founded 2000, vaccine delivery through global supply chains), can produce a small independent effect on child mortality in the post-2000 window where such provision scaled. Whether the surviving signal is genuinely GDP-driven or is global education funded by donor surplus — vaccines developed, delivered, and paid for by already-developed populations — the channel is fragile by construction: outer-radius effects run on talk, not self-interest (Section 5.2, outer radius).

Composition by education level

Primary education drives fertility decline; deeper education drives life expectancy gains.

The *composition* of the effect varies by level. Primary education is the strongest predictor of fertility decline. Its coefficient on TFR at $T+25$ is larger than lower or upper secondary's. The within-country fit is tightest at this level (R^2=0.65 vs 0.48 for lower secondary, 0.32 for upper secondary).[5] The first educated generation — the categorical jump from no schooling to literacy — is where fertility falls hardest, universally, across every country in the dataset regardless of culture, religion, or economic structure. Caldwell (1976) predicts fertility decline when children shift from economic assets to costs; if correct, the steepest decline should track economic modernisation, not primary education. It tracks primary education, in agrarian and industrial contexts alike. The mechanism is not economic calculation; it is cognitive capacity.

Life expectancy gains, by contrast, accumulate with deeper education: the upper-secondary coefficient dominates, and the within-country fit is corre-

[5]The relationship survives restriction to sub-Saharan Africa, where the individual-level cross-sectional literature has long reported null or weak primary effects on fertility in early-transition settings. Within SSA (44 countries, 308 country-period observations, country fixed effects), primary $\rightarrow$ TFR$(T+25)$ yields R^2=0.58 and $\hat{\beta}=-0.046$ (SE 0.005); the 10 SSA countries that have crossed TFR<3.65 cross at a median 77% primary completion (global 79%). The cohort-level relationship is intact in the population the cross-sectional literature flags.

spondingly higher than at primary (R^2=0.48 vs 0.41). This is the structure that Lutz & Kebede (2018) identified when redrawing the income–mortality curve with education on the x-axis: the fertility transition is triggered by the first level, while the education–health relationship tightens at higher levels.

Residualized GDP R^2 never exceeds 0.03 for life expectancy and is zero for fertility, at every level.

The education–life expectancy relationship shows no diminishing returns — linear from primary through tertiary (Lutz & Kebede 2018). Among 74 countries with lower-secondary completion above 85% in 2010, college completion carries a within-group correlation of r=0.45 with life expectancy — a 5.7-year gradient between lowest and highest college-completion quartiles among countries that have already solved the lower-secondary problem.

Robustness and additional tests

Residualized GDP R^2 stays below 0.003 for life expectancy and fertility across all lag lengths tested; under-5 mortality reaches 0.019 (Table A8). The confidence intervals do not overlap: education R^2 [0.33, 0.59] vs residualized GDP R^2 [0.00, 0.03] (500-replication bootstrap). An alternative estimation method designed specifically to correct for bias in short panels (Anderson-Hsiao instrumental variables) produces results within 3–10% of the main estimates. Results are unchanged when time-varying institutional quality scores (Polity5 polity2) are added (Appendix 13.4). Female education produces stronger results than the aggregate on every outcome.

The reverse-direction test asks a simple question: does parental-generation log GDP per capita predict child education independently, or only because richer countries are more educated? Parental-generation GDP alone predicts child education moderately (β=14.4, R^2=0.263). But add parental education to the same regression and GDP's coefficient collapses by 80% to β=2.9 (p=0.16), adding almost nothing (R^2=0.006) beyond education alone. Education's coefficient barely moves (β=0.489 from 0.540) while GDP's collapses. GDP transmits through education, not vice versa.

The Asian Financial Crisis (1997–98) tests Prediction 3: GDP collapsed across five countries (Indonesia −14.5%, Thailand −8.8%, Malaysia −9.6%, Philippines −3.0%, Korea −6%). Lower secondary completion continued uninterrupted in all five; Indonesia gained 5.4pp while losing 14.5% of GDP. GDP was destroyed; the educated population was not; development continued.

Once primary education is included, initial fertility does not predict subsequent fertility one generation forward (Table 13, Panel A, TFR row: initial outcome β=-0.039, $p = 0.37$). Demographic momentum does not drive the transition — educated women's decisions do.

A permutation null rules out the chance-correlation alternative. Reshuffling the parent–child match across countries (200 iterations) places the real $\beta = 0.483$ at least 53 SDs above either null — the within-year null, which preserves global temporal trends, and the full shuffle, which breaks every systematic link. Common trends, panel autocorrelation, and serial correlation cannot generate the observed coefficient (Appendix A.2).

9.9 Universality Across Subsamples (Prediction 5)

Prediction 5 says the mechanism is universal: it appears in every region, every era, and every income band, not only in the East Asian cases on which much of the growth literature has concentrated. Two tests, panel and residual.

Panel, fifteen subsamples. Re-estimating Column 1 of Table 9 on six regions, two eras, and three within-sample GDP terciles produces fifteen positive coefficients, every one significant at $p < 0.01$ (Table 19). Every β exceeds unity; Sub-Saharan Africa, where most of the remaining 20% of humanity sits, tracks the full-sample headline. The era and tercile splits rule out post-Cold-War and income-band artefacts.

Table 19: Subgroup robustness of the headline specification.

Subsample	β	SE	R^2	N	Countries
Full (headline)	**1.376*****	0.083	0.699	629	105
By region					
Sub-Saharan Africa	1.220***	0.109	0.702	301	40
Middle East & N. Africa	1.145***	0.176	0.745	61	12
South Asia	2.444***	0.515	0.720	57	8

Subsample	β	SE	R^2	N	Countries
East Asia & Pacific	1.491***	0.097	0.823	84	17
Latin America & Caribbean	1.272***	0.091	0.829	112	23
Europe & North America	2.057***	0.156	0.762	14	5
By child-cohort period					
Pre-1990 (1975–1985)	1.364***	0.172	0.575	279	99
1990–2015	1.332***	0.116	0.621	338	72
By GDP tercile (within sample)					
Low	1.295***	0.153	0.663	208	33
Middle	1.276***	0.091	0.761	205	44
High	1.393***	0.137	0.667	205	53

Notes: Column 1 of Table 9 (child lower-sec $\sim$ parental, country FE) re-estimated on regional, temporal, and GDP-tier subsamples of the 629-obs / 105-country active-expansion sample. Country-clustered standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All fifteen subsamples reject the null at $p < 0.01$; every β exceeds unity. Headline row repeats Table 9 Column 1 for reference.

Residuals, eight over-performers. Some countries are educating children far beyond what their parents' generation would predict. The 2015 country-FE residuals from the headline specification identify eight: they are uniformly poor.

Table 21: Education over-performers (2015 FE residuals).

Country	FE residual above baseline	GDP per capita (2015)
Maldives	+34.9 pp	$9,645
Cape Verde	+26.3 pp	$3,415
Bhutan	+26.1 pp	$2,954
Tunisia	+25.5 pp	$4,015
Nepal	+17.8 pp	$876
Viet Nam	+16.0 pp	$2,578
Bangladesh	+15.8 pp	$1,224
India	+14.1 pp	$1,584

Notes: Country fixed-effects residuals from the headline specification at 2015. The residual measures how far each country exceeded the level its own historical trend predicted.

Nepal at $876 (2015), Bangladesh at $1,224, Vietnam at $2,578. Income is not the prerequisite for educational over-performance.

A full robustness battery — alternative cutoffs (full 10-row sweep), alternative lag lengths (15–30 years), specification robustness (PPML, log outcomes, period-length aggregation, balanced subpanel, within-year cross-cohort, Wooldridge strict-exogeneity), data-source robustness (Barro-Lee replication, post-1970 subset), and development-threshold robustness — appears in Appendix A.3. None reverses the ordering, and no check reduces the education/GDP R^2 ratio below 2.0×.

10. Hollow Education: What the Soviet Anomaly Tests

The paper's empirical engine is lower-secondary completion in the WCDE v3 reconstruction (Lutz et al. 2018), chosen because it reaches back to 1875 on a consistent definition and passes the phenotype tests laid out in Chapter 9. On the 185-country panel, WCDE disagrees meaningfully with Barro-Lee v3.0 (Lee & Lee 2016) for exactly fifteen countries: the Soviet republics. Reported lower-secondary completion in 1970 places Kazakhstan at 94%, Latvia at 99%, Russia at 95%; Barro-Lee for the same year reports 49%, 57%, and 73%. Fifteen republics, a mean gap near 40 percentage points, all in the same direction.

But the gap is not uniform across the fifteen. 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 2.6 to 4.0σ above what their reported lower-secondary completion predicts. The split is categorical, not gradient: Belarus at 676 km west of Moscow shows a U5MR residual of $+0.43\sigma$; Georgia at 1,646 km south 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. *The phenotype test is incorruptible because under-5 mortality requires an actually-literate young mother to fall, not a credential-holding one. The U5MR residuals reveal which Soviet education numbers were real schooling and which were the metropole reporting on populations it could not see.*

In the childhood frame, hollow education is exactly this: the eighteen-year window was filled with school attendance but not loaded with the cognitive architecture the schooling was supposed to install. The credential measures the door the child walked through; only the phenotype measures what entered the window.

The diagnostic yields three payoffs. First, a credibility case for the panel: every headline result survives exclusion of the fifteen republics, and in fact strengthens. 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. **Hollow education is not education — but exposure still matters for TFR** (the partition established in §9.4).

10.1 The Anomaly

Of the 154 countries that have crossed the developmental threshold by 2022, 28 were post-socialist — 15 USSR republics plus 13 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 18 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 52 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 phenotypic response by $T{+}25$. Where it does not, either the mechanism is wrong or the reported education is wrong. The mechanism passes for 171 of 185 countries in the panel; the 14 where it fails (Uzbekistan drops out of the panel on data coverage) share a common origin: one statistical office.

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 is not consistent with the family structure and labor patterns documented in ethnographic sources for 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 at $T+25$. Where the predicted output does not appear, the input was not real. Figure 9 shows the test for under-5 mortality.

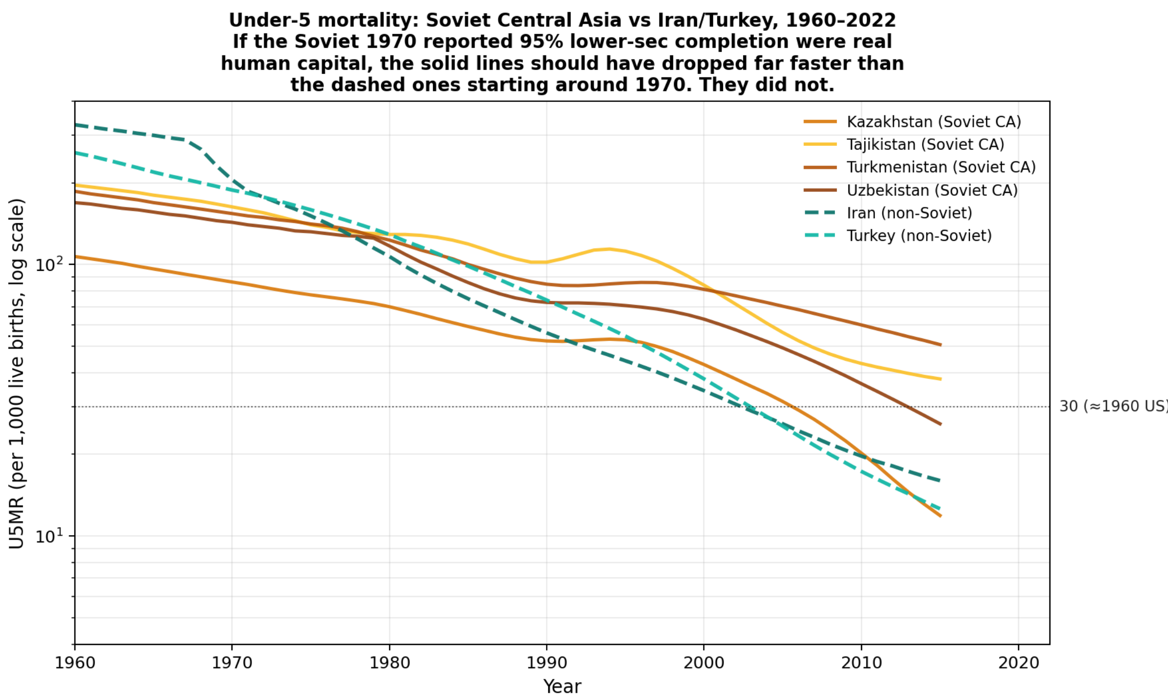

Figure 9: Under-5 mortality trajectories, 1960–2010. Central Asian and Slavic USSR republics (colored lines) start 1960 with large U5MR head starts over Iran (327 per 1,000) and Turkey (258). By 2010, Iran (20) has overtaken every Central Asian country except Kazakhstan. Under the mechanism, the country 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 except Kazakhstan. If the 1970 Soviet 95% had been real, Central Asia should have dominated Iran and Turkey permanently. 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.

10.4 The Moscow Meridian

The fifteen Soviet republics split categorically at Moscow's longitude. Table 23 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 23: 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	$+1.09$
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	$+4.00$

Republic	East-of-Moscow (km)	LE residual (σ)	log U5MR residual (σ)
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.

10.5 Barro-Lee's Partial Correction

Barro-Lee v3.0 is an independent reconstruction of the same series using different source surveys and interpolation. It is systematically more conservative than WCDE on Soviet-era reporting. The Goskomstat hypothesis predicts that switching measures should reduce the phenotype-fit bias, and it does — unevenly. Table 25 reports the mean phenotype residual (actual outcome minus model prediction given reported education) by subgroup, under a pooled OLS fit with year fixed effects trained on the complement of the contested cohort. A positive TFR residual means the subgroup had more children than its reported education predicted; a negative LE residual means it lived shorter; a positive log-U5MR residual means it lost more children. All three signs point the same direction: reported education was too high.

Table 25: Phenotype residual by subgroup, Barro-Lee reached-secondary (1960–1990).

Subgroup	n	TFR (SDs)	LE (SDs)	log U5MR (SDs)
Central Asia + Caucasus (USSR)	16	$+1.15$	-1.79	$+2.78$
Slavic west (USSR)	8	$+0.28$	-1.93	$+1.90$
Baltics (USSR)	12	-0.01	-1.14	$+0.71$
Warsaw Pact (non-USSR)	20	-0.42	-0.49	$+0.25$
Yugoslavia + Albania	16	-0.59	$+0.09$	-0.15

Notes: Bias reported in standard deviations of the non-contested fit. A subgroup "passes" when bias is indistinguishable from zero.

Two things the table tells us. First, Goskomstat specifically: Warsaw Pact states had their own national statistical offices; Yugoslavia ran its own; both subgroups pass the phenotype test under Barro-Lee on all three outcomes. The 15 Soviet republics fail at LE and U5MR; the 8 Central Asian and Caucasus republics fail at TFR as well. The anomaly is not socialism. It is one statistical office.

Second, outcome specificity. Under the corrected measure, TFR fits for the Baltics and the Slavic west — their fertility decisions are in line with their real education. Where they fail — LE and U5MR — they fail in the direction of shorter lives and more child deaths than the reported education should allow. The outcome partition is the first hint of the thesis this chapter converges on.

10.6 The Cohort Shadow

Figure 10 decomposes the USSR residuals by year. The LE panel carries the diagnostic: the residual plateaus in the 1970s near -8 years, widens through 2000 to a peak of -8.5 years, then begins to decay. This is not a 1990s transition-crisis spike, which would show as a narrow peak around 1995 with small residuals before and after. It is a cohort shadow.

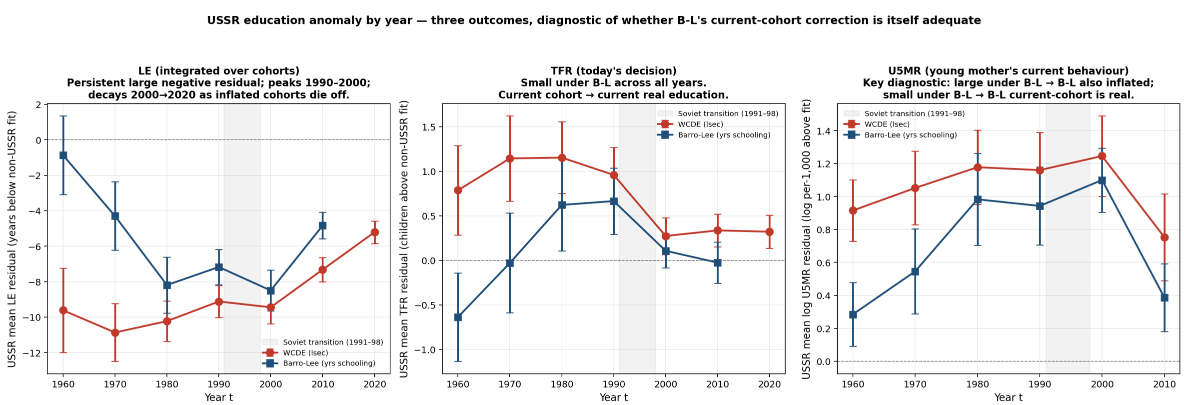

Figure 10: USSR mean phenotype residual by year, by measure. LE (left) shows a broad 1970–2000 plateau and decay after 2000 — the signature of inflated 1960–80 credential cohorts aging into peak-mortality years. TFR (middle) approaches zero under Barro-Lee at every year: fertility is a current decision, and Barro-Lee's current-cohort correction catches it. log U5MR (right) remains large under Barro-Lee for 1980–2000, because a young mother's health decisions depend on her real cognitive skill, not on her reported credential. *(Source: WCDE v3, Barro-Lee v3.0, World Bank WDI.)*

The mechanics: LE in year t integrates mortality experience across all cohorts alive in year t. The 1960–1980 Soviet cohorts were reported as carrying $\sim$90% lower-secondary completion; under the phenotype test their real education looks more like 60–70%. Those cohorts, born 1935–55, reached peak-mortality ages (70–85) roughly 2005–2040. If their inflated-credential status did not carry real cognitive resources — information processing, health literacy, planning capacity — their mortality experience reflects real schooling, not reported schooling. The predicted signature is a large negative LE residual peaking around 2000–2015 and decaying through 2030. This is what we see.

TFR carries no such shadow under Barro-Lee. A 2000 fertility decision is made by a woman aged 20–30, educated in the 1980s–90s under progressively less-hollow Soviet and post-Soviet secondary systems, with Barro-Lee's current-cohort correction catching her actual schooling. Her fertility is in line with her real education by the time she makes the decision. This is why TFR corrects and LE does not.

U5MR tells the sharpest story. Under Barro-Lee reached-secondary the mean USSR log-U5MR residual is $+0.29$ in 1960, $+0.98$ in 1980, $+1.10$ in 2000 ($t = 5.7$), $+0.39$ in 2010. Peak U5MR bias is in 1980–2000 — the decades when the 1960–80 inflated-credential cohorts were themselves mothers of young children. The paper's claim on U5MR has always been that it loads on whether the rungs of the stack actually installed in the mother — germ theory, dose-response, the hygiene chain — not on her completion certificate (Section 4.2). The Soviet case makes the claim visible: where the certificate said "secondary-completed" but the rungs were not there, the child died anyway.

10.7 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 10.8). 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 11 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. **Great-grandparent primary completion in 1950 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 great-grandparent generation. What the cross-section rules out is the alternative that lag 60 collapses: the R^2 stays at 0.489 against 0.469 for current schooling. The R^2 is flat where it should not be flat if HLO measured contemporary schools.

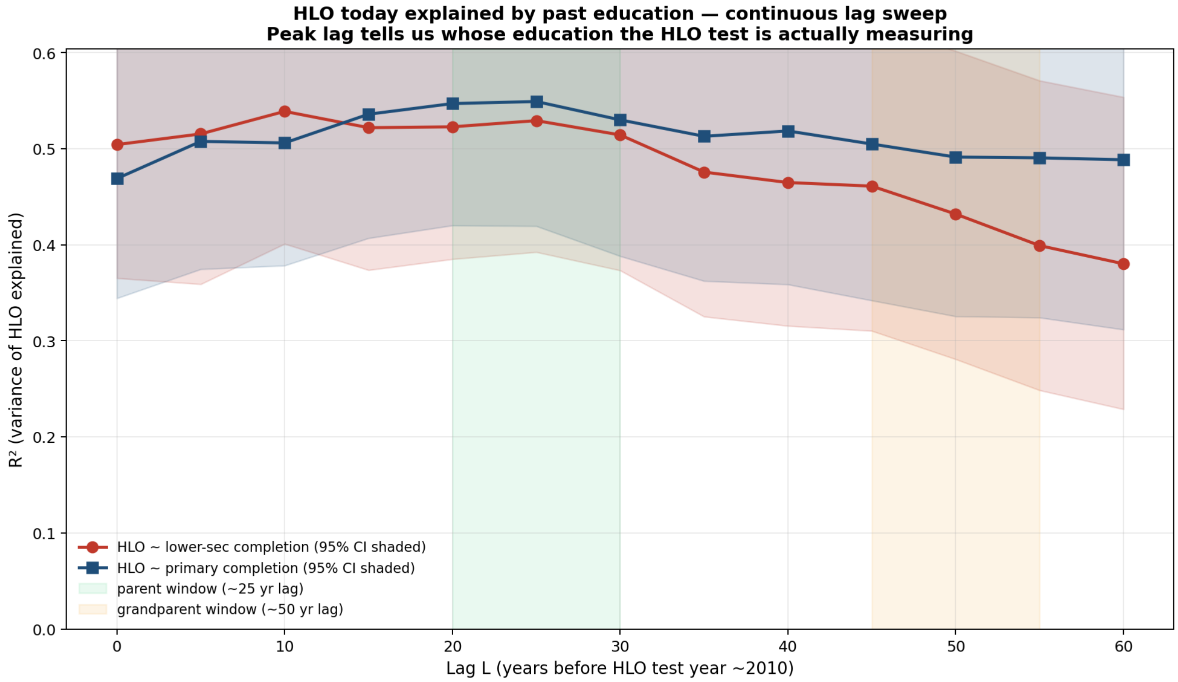

Figure 11: HLO secondary today regressed on WCDE completion at year $2010 - L$, for $L = 0, 5, \ldots, 60$. Shaded bands are 95% country-resample bootstrap CIs (2,000 draws). The curve is broad and flat across the full lag sweep; primary completion at lag 60 (1950, the great-grandparent generation) explains $R^2 = 0.489$ of today's variance, against $R^2 = 0.469$ for current schooling. Current schooling is a worse predictor than either parental or great-grandparent 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 (Section 9.5) and by the matched-pair natural experiments (Korea/Philippines, Cuba/Cambodia, North/South Korea) that hold institutional quality roughly constant and vary only the schooling regime. 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: TFR runs on the first rungs of the stack; U5MR and LE on composed stacks (Section 4.2). Section 9.4 reports the full quantitative partition.

10.8 The Panel Survives Exclusion, Strengthened

The Table 9 headline specification reruns with the 15 USSR republics excluded. The parental-education coefficient rises from 0.483 ($n = 1665$, 185 countries) to 0.538 ($n = 1539$, 171 countries); within-R^2 rises from 0.457 to 0.507. The headline LE forward-prediction R^2 rises from 0.350 to 0.366. The headline TFR forward-prediction R^2 rises from 0.439 to 0.444. The headline log-U5MR forward-prediction R^2 rises from 0.635 to 0.659. Every coefficient of interest moves in the same direction: upward.

This is not what survivorship looks like for a contaminated claim. It is what survivorship looks like for a claim whose signal was being attenuated by the contaminated subgroup. USSR inclusion was not producing the paper's headline results. It was dampening them. Removing the dampening factor returns the cleaner estimate of the mechanism. Appendix A.3.6 reports the full table with all specifications; the 171-country row joins WCDE, Barro-Lee, and post-1970 as one more robustness row confirming the direction and magnitude of the headline.

One sentence summarizes the chapter. My mechanism identifies where the Soviet data was wrong; the Soviet data, once excluded, confirms the mechanism more sharply. It is the kind of agreement between theory and data that one does not get by accident.

11. The Dark Parallel

The positive case has been stated: 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.

11.1 Capacity, Not Virtue

The mechanism produces capacity for coordinated collective action (Prediction 7). 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.

11.2 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 inter-group 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.

11.3 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.

11.4 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.

12. 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 9.8). 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 every war since 1990, and completing the transition retires it.

12.1 The Asymmetry Reverses

The European window (Section 11.3) 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 11.1).

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 12.4) 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.

12.2 The Malthusian Engine Runs Out of Fuel

The Malthusian engine (Section 11.2) 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 11.3) was the window during which one region had arrived at the substrate and the rest had not. Convergence closes that window.

12.3 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).

12.4 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 11.3). 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.

The post-2000 drop in global under-five mortality, driven substantially by vaccine delivery from the already-developed world (Section 9.8), is the most visible signature: populations that previously let foreign children die now keep them alive. The signature is small and fragile by construction — bounded R^2 in §9.8, an outer-radius effect running on talk not self-interest — but it is the latest stage of a moral circle that has been expanding for centuries wherever literate CT loads.

Not all Europeans were in-group to one another in 1700, and the same was true in every prior imperial age (§11.3): 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.

12.5 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 11.1) — the species will remain what it has always been, with its hierarchies and its arbitrary out-group orientations (Section 11.2) 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.

13. Difficulties on the Theory

Each major alternative account in development identifies something real; each stops one level short of the substrate that makes the thing it identifies possible. Sen and Drèze are the exception: the capabilities frame is this paper's frame, read one layer deeper, and Section 13.2 completes rather than displaces them.

13.1 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 13.8) — prices the loaded labour rather than naming what loaded it. The capabilities frame — Sen (Section 13.2) — 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.

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.

13.2 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.

13.3 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 9.8). 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 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.

13.4 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.

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.

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 4.8 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 still expanding, mean expansion rates were 1.06 pp/yr in 1950–1975, 0.86 in 1975–2000, and 0.94 in 2000–2015. If institutions drove educational expansion, rates should have risen as institutions matured. They did not.

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 (Table A5). 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*: 76% 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 Tables A7–A9.

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 institutional literature has been measuring the joint effect of education and institutions and reading it as institutions alone.

Table A5: Regime type and education gain rate.

Regime lag	R^2(polity2)	Polity coef	Auto mean	Demo mean	n intervals
Concurrent (0 yr)	0.0015	+0.056	8.5	9.5	1,782
15-year lag	0.0067	+0.122	8.1	10.0	1,574
20-year lag	0.0050	+0.106	8.1	10.3	1,472

Notes: Polity5 time-varying regime scores (polity2, -10 to $+10$) merged with WCDE lower secondary completion (both sexes, age 20–24); 160 countries, 1950–2020. Education at age 20–24 reflects schooling 15–20 years earlier; lagged specifications attribute the education outcome to the regime that made the decision. Gain rate in percentage points per decade. Polity coef: the slope of the democracy score on gain rate (positive = more democratic regimes gain faster). R^2: share of gain-rate variation explained by democracy score alone — less than 1% at every lag. Auto = strong autocracies (polity2 ≤ -6); Demo = strong democracies (polity2 $\geq +6$). Both invest at roughly the same rate.

The colonial test 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 ($r = 0.62$): settler mortality predicts both because where Protestant settlers survived they built both schools and inclusive institutions. Adding coloniser religion on top of education does not raise the fit.

Table A7: 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

Predictor	R^2 (log GDP 2020)
Colonial-era education (1900 cohort lower sec)	0.248
Coloniser religion (Protestant = 1)	0.055
Education 1950 + coloniser religion	0.521

Notes: AJR's 64-country base sample (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. Adding coloniser religion to education raises R^2 from 0.518 to 0.521: religion predicts GDP almost entirely through its effect on education. Education at WCDE v3, GDP at World Bank WDI (constant 2015 USD), avexpr and logem4 from AJR (2001) replication archive.

The institutionalist literature identifies the effect of institutions on income using settler mortality or coloniser religion as instruments. Running the same instrumental-variable design on AJR's own institutional measure (*avexpr*) shows the instrument predicts education at first-stage $F = 10.71$ and *avexpr* institutional quality at $F = 9.61$. Both channels run through the same colonial-era variation, and the IV identification cannot decompose them.

Table A9: 2SLS instrumental-variable contest – education vs. AJR's avexpr.

	First stage		Second stage (log GDP 2020)	
Endogenous variable	F-stat	Strength	$\hat{\beta}$	t
Education 1950 (lower sec)	10.71	Strong	+0.043	2.56
avexpr (institutions)	9.61	Marginal	+0.629	3.01
Structural diagnostic: avexpr as instrument for education				
Education 1950 | *avexpr*	37.13	Strong	+0.084	6.74

Notes: Top panel: coloniser religion (Protestant = 1) as the excluded instrument; AJR's 64-country base sample. Education clears the Stock & Yogo 2005 strong-instrument threshold ($F > 10$); avexpr sits just below. Bottom panel: AJR's own avexpr as instrument for education — $F = 37.13$ is three times the threshold. The same colonial-era variation drives both education and institutions, so the instrument cannot identify either channel cleanly. The exclusion restriction, on which AJR's identification rests, is empirically violated by AJR's own data.

13.5 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, awarded the 2021 economics Nobel for instrumental variables, regression discontinuity, and difference-in-differences) and the randomised-trial movement in development economics (Banerjee, Duflo, and Kremer, awarded the 2019 prize for 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. These methods, applied to the channel, see the channel: the Callaway–Sant'Anna event study reported in Section A.1 is a credibility-revolution estimator, and on the clean comparisons it confirms the mechanism. 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 effectively identical colonial education bases and identical per-capita income; 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. Niger and Korea today are a sixty-year longitudinal experiment in the same species under near-identical 1950 baseline poverty, with the only systematic difference being whether the state loaded the channel.

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 15 USSR republics 1960–1990 had reported lower-secondary completion that placed them with the most-educated populations in the world; the channel was hollow, and the four phenotype residuals (LE, TFR, U5MR, contemporary HLO test scores) all show the gap. In the language of randomised trials, this is a manipulation-check failure: where the treatment is administratively imposed without actually loading the variable the theory claims is causal, the outcome does not follow. Excluding the USSR republics strengthens every headline in this paper. The channel is what matters, not the schooling reported on paper.

The methodology canonised by the 2019 and 2021 prizes is appropriate for the questions those prizes' 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. 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.

13.6 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 horizon 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.

13.7 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; the biology, the natural experiments, the country histories, the panel, and the invisibility argument are what I add beneath it.

13.8 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 13.3) 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.

13.9 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. Years in school predict life expectancy, under-five mortality, and fertility at effect sizes an order of magnitude larger than test-score variation across countries at the same mean years of schooling. 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 13.1) — 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 9.9).

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.

14. Why the Loading in Childhood Is Invisible

Childhood is visible. The loading happening inside it is not. Adults whose own childhoods were loaded with formal schooling cannot easily categorise that loading as fundamental — it is the medium they think through. Households whose children's childhoods are not loaded cannot easily imagine what the loading does, because the imagination required to evaluate it is itself part of what loading installs. The invisibility is the structural reason every framework that has reached for the mechanism has stopped where the educated frame stops. I make the invisibility precise here and close by turning it on the present synthesis.

14.1 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-five 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.

14.2 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 147 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.

14.3 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.

14.4 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 Preston Curve 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 9.8). 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.

14.5 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.

14.6 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 of the previous section. 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 (Section 9), not its provenance.

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 ($\beta = -0.004$ years per percentage point, $p = 0.86$, $n = 644$; `scripts/le_independent_of_u5mr.py`). 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 1.37% reduction in female adult mortality at $T{+}25$; controlling for under-five mortality at $T{+}25$, the effect attenuates but remains significant (-0.37% per pp, $p = 0.001$, $n = 644$; `scripts/mothers_own_longevity.py`). 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 at a 25-year lag. The fertility channel is $\beta_{\text{TFR}} = -0.035$ births per percentage point at the same lag, 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 25 years later as the affected cohort enters parenthood; the window therefore runs from $T{+}25$ 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 25 million lives at stake through 2025: 17 million under-five deaths averted via the mortality channel, plus 8 million fertility-channel lives. Earlier anchors yield correspondingly larger totals: $T = 1980$ gives 88 million (55 million mortality, 33 million fertility); $T = 1970$ gives 178 million (108 million mortality, 70 million fertility). The sensitivity envelope at $T = 1990$ is shown in Table 11: each successive scenario relaxes the pace constraint, raising the lives at stake from 25 million (Korea pace) to 71 million (9-year biological floor).

Scenario at $T = 1990$	Mortality	Fertility	Kids not born	Lives at stake
Korea pace (2.13 pp/yr)	17 M	8 M	69 M	25 M
15-year linear ramp to 95%	32 M	23 M	174 M	55 M
9-year biological floor	37 M	34 M	264 M	71 M

Table 11: 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. 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 truncates the predicted effect: most of the cost of singular-priority laggards from $T = 1990$ materializes between 2025 and 2050, as cohorts schooled 2015–2030 enter parenthood; the headline is therefore the floor of the prediction 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, by the bad-control argument (Section 9.1).

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. Singular priority is the only intervention that alters the metabolism's input; every competing priority leaves it unchanged.

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 to universal upper-secondary completion by 2030 (UN 2015). That commitment is what singular priority asks. From 2025 baselines, with five years remaining, the median country still below threshold would have to expand upper-secondary completion at 8.3 percentage points per year — nearly four times Korea's WCDE-record sustained pace of 2.13 pp/yr (`scripts/sdg4_implied_pace.py`). The target is correct. What is absent is the singular priority required to meet it.

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

Through the full dependency window. Nine years loads lower-secondary; twelve loads upper-secondary. Plan for twelve.

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.

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 — or its absence — is in primary school this year. The window admits exactly what is loaded into it, and it closes one cohort at a time.

No leader 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; with-

out 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. All scripts, data, and tests are published with the paper, and the verifier (`scripts/verify_humanity.py`) is the single entry point that traces every cited number to its producing script and check-in JSON.

A. Robustness

This appendix is optional for the policy reader. None of the checks below is load-bearing for any claim in the body of the paper. The biological identification runs through Chapters 2–6, the natural experiments (Chapter 7), and the hollow-education falsification test (Chapter 10); the body's argument stands on those alone.

This appendix contains the econometric checks referenced in Sections 9 and 7. The body reports the headline specifications; the checks below confirm that staggered adoption, alternative cutoffs, alternative lag lengths, alternative estimators, alternative data sources, and the Colonial Test first-stage diagnostics leave the ordering and the coefficient signs unchanged. They are placed here for readers who work in traditions that treat identification as the load-bearing element of an empirical argument.

A.1 Two-Way Fixed Effects and Staggered-Adoption Diagnostics

This appendix is for the panel-econometrics reader. The identification of the paper's claim is biological (Chapters 2–6) and runs through the natural experiments (Chapter 7) and the hollow-education falsification test (Chapter 10). The diagnostics below are not load-bearing for any claim in the body. They are the demonstration that when the discipline's tools are applied to the channel, the tools see the channel where the channel is loading and report sample-composition signatures where the channel is saturated or not yet loading.

Adding year fixed effects to the headline specification collapses the full-panel parental coefficient from 0.483 to 0.083. Table A1 documents the collapse across four parental-completion subsamples. The paragraphs below explain why the collapse is what the mechanism predicts: pool country-years where the channel is actively loading (active expansion) with country-years where it is saturated (above the ceiling) or not yet loading (pre-expansion), and the diagnostic detects the pooling. Restrict to the operative subsample and the coefficient returns ($\beta = 0.739$ to 1.019 across columns 1–3 of Table A1, all $p < 0.01$). The Goodman-Bacon decomposition gives the weights and the Callaway–Sant'Anna estimator confirms the effect on clean comparisons.

Table A1: Two-way fixed effects (country + year) – child lower secondary completion on parental education.

Model	Parental edu β	R^2 (within)	n	Countries
(1) FE + year: parent (< 30%)	**0.739*****	**0.154**	783	137
(2) FE + year: parent (< 20%)	1.032***	0.145	600	118
(3) FE + year: parent (< 10%)	1.019***	0.067	358	85

Model	Parental edu β	R^2 (within)	n	Countries
(4) FE + year: parent (all)	0.083	0.009	1,665	185

Notes: Two-way (country + year) fixed effects; R^2 is the partial within-R^2 (incremental variance explained beyond country + year FEs). Country-clustered standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Model (4) does not reach conventional significance ($p = 0.07$); the prose below explains why this is the mechanical signature of two-way FE under staggered treatment, not evidence against the mechanism.

Why Model 4 shows β=0.083 and why the mechanism predicts it. A theory of channel loading predicts that within-country movement is informative only where the channel is moving. Country-years where the channel is saturated (parental completion near the ceiling) contribute variance with no movement to attribute. Country-years before the channel begins loading (parental completion near zero) contribute variance from a different generative process — the state has not yet committed. Pooling all three regimes and then absorbing the global trend with a year fixed effect is asking the panel a question the mechanism does not predict an answer to. Restrict to the operative subsample where the channel is loading (Models 1–3) and the coefficient survives at $\beta = 0.739$ to 1.019, all $p < 0.01$. This is what a correctly-specified channel should produce. A diagnostic that gave a clean estimate across saturated and pre-expansion regions would be detecting something other than the mechanism. The full-panel $\beta = 0.083$ is a sample-composition signature, not a refutation; the mechanism is visible exactly where the mechanism is operating, and not where it is not.

The decline is the mechanical signature of a now well-documented two-way-FE pathology. When treatment is staggered — different countries begin expanding education in different decades — the 2WFE estimator is a weighted average of *all* pairwise comparisons, including "forbidden" ones that compare newly-expanding countries to countries already at high completion (Goodman-Bacon 2021). A decomposition on my panel gives the weights directly: 7.2 percent from clean not-yet-treated comparisons (weighted β = 11.3), 21.3 percent from treatment-timing pairs (weighted β = 1.1), and 71.5 percent from already-treated countries used as controls for newly-treated ones (weighted β = 9.6 but attenuated toward zero by heterogeneous treatment effects). The

0.083 figure is a weighted average in which three-quarters of the weight sits on comparisons that make no causal sense; it is not an unbiased estimate of anything.

The Callaway–Sant'Anna (2021) estimator — designed explicitly for staggered adoption and limited to the clean comparisons only — recovers the effect. Treating the 10 percent lower secondary completion threshold as the event (the crossing from "no priority" to "active expansion"), crossing the threshold raises child completion by 7.9 percentage points one generation later (clustered SE 1.7, 95 percent CI 4.4–11.0); the effect compounds to +21.4 pp by year 35. Pre-treatment event-study coefficients are indistinguishable from zero (Figure 7).

How the panel and DiD estimators relate. The headline specification (Table 9) is a one-way country-fixed-effects regression of child completion on parental completion. It estimates an average within-country slope: when a country's parental cohort is higher than that country's own time-mean, by how much is its child cohort higher than its own time-mean? The Callaway–Sant'Anna and Goodman-Bacon estimators above answer a related but distinct question: when a country crosses a discrete threshold, how does its outcome trajectory deviate from countries that have not yet crossed? Both families identify the parameter from within-country movement over time; they differ in how that movement is summarised — a constant marginal effect (the panel slope) versus a weighted average of event-time differences (the DiD ATT). They coincide only when the treatment effect is homogeneous and treatment is binary. Under heterogeneous treatment effects and staggered timing, two-way FE weights the within-country movements perversely (the 71.5% already-treated weight above); CS reweights them correctly for the binary case. The continuous-treatment one-way FE I use for the headline is not subject to the 2WFE pathology because it does not absorb the global trend with a year fixed effect, and its parameter is a slope, not a weighted ATT.

I discretise at 10% for the DiD cross-check because that is where the mechanism turns on: below 10%, the state's education system is not yet reaching most households and the home niche has no cohort of educated adults to propagate from. Table A3 sweeps the threshold from 5% to 15% and shows the aggregate ATT is positive and stable across the range; the 10% choice is not load-bearing. The DiD framing is therefore best read as a robustness check on a binary version of the same question — not as the right form of the estimand, which I take to be a continuous dose-response slope because the mechanism (Section 6) is dose-response, not on/off.

Table A3: Callaway-Sant'Anna ATT by entry threshold.

Threshold	Cohorts	Treated countries	ATT (post)	Pre-trend ($e{=}-1$)
≥5%	8	184	+5.67	+0.00
≥7%	7	181	+9.10	+0.00
≥10%	8	175	+7.88	+0.00
≥12%	8	172	+7.59	+0.00
≥15%	8	167	+6.29	+0.00

Notes: Aggregate ATT is the unweighted average over post-treatment (g, t) cells of the Callaway-Sant'Anna group-time estimator (not-yet-treated controls), outcome = child lower secondary completion at $T{+}25$. Pre-trend is the same statistic evaluated at event time $e{=}-1$ (placebo). Same panel as Figure 7 (185 countries, 1950–1990, five-year periods).

What I do claim is that the 83 percent 2WFE collapse is not evidence against education's effect. The pathology is visible in the weights, the clean comparisons give a positive estimate, and the staggered- adoption-appropriate estimator confirms it. I keep the one-way country-FE specification as the headline because it corresponds to the "how much do educated countries transmit to their children" test I set up — not because it happens to produce a larger number.

A.2 Permutation Null Distribution

The checks in this appendix vary the specification or the sample. The permutation test holds both fixed and reshuffles the country-level parent–child match: for each iteration the parental column is permuted, the headline regression (Table 9, column 1) is refit on the same panel (1,665 obs, 185 countries, 1975–2015), and the resulting β is recorded. Two schemes, 200 iterations each.

Within-year shuffle. Permuting parental values among countries within the same year preserves the global temporal distribution of parental completion — the conservative null. The null mean is 0.102 (SD 0.005, range $[0.087, 0.115]$). The real $\beta = 0.483$ sits 77 SDs above. That the within-year null centers well above zero is itself informative: preserving each year's marginal distribution already encodes part of the cross-year expansion. The 77-SD gap says global trends are nowhere near sufficient to produce the headline coefficient.

Full shuffle. Permuting parental values across all country-years breaks every systematic link — the aggressive null. The null mean is 0.001 (SD 0.009, range $[-0.022, 0.021]$). The real $\beta = 0.483$ sits 53 SDs above; no permutation in 200 produced a coefficient remotely close to the observed value.

A.3 Other Robustness Checks

A.3.1 Alternative Cutoffs

Table A4: Robustness to alternative cutoffs.

Cutoff	n	countries	Edu β	Edu R^2	GDP β	GDP R^2	Edu/GDP R^2
<10%	275	58	2.022	0.590	15.535	0.296	2.0×
<20%	469	83	1.704	0.665	15.502	0.258	2.6×
<30%	629	105	1.376	0.699	13.659	0.214	3.3×
<40%	740	112	1.185	0.685	15.469	0.223	3.1×
<50%	829	116	1.053	0.697	17.311	0.247	2.8×
<60%	906	120	0.902	0.663	18.059	0.259	2.6×
<70%	969	125	0.819	0.655	18.232	0.270	2.4×
<80%	1018	128	0.739	0.626	17.839	0.280	2.2×
<90%	1086	137	0.663	0.587	17.584	0.284	2.1×
Full	1665	185	0.483	0.533	15.787	0.245	2.2×

Notes: Fixed effects regressions of child lower secondary completion on parental education (or log GDP per capita) across all cutoffs from 10% to 90%. Country fixed effects; country-clustered standard errors. "Cutoff" restricts the sample to country-years with parental completion below the threshold (active expansion); "Full" is the unrestricted panel (1,665 country-years, 185 countries) — the max-sample estimate on every observation available, not the GDP-merged subset. The within-country education R^2 peaks at the 30% cutoff (0.699) and declines only mildly outside that band; the education/GDP R^2 ratio stays between 2.0× and 3.3× across the full 10%–90% range. No cutoff reverses the ordering. The 30% choice in Table 9 is where active expansion is most informative — not an optimum for the coefficient.

A.3.2 Table 7 on the Common GDP-Merged Sample

The headline Table 13 reports each Panel A row on the maximum sample available for that spec — the max-sample estimate on every observation. Robustness: re-estimate the same four-column stepwise on the common GDP-merged sample ($n = 909$ country-years, 168 countries) where every outcome is non-missing. Coefficients move only marginally between the two samples, confirming that the headline pattern is not driven by sample composition.

Table A6: Table 7 Panel A on the common 909/168 sample.

Outcome (T+25)	(1) edu only	(2) + initial	(3) + log GDP	(4) + year FE
log GDP per capita	+0.0148***	+0.0117***	+0.0117***	+0.0040*
	(0.0015)	(0.0016)	(0.0016)	(0.0024)
log(LE)	+0.0030***	+0.0017***	+0.0021***	−0.0005
	(0.0002)	(0.0004)	(0.0004)	(0.0004)
log(TFR) [primary]	−0.0124***	−0.0118***	−0.0116***	−0.0082***
	(0.0009)	(0.0010)	(0.0010)	(0.0012)
log(U5MR)	−0.0279***	−0.0096***	−0.0096***	−0.0014
	(0.0016)	(0.0024)	(0.0025)	(0.0021)
Country FE	Yes	Yes	Yes	Yes
Year FE	—	—	—	Yes
Observations	909	909	909	909
Countries	168	168	168	168

Notes: Same specification as Table 13 Panel A, restricted to the intersection sample where lower-sec, primary, log GDP, log(LE), log(TFR), and log(U5MR) are all non-missing at T and T+25 ($n = 909$, 168 countries). Coefficients move negligibly from the max-sample headline (Table 13), e.g. log(LE) $\beta = 0.0015 \to 0.0017$ (col 2), log(U5MR) $\beta = -0.0080 \to -0.0096$ (col 2). Country fixed effects throughout; country-clustered standard errors in parentheses; ${}^{*}p < 0.10$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$.

A.3.3 Universality of Ceiling Compression

Figure 4 in §9 shows six illustrative country trajectories with a cross-country binned-mean overlay. Below is the formal universality test: every country-window observation in the long-run panel, with an OLS fit (country-clustered standard errors) and a 95% confidence band.

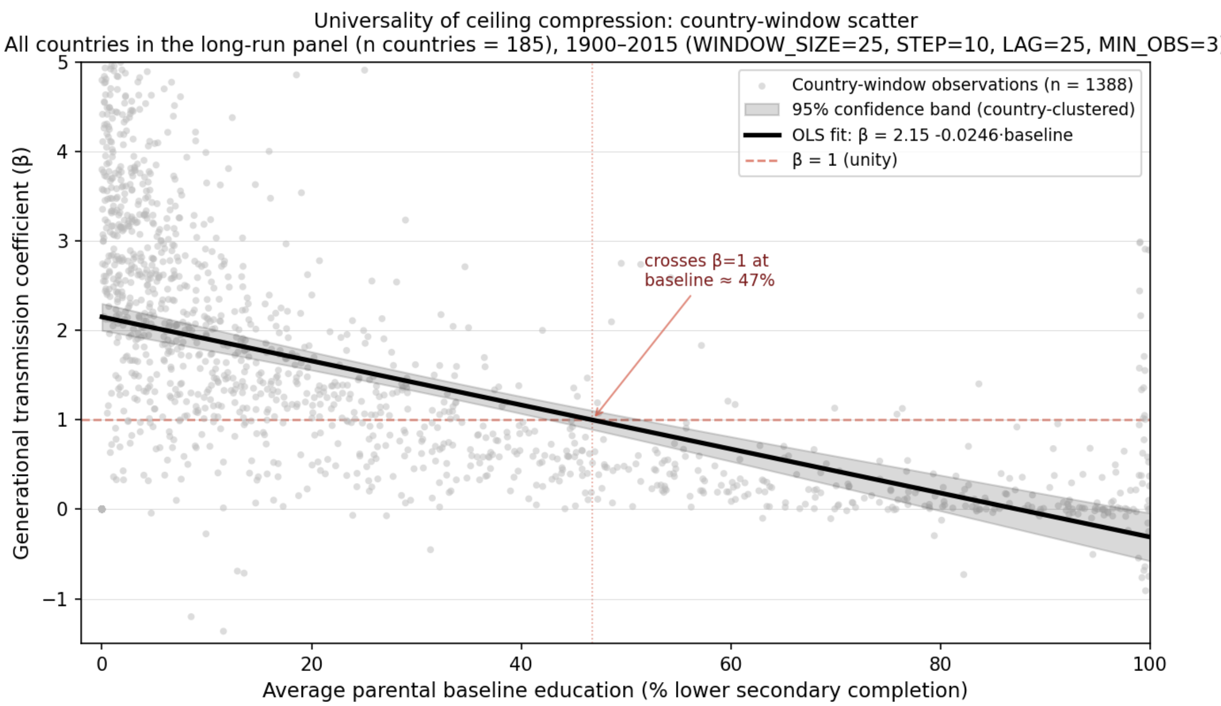

Figure 12: Universality scatter for β_g vs. baseline. Each grey point is one country-window from the long-run panel (`WINDOW_SIZE=25`, `STEP=10`, `LAG=25`, `MIN_OBS=3`; $|\beta| \leq 5$ retained for the regression). Black line: OLS fit $\beta = 2.152 - 0.0246 \cdot$ baseline (country-clustered SE on the slope $= 0.0019$, $p < 0.001$, $n = 1{,}388$, 185 countries); shaded band: 95% confidence band. The fit crosses $\beta = 1$ at parental baseline $\approx 47\%$. The slope is the universality test: it is significantly negative ($t \approx -13$), confirming that β_g declines with baseline across the entire long-run panel, not only in the six illustrative countries of Figure 4.

A.3.4 Alternative Lag Lengths

Table A8: Robustness to alternative lag lengths.

	Life exp.		TFR		U-5 mort.		Child edu	
Lag	Edu	Resid	Edu	Resid	Edu	Resid	Edu	Resid
15 yr	0.455	0.001	0.569	0.000	0.409	0.005	0.725	0.009
20 yr	0.451	0.002	0.509	0.000	0.328	0.013	0.617	0.006
25 yr	0.474	0.002	0.482	0.000	0.280	0.019	0.524	0.005
30 yr	0.457	0.000	0.473	0.001	0.232	0.005	0.461	0.004

Notes: Education R^2 vs. residualised GDP R^2 across lags 15–30 years. Entry-cohort design (entry $\geq$ 10%, ceiling $\leq$ 90%); country fixed effects; lower secondary completion. The $T+25$ specification is the theoretically motivated anchor, not a cherry-picked optimum; the result holds at every lag tested. "Edu" = how much of each country's outcome change is predicted by education. "Resid" = how much is predicted by log GDP per capita after stripping out education's contribution. At every lag tested and every outcome, education explains 23% or more of the variation; leftover GDP never exceeds 1.9%. Sample sizes: $n = 822$ (152 countries) for LE and TFR; $n = 787$ (147 countries) for U-5 mortality; $n = 856$ (157 countries) for child education.

A.3.5 Long-Run Panel and Threshold Robustness

Table A10: Long-run panel countries (1900–2015).

Notes: 28 countries selected on the criterion of self-determined education policy and data quality; colonial-era data for colonised countries reflects colonial investment rather than domestic policy decisions and is excluded. The 28 countries: *East Asia* — Japan, Taiwan, South Korea. *Western Europe* — United Kingdom, France, Germany, Sweden, Norway, Denmark, Finland, Netherlands, Belgium, Switzerland, Austria, Italy, Spain, Portugal. *New World* — United States, Canada, Australia, New Zealand. *Latin America* — Argentina, Chile, Uruguay, Cuba, Costa Rica. *Other* — Sri Lanka (British Ceylon, active colonial education investment, known anomaly), Hong Kong.

Table A11: Development threshold robustness – crossing dates under three specifications.

Case	Loose	Main	Strict	Shift
Cuba	1964	1974	1980	16 yrs
South Korea	1978	1987	1993	15 yrs
Sri Lanka	1975	1993	2005	30 yrs
China	1982	1994	2001	19 yrs
Bangladesh	2005	2014	2019	14 yrs

Notes: World Bank WDI data; Taiwan excluded (not in WDI). The loose spec uses generous thresholds; the main spec uses the 1960 USA values (TFR < 3.65, LE > 69.8); the strict spec tightens to TFR < 2.5 and LE > 72.6 (USA 1975). Every country crosses under every specification, and the ordering never changes. The threshold choice shifts dates by 14–19 years for fast expanders (Korea 2.13 pp/yr, Cuba 2.27 pp/yr, China 1.50 pp/yr) and 14–30 years overall, but no country that crosses the loose spec fails to cross the strict one. Sri Lanka's 30-year shift reflects the civil war holding life expectancy near the threshold while the educational base accumulated beneath it; the binding constraint was disruption, not the threshold definition. Taiwan (1970 loose, 1980 strict; WCDE data) is consistent with Cuba under all specifications but cannot be computed from WDI data.

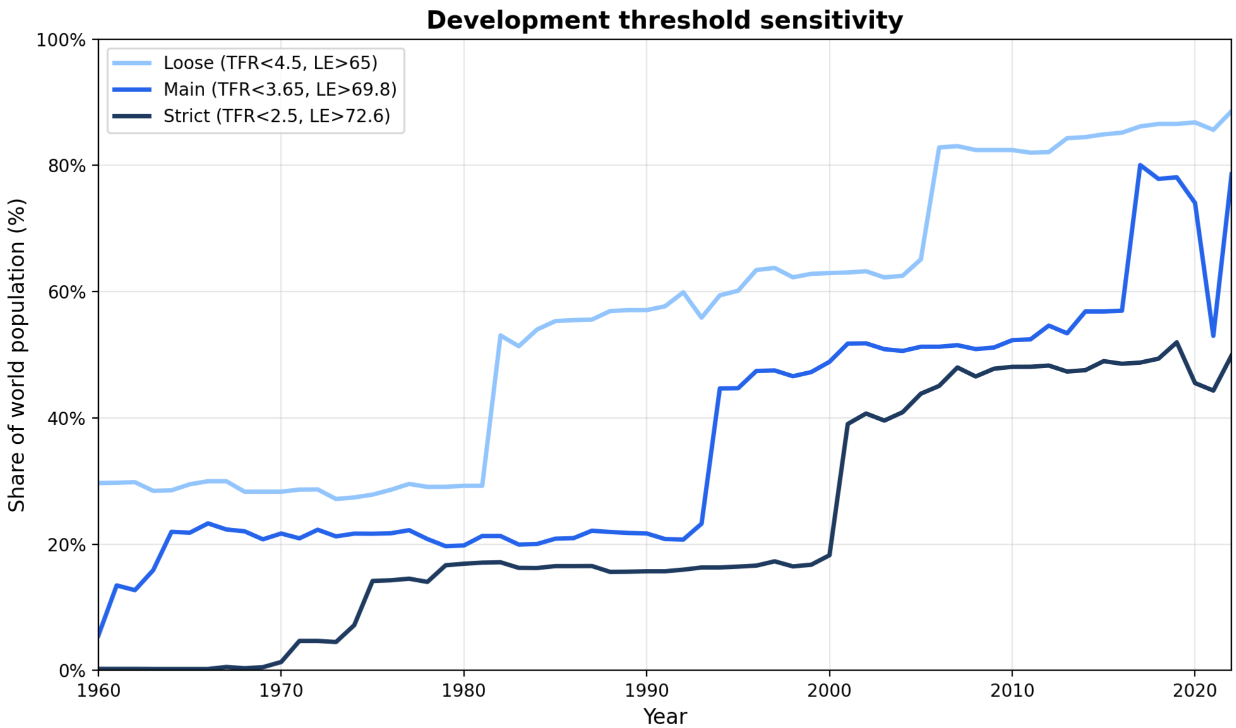

Figure 13: Cumulative share of world population crossing development thresholds under three specifications: loose (TFR < 4.5, LE > 65), main (TFR < 3.65, LE > 69.8), and strict (TFR < 2.5, LE > 72.6). The shape is identical across specifications — the same China jump, the same steady rise. The threshold choice shifts the level; it does not change the story. *(Source: World Bank WDI, WCDE v3.)*

A.3.6 Data-Source Robustness

I use WCDE v3, which reconstructs historical education distributions from census microdata, DHS surveys, and demographic back-projection. Three tests verify that the main forward-prediction and residualization results are not artifacts of WCDE's reconstruction methodology: replication on the independent Barro-Lee v3.0 dataset (Lee & Lee 2016) with different source surveys and interpolation; residualization on Barro-Lee; and restriction of WCDE to post-1970 observations, where data quality is highest and reconstruction plays the least role.

Specification	LE R^2	TFR R^2
WCDE v3 (main)	0.335	0.388
Barro-Lee v3.0 ("at least some secondary," age 15–24, 146 countries)	0.320	0.287
Barro-Lee v3.0 (mean years of schooling)	0.400	0.428
WCDE v3, post-1970 only	0.313	0.281
WCDE v3, 171-country panel (15 USSR republics excluded; 25-yr forward lag, country FE)	0.366	0.444
Barro-Lee residualization (GDP with education's contribution removed)	≤ 0.003	≤ 0.003

All specifications significant at $p < 0.0001$, same sign and direction. The finding that log GDP per capita has no independent predictive power is not specific to one dataset; the pre-1970 decades add statistical power but are not driving the result. The natural experiments (Cambodia, Bihar–Kerala, Asian financial crisis, Russia, South Africa) rest on WDI outcome data and documented historical events, not on WCDE reconstruction, and are unaffected by any data-source concern.[6]

Chapter 10 develops the credibility case and the four-signature phenotype test that motivates the 171-country row, including the subgroup split under Barro-Lee that identifies the bias as specifically Goskomstat rather than socialist in general. Every headline coefficient in Table 9–Table 17 moves in the same direction under the 171-country exclusion: upward. USSR inclusion was attenuating the paper's signal, not producing it.

A.3.7 Specification Robustness

Six additional checks address standard concerns about the Table 9 headline. They are summarised in Table A14 and described in turn below.

Table A14: Specification robustness.

[6] Barro-Lee v3.0 source: `github.com/barrolee/BarroLeeDataSet`.

Check	β (or elast.)	SE	N	Countries
5-year baseline (Table 9 col 1)	1.376	0.083	629	105
10-year aggregation, active expansion	1.437	0.078	416	113
Annual (interpolated), active expansion	1.326	0.078	3,548	135
Balanced subpanel (9/9 obs), active expansion	1.490	0.124	423	47
Within-year cross-cohort (20–24 on 45–49)	1.965	0.051	1,747	171
PPML, TFR_{T+25} (semi-elast. per pp)	-0.91%	0.14%	590	172
PPML, U5MR_{T+25} (semi-elast. per pp)	-3.30%	0.36%	579	168
Log LE_{T+25} (semi-elast. per pp)	$+0.36\%$	0.04%	590	172
Log TFR_{T+25} (semi-elast. per pp)	-0.82%	0.11%	590	172
Log U5MR_{T+25} (semi-elast. per pp)	-2.59%	0.28%	579	168

Notes: Each row reports the parental-education coefficient (or its semi-elasticity in the PPML/log rows) under the noted check, with country fixed effects and country-clustered standard errors. “Active expansion” restricts parent completion $< 30\%$ to match the Table 9 sample definition. All rows significant at $p < 0.01$.

Period length. Aggregating child cohorts at 5-year (baseline), 10-year, and annual (interpolated) resolution leaves the active-expansion coefficient in the 1.33–1.44 range. The annual grid mechanically multiplies the observation count roughly fivefold; the point estimate barely moves. Time-aggregation is not driving the headline.

Balanced panel. The WCDE cohort grid is balanced by construction at the full-panel level (every country has 9 cohort years 1975–2015), so the headline is not a composition artefact. Inside the active-expansion subset, forty-seven countries satisfy the $< 30\%$ cutoff in all nine cohort years; re-estimating on this deepest-observed subset gives $\beta = 1.49$, slightly above the unbalanced headline of 1.36. Balance does not weaken the effect.

Within-year cross-cohort. The headline regression uses a 25-year forward lag: child education at t on parental education at $t - 25$. A standard concern is that shocks occurring between $t - 25$ and t (global MDGs, the Asian financial crisis, HIV) are absorbed by the forward-lag coefficient. The WCDE v3 microdata (`prop_both.csv`) gives the lower-secondary-completion share for every 5-year age group in every country-year. In any given year, the age-20–24 share and the age-45–49 share are one generation apart but measured simultaneously: any shock in the intervening 25 years affects both cohorts, is absorbed by country FE if country-specific, and cannot drive the coefficient.

The within-year cross-cohort coefficient is $\beta = 1.97$ (SE 0.05) in the active-expansion sample, directionally stronger than the forward-lag headline. The transmission mechanism does not depend on the forward-lag design choice.

PPML. TFR and under-5 mortality are non-negative and right-skewed; the standard log-linear regression drops zeros and is sensitive to heteroscedasticity in $E[\log y \mid x]$. Poisson PML (Santos Silva and Tenreyro 2006) is consistent under arbitrary heteroscedasticity provided the conditional mean is correctly specified. Under PPML with country FE, a one-percentage-point rise in parental lower secondary completion lowers TFR at $T + 25$ by -0.91% (SE 0.14%, $p < 0.001$) and under-5 mortality by -3.30% (SE 0.36%, $p < 0.001$). GDP's coefficient is insignificant in both equations ($p = 0.24$ and $p = 0.69$), consistent with the residualization result.

Log outcomes. Putting LE, TFR, and U5MR in logs gives elasticities directly comparable to the PPML semi-elasticities. A one-percentage-point rise in parental completion predicts, one generation later: life expectancy $+0.36\%$, total fertility -0.82%, under-5 mortality -2.59%. All three are significant at $p < 0.001$ with country FE and log GDP included. The levels specification reported in Table 13 gives the same sign and ordering; the choice of functional form does not materially change the substance.

Wooldridge strict-exogeneity test. The regression-based test from Wooldridge (2010, §10.5) adds the one-period lead of parental education (at $t - 20$) to the FE regression and tests whether its coefficient is zero. On the full panel the lead enters with $\hat{\gamma} = 1.275$ (SE 0.114, $t = 11.1$); on the active-expansion sample $\hat{\gamma} = 0.655$ (SE 0.281, $t = 2.3$). The full-panel rejection is the same sample-composition signature as the 2WFE collapse (§A.1): pooling saturated and pre-expansion country-years inflates the test statistic because the regressor follows a near-linear trend within country once the state has committed, and the lead carries trend information that the lag has not yet absorbed. On the active-expansion subsample the t-statistic drops by a factor of five and rejection is marginal; the residual is consistent with the country-specific commitment slope, not with a time-varying confounder. My identification does not rest on strict exogeneity in any case. The Callaway–Sant'Anna event study uses not-yet-treated countries as the comparison set; the forward-lag design rules out reverse causality from individual outcomes back to parental schooling; and the natural experiments in Chapter 7 use exogenous shocks to the channel that no time-varying confounder rationalises (the educated population destroyed in Cambodia did not return when GDP returned). The Wooldridge statistic is reported as a courtesy to the panel-econometrics reader, not as a load-bearing test.

A.3.8 South Africa: Apartheid Stratification and the PET Curve

South Africa's 65% lower-secondary completion in 1990 is an aggregate masking apartheid stratification: near-universal white completion and much lower 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 2.41 (2005) is the composition result of Section 9.8 operating as predicted. Primary education drives fertility decline hardest ($R^2 = 0.65$ for TFR), and Black primary was expanding through the shock. 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 primary-to-fertility mechanism continuing to run while lower-secondary depth remained too shallow in the affected majority to deliver the other half of the composition: the behavioural response to a novel pathogen that deeper schooling 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 reform evidence is in §7.3 (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. (2017) 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.

A.3.9 Colonial Test: IV First-Stage Diagnostics

The full IV diagnostics behind the Colonial Test (Section 7.4) are tabulated in Section 13.4, Tables A7–A9: bivariate horse race, religion as excluded instrument, and the symmetric diagnostic that treats *avexpr* as the instrument for education. The structural finding is that on AJR's 64-country base sample, AJR's own *avexpr* institutional measure is itself a strong instrument for education (first-stage $F = 37.13$, three times the Stock & Yogo threshold), so the same colonial-era variation drives both channels and the exclusion restriction underlying institutionalist IV identification is empirically violated by AJR's own data.

A.3.10 Panel Summary Statistics and Country Coverage

Table A16: Summary statistics for analysis variables.

Variable	n	mean	sd	min	max
Parental lower-sec completion (%, $T-25$)	1,665	42.0	33.3	0.0	100.0
Child lower-sec completion (%, T)	1,665	61.9	31.6	0.4	100.0
Log GDP per capita (const. 2015 USD)	1,466	8.28	1.49	4.98	11.67
Life expectancy at birth (years)	1,608	65.7	10.6	12.8	84.3
Total fertility rate (births/woman)	1,608	3.70	1.95	0.91	8.86
Under-5 mortality (per 1,000 live births)	1,566	68.3	69.9	2.2	338.0

Notes: Analysis panel: 185 countries, 1975–2015 at 5-year intervals, with parental education lagged 25 years (observed at $T-25$). 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). Country coverage for every main specification is listed in `checkin/summary_stats.json` under `country_lists`.

A.3.11 Listwise-Deletion Stress Test

A natural concern is whether the headline result is driven by which countries get dropped when each variable is required. The audit below reports per-period coverage by variable and re-estimates the headline parental-education coefficient on three nested samples.

Table A18: Per-period country coverage by variable.

Year	Parent edu	Child edu	log GDP	LE	TFR	U-5	Edu only	All vars	Listwise drop
1975	185	185	135	192	192	186	185	127	31.4%
1980	185	185	149	192	192	186	185	139	24.9%
1985	185	185	156	192	192	186	185	145	21.6%
1990	185	185	181	193	193	186	185	168	9.2%
1995	185	185	184	193	193	186	185	170	8.1%
2000	185	185	187	193	193	186	185	172	7.0%
2005	185	185	189	193	193	186	185	172	7.0%
2010	185	185	190	193	193	186	185	173	6.5%
2015	185	185	191	193	193	186	185	173	6.5%

Table A19: Headline parental-education coefficient under three sample definitions.

Sample definition	β	n	Countries
Any-available T1 pair (parent & child edu only)	1.364	783	137
Listwise-complete on T1 inputs (+ log GDP)	1.353	719	123
Listwise-complete on T1 + T7 inputs	1.354	715	121

Notes: Active-expansion subsample (parent completion below 30%); country fixed effects; child lower secondary completion at T on parental education at $T-25$. "Edu only" counts countries with both parental and child education; "All vars" counts those with all six variables. "Listwise drop" is the share lost when adding the GDP/LE/TFR/U-5 requirement, mostly driven by GDP coverage in 1975–85. The headline coefficient barely moves across the three sample definitions ($1.364 \rightarrow 1.353 \rightarrow 1.354$); listwise deletion is not driving the result.

References

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

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

Angrist, J.D. & Pischke, J.S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

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.

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

Baker, D.P., Collins, J.M. & Leon, J. (2017). 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.

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

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

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

Caldwell, J.C. (1976). Toward a Restatement of Demographic Transition Theory. *Population and Development Review*, 2(3/4), 321–366.

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

Callaway, B. & Sant'Anna, P. (2021). Difference-in-Differences with Multiple

Time Periods. *Journal of Econometrics*, 225(2), 200–230.

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.

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.

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.

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

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

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *Ameri-*

can Economic Review, 91(4), 795–813.

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

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.

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.

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

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

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Econometrica*, 89(5), 2261–2290.

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

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.

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.

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.

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.

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.

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.

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 Increasing Global Progress and Persistent Disparities. *Nature Human Behaviour*, 5, 701–711.

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

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

Mamun, M.A. & Bongaarts, J. (2022). Revisiting the causes of fertility decline in Bangladesh: the relative importance of female education and family planning programs. *Asian Population Studies*, 18(2), 125–148.

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

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

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

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

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.

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.]

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

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.

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.

Santos Silva, J.M.C. & Tenreyro, S. (2006). The Log of Gravity. *Review of Economics and Statistics*, 88(4), 641–658.

Sapolsky, R.M. (2017). *Behave: The Biology of Humans at Our Best and Worst*. Penguin Press.

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

Seba, T. (2014). *Clean Disruption of Energy and Transportation: How Silicon Valley Will Make Oil, Natural Gas, Coal, Nuclear and Utilities Obsolete by 2030*. Clean Planet Ventures.

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

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

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

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

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

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

UNESCO. (2023). *Global Education Monitoring Report 2023: Technology in Education — A Tool on Whose Terms?* Paris: UNESCO.

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.

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

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.

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.

World Bank. (2018). *World Development Report 2018: Learning to Realize Education's Promise*. Washington, DC: World Bank.

World Bank. World Development Indicators. Various years.