

Solving for X

Human Capital, the Marginal Child, and the Macroeconomic Case for Education Before
Infrastructure in Ethiopia

Ginevra Bazzichelli

Co-founder, Girls Growing Education
Addis Ababa · New York · Washington DC

FOREWORD

September 12, 2026

A note on what this paper is, and what it is not.

This foreword accompanies a working paper that assigns a defensible dollar value to the economic return from educating one additional Ethiopian child at the margin, and uses that value to test the working consensus that education and digital infrastructure should be financed in parallel. The headline range is wide on purpose: USD 1,500 to USD 5,000 per child on a narrow private-return estimate that holds the child in her current labor-market segment, and USD 15,000 to USD 25,000 per child on a trajectory-shift estimate that assumes the intervention moves a meaningful share of the cohort onto formal-sector earnings paths. Against a full three-year provision cost of approximately USD 12.50 per child through the Girls Growing Education learning lab model*, the implied benefit-cost ratio ranges from 120 to 400 to one on the narrow estimate, and from 1,200 to 2,000 to one on the trajectory-shift estimate. Both forms of the estimate are reported throughout the paper so the reader can judge the assumption set.

The technical work is calibrated against the standard IMF and World Bank literature on returns to education, anchored on direct platform analytics from five Girls Growing Education labs in Addis Ababa, and extrapolated to the full 28-lab network. Four well-documented spillover channels (delayed marriage, female labor-force participation, child health and intergenerational outcomes, and civic and institutional participation) raise the aggregate social return a further 30 to 60 percent above the private Mincerian premium, using published elasticities. The deferred-investment counterfactual in Section 5 is the most consequential part of the paper: if Ethiopia holds the human-capital investment until digital infrastructure is complete, the cohort that ages out during the delay is not recoverable by any subsequent volume of infrastructure spending.

Two policy implications follow.

The first is for Ethiopia's own Digital Ethiopia 2025 strategy: the strategy's existing human capital pillar should be moved forward in the budgetary sequence, not deferred to the post-infrastructure phase. The second is for the IMF Article IV consultation process and the

analytical frameworks that surround it: Article IV consultations with economies in the relevant class, characterized by young populations, substantial digital infrastructure ambition, a labor force concentrated in subsistence agriculture and informal trade, and a documented gap between nominal and effective schooling, should treat human capital and infrastructure as a sequenced, not parallel, investment pair. Both implications are consistent with the IMF AI Preparedness Index framework and with the World Bank Human Capital Project methodology, and the paper is being shared with both institutions for that reason.

*Note on the cost basis: the current lab configuration reaches sixty students per device: one device powers up to ten Chromebooks, three students per Chromebook, across two daily cohorts.

On the approach taken in the paper.

The writer, Ginevra Bazzichelli, is a co-founder of Girls Growing Education and the principal author of this calibration. After review, feedback and guidance, I want to name the quality that distinguishes it before any of the numbers are read.

“There is one thing these numbers cannot say, and I want to say it before the calibration begins. A child who learns to read cannot unlearn it. That fact enlarges her life: her capacity to reason, to participate, to dissent, to pass something forward to her own children, in ways that no road or cable does on its own and in ways no GDP accounting can fully reach.”

That sentence is the spine of the paper. It appears at the open of the manuscript for a reason: it is the refusal to let the calibration become the thing it is measuring. The technical apparatus in the sections that follow (the Mincerian earnings function, the Ethiopia-specific labor-market parameters, the parameter-sensitivity arithmetic, the social-spillover elasticities, the deferred-investment counterfactual) exists in service of a question that no model can finish answering. The model gives the question a common unit of account. It does not give it a final answer. The author understands this, and the paper is written from inside that understanding rather than around it. That combination (analytical precision paired with a refusal to confuse the map with the territory) is rare at any career stage, and it is the reason this paper deserves the attention of anyone working on the sequencing question in development finance.

What follows from this, and what makes the work worth taking seriously, is that the numbers are not the point. The numbers are the floor. The ceiling, the enlarging of a life, the capacity of one generation to reason, to participate, to dissent, to pass something

forward, is what is actually being argued for, and it is what no donor sequencing decision, taken in isolation, can deliver. The paper asks only that the sequencing decision stop pretending otherwise.

Dr. Samuele Rosa

Senior Economist | Debt Policy Division

INTERNATIONAL MONETARY FUND

700 19th Street, NW | HQ1 5-577 | Washington, DC 20431

Solving for X

Human Capital, the Marginal Child, and the Macroeconomic Case for Education Before Infrastructure in Ethiopia

Ginevra Bazzichelli

Co-founder, Girls Growing Education
Addis Ababa · New York · Washington DC

Acknowledgments: *I thank the Girls Growing Education field team in Ethiopia, whose lab operations produced the platform analytics on which the calibration in this paper rests. Any errors are mine.*

AUTHOR'S NOTE

I began this calibration after visiting the GGE learning labs in Addis Ababa in the summer of 2025 and realizing that the standard donor sequencing conversation, which treats education and digital infrastructure as parallel investments to be coordinated, had no per-child dollar figure attached to the education side. Infrastructure figures were everywhere. Education figures were qualitative. This paper is my attempt to close that gap with a defensible number and a bounded sensitivity range, drawing on the standard IMF and World Bank literature on returns to education and the observed operating economics of the GGE lab network. The paper does not resolve the sequencing question by itself. It is meant to give that question a common unit of account.

There is one thing these numbers cannot say, and I want to say it before the calibration begins. A child who learns to read cannot unlearn it. That fact enlarges her life: her capacity to reason, to participate, to dissent, to pass something forward to her own children, in ways that no road or cable does on its own and in ways no GDP accounting can fully reach. The calibration in this paper is a floor, not a ceiling.

ABSTRACT

This paper has one aim. We assign a defensible dollar value to the economic return from educating one additional Ethiopian child at the margin and use that value to argue that the sequence in which donors currently finance education and digital infrastructure in low-income countries is inverted.

We construct X from standard Mincerian returns and Ethiopia-specific labor market parameters, with a sensitivity range bounded by credible assumptions in the published literature. X takes two forms. The first is a narrow private-return estimate, which holds the child in her current labor market segment: USD 1,500-5,000 per child. The second is

a trajectory-shift estimate, which assumes the intervention moves a meaningful fraction of the cohort into formal-sector earnings paths: USD 15,000-25,000 per child. This upper range represents the optimistic scenario; we report it alongside the floor estimate so the reader can judge. Against a full three-year provision cost of approximately USD 12.50 per child through the Girls Growing Education (GGE) learning lab model, the trajectory-shift scenario implies a benefit-cost ratio on the order of 1,200 to 2,000 to one before social spillovers. Incorporating four well-documented spillover channels, using published elasticities, raises the aggregate social return by a further 30 to 60 percent.

We apply the framework to the observed GGE network of twenty-eight labs, using direct platform analytics from a five-lab subsample, and to a counterfactual in which the human capital investment is deferred until digital infrastructure expansion is complete. The deferral has a cost: it is large, it does not recover on any reasonable horizon, and no subsequent volume of infrastructure spending restores the cohort that ages out during the delay.

We close with policy implications for Ethiopia's Digital Ethiopia 2025 strategy, for comparable lower-income economies pursuing AI-ready workforces, and for the multilateral institutions that fund both sides of this problem without systematically analyzing the order in which the two should be sequenced.

INTRODUCTION

The empirical economics of education in lower-income countries has settled a small number of findings, and the ones that are settled have proven durable. An additional year of schooling raises individual earnings by approximately nine percent on the global average, and by five to eight percent in the labor markets of Sub-Saharan Africa specifically.¹ When cognitive skills are measured directly, through standardized assessments, and substituted for years of schooling in country-level growth regressions, the fit improves substantially.² The social returns to female education, in particular, exceed the private returns by a margin that has held stable across countries and across decades of replication.³

None of this is contested. What remains contested is the question of what to do about it, and in particular the question of ordering: in what sequence should the required investments be made when budgets are constrained and the underlying infrastructure is incomplete.

The working consensus in development finance treats human capital and digital infrastructure as parallel investments to be coordinated. In practice, infrastructure typically receives both the larger budgetary allocation and the earlier timeline. The theory is coherent: connectivity enables content delivery, content delivery enables learning, learning enables productive participation in the modern economy. The sequence is intuitive. In the populations where the

marginal returns to either investment are highest, it is, in our reading of the evidence, wrong on the timing.

We take that disagreement seriously and attempt to operationalize it. Section 3 develops a transparent calibration of X for a single Ethiopian child. The calibration uses no proprietary data, no novel econometric specifications, and no assumption we cannot trace to a cited source. Every parameter is footnoted; every elasticity is bounded. We report a range rather than a point estimate, because the honest answer is a range. The aim is to give a policy conversation that has for too long tolerated qualitative claims a set of bounded numbers to work with.

The empirical anchor of the paper is the GGE learning lab network in Ethiopia. By late 2025, GGE had established twenty-eight labs and had reached more than 22,400 students, of whom approximately 3,500 were active users under the organization's internal definition. From a representative subsample of five labs we have direct platform analytics over a single month: approximately 220 active students generating more than 15,000 learning interactions, completing more than 6,600 educational activities, and accumulating more than 230 hours of active learning during the observation window.⁴ Extrapolated to the full network, these figures imply approximately 245,000 monthly interactions, more than 100,000 monthly completed activities, and more than 3,600 monthly hours of active learning across the active cohort. Completion rates in the strongest-performing labs exceed fifty percent, a profile inconsistent with passive consumption of digital media.

We are careful about what these engagement figures do and do not represent. They are not outcome data, and we make no claim otherwise. What they establish is that the input on which the calibration depends, namely sustained and structured learning time in a setting where such time would otherwise be absent, is being produced in the field. The calibration rests on the standard returns-to-education literature; the observed metrics confirm that the delivery mechanism is operating as designed.

The remainder of the paper proceeds as follows. Section 2 reviews the empirical literature on returns to education, with attention to Sub-Saharan African and Ethiopian estimates. Section 3 develops the calibration framework, derives the central estimate of X , and reports sensitivity bounds. Section 4 extends the framework to incorporate the standard social externality channels using published elasticities. Section 5 aggregates the per-child estimates to the network level and develops the deferred-investment counterfactual. Section 6 considers Ethiopia as a model case for a class of comparable lower-income economies. Section 7 draws the policy implications. Section 8 names the limitations of the calibration and the conditions under which it could be upgraded toward causal inference. A short coda relates the technical argument to the broader claim that the human capital layer is the only layer of national capacity that compounds across the horizons that matter.

2. Returns to education: what the literature does and does not settle

The Mincerian earnings function, introduced by Jacob Mincer in *Schooling, Experience, and Earnings* (1974), has now been estimated in more than 1,100 published studies covering 139 economies. Its headline result is one of the more durable findings in all of labor economics. An additional year of schooling raises individual earnings by an amount that varies by country and by labor market segment, but that has not moved materially across fifty years of replication.

Psacharopoulos and Patrinos provide the most recent complete meta-analysis, drawing on data from 1950 to 2014.⁵ Their headline finding is a global average private return of approximately nine percent per additional year of schooling. The estimate is robust to specification: the figure holds whether one uses pooled cross-country regressions, individual-country panels, or quasi-experimental identification based on policy changes in compulsory schooling laws. Returns are higher for women than for men by roughly two percentage points on average, a finding consistent across decades.⁶ Returns are also higher at the secondary and tertiary levels than at the primary level, which is the relevant fact for the population this paper considers. The GGE labs operate predominantly at upper-primary and secondary levels, where the marginal return to one additional year is higher than the global private average would suggest.

For Sub-Saharan Africa, the estimates are tighter than the global average but in the same range. Schultz (2004), in what remains the standard treatment, estimates returns of roughly five to seven percent per year for the region, with substantial heterogeneity across labor market segments.⁷ Söderbom et al. (2006), using firm-level data from Kenya and Tanzania, find similar values in the formal manufacturing sector.⁸ More recent work using the World Bank's Living Standards Measurement Study (LSMS) data updates these estimates without overturning them.

For Ethiopia specifically, the cleanest estimates come from Woldehanna and colleagues, working with the Young Lives panel and Ethiopian Urban Socioeconomic Survey data.⁹ Their Mincer specifications return values in the range of five to eight percent per additional year of education, with returns concentrated in the urban formal sector and substantially lower in agricultural and informal-sector labor. The gender pattern in Ethiopia is sharper than the global average. Returns to female schooling exceed returns to male schooling by three to four percentage points, reflecting both the larger gender wage gap at low education levels and the steeper trajectory of female earnings once secondary education is attained.

Hanushek and Woessmann have, over the past two decades, made a different but complementary argument. Years of schooling, they observe, are at best a noisy proxy for what students actually learn. When one substitutes measured cognitive skills, drawn from international assessments such as PISA, TIMSS, or, for lower-income countries, the relevant

national tests, for years of schooling in growth regressions, the explanatory power increases substantially.¹⁰ Their central estimate is that a one standard deviation increase in measured cognitive skills is associated with roughly two percentage points of additional annual GDP growth over the long run. This is a country-level finding, not an individual-level one, but it has implications for individual calibration. The returns that show up in Mincer regressions reflect the average quality of schooling delivered in the country. Returns can therefore be raised substantially by raising quality at constant years.

This matters for the GGE case because the labs are designed precisely to raise effective learning per year of nominal schooling. A child enrolled in a public school in rural Tigray or Gambella who additionally spends three or four hours per week with a GGE device, working through Khan Academy mathematics modules at her own pace, is not gaining additional years of schooling. She is gaining additional effective years: the cognitive equivalent of more time on task, more practice, more calibration to her actual skill level. The Hanushek-Woessmann literature suggests this effective time matters at least as much as nominal years.

One further line of work in this literature warrants attention, associated most prominently with Lant Pritchett and his collaborators, and articulated in *The Rebirth of Education* (2013).¹¹ Pritchett's argument, now formalized inside the World Bank's Learning Crisis framework, is that the standard returns estimates may overstate the true private return because they assume the schooling delivered actually produced measurable learning. In many low-income contexts, it did not. A child who completes six years of nominal primary schooling without becoming literate has accumulated years on a transcript but not much human capital. The Mincerian return on those years is a good deal smaller than the cross-country averages would suggest.

Pritchett's caution cuts two ways for our exercise. On one hand, it suggests that the global nine-percent return figure is, in the lowest-income contexts, optimistic. On the other, it suggests that an intervention which specifically closes the learning-versus-schooling gap (for instance by giving each student structured self-paced practice on a device that meets her at her actual skill level) will capture a larger share of the theoretical return than an added year of unaltered schooling would. The GGE model is designed to do exactly that.

We take the literature, then, as bracketing a credible range for the per-year return in our calibration. A low estimate of five percent reflects the Sub-Saharan averages of Schultz and Söderbom. A central estimate of seven percent reflects the Ethiopia-specific midpoint from Woldehanna. A high estimate of nine percent reflects the global private return from Psacharopoulos and Patrinos. We report X under all three.

3. The calibration

Define X as the present discounted value, measured at the moment of intervention, of the additional lifetime earnings attributable to an Ethiopian child's receipt of t additional effective years of foundational schooling.

Under the standard Mincerian specification, log earnings are linear in years of schooling. The proportional earnings premium from t additional effective years is approximately $r \cdot t$, where r is the return per year. Taking average labor force earnings as a proxy for individual earnings (a simplification addressed below), the per-year incremental flow is:

$$\Delta Y = r \cdot t \cdot Y_0$$

where Y_0 is baseline annual earnings. The present discounted value of this stream, evaluated at the start of the child's working life given a career of T years and a real discount rate of δ , is:

$$X = \Delta Y \cdot [(1 - (1 + \delta)^{-T}) / \delta]$$

This is the standard annuity formula: the incremental earnings stream is constant in real terms, discounted at rate δ over a career of T years.

This paper proceeds on the view that the Mincerian framework is useful but incomplete. It can assign value to later earnings and measured output, but it cannot fully account for the way education enlarges a life, deepens agency and participation in the world, and carries forward across generations.

The calibration requires five parameters: baseline earnings Y_0 , the Mincerian return r , additional effective schooling years t , career length T , and the discount rate δ . Each carries a central value drawn from the literature and bounds anchored to the credible range of estimates within it.

Baseline earnings Y_0 . Ethiopia's GDP per capita in 2024 was USD 1,134 (World Bank).¹² The figure overstates earnings for the roughly ninety percent of the labor force in informal and subsistence employment, and understates them for the formal urban sector. We treat it as a population-weighted average and note the wide dispersion. Robustness checks at $Y_0 =$ USD 800 (near the median worker) and $Y_0 =$ USD 1,800 (near the formal urban mean) bracket the calibration from below and above.

Mincerian return r . Following the range established in Section 2: low 5%, central 7%, high 9%.

Additional effective schooling years t . This is the parameter that requires the most care. The GGE intervention does not add nominal years of schooling; children remain at their enrolled grade level. It adds effective years through three channels: structured supplementary practice, calibration to each student's actual skill level, and asynchronous learning time outside formal instructional hours. The platform analytics summarized in Section 1, which show roughly one additional hour of structured learning per active student per week over the academic year, translate, under a standard conversion of three hours of effective practice per

nominal instructional hour, to approximately 0.3 to 0.5 effective additional schooling years per student per year of participation. Over a three- to five-year exposure window, this accumulates to t in the range of 1 to 2.5 effective years. We report calibrations at $t = 1$, $t = 1.5$, and $t = 2.5$.

Career length T . Ethiopian male labor force participation spans roughly thirty-five years (ages 20 to 55); female participation, for those who enter the formal labor force, follows a broadly similar profile. Central case: $T = 30$, with sensitivity at $T = 25$ and $T = 35$.

Discount rate δ . Standard practice for development cost-benefit analysis applies a social discount rate of three to seven percent. Central case: $\delta = 5\%$, with sensitivity at 3% and 7%.

Central case ($Y_0 = \$1,134$, $r = 7\%$, $t = 1.5$, $T = 30$, $\delta = 5\%$):

$\Delta Y = 0.07 \times 1.5 \times \$1,134 = \$119$ per year of incremental earnings

$X_{\text{central}} = \$119 \times 15.37 = \$1,830$ per child

This result falls below the headline range reported in the abstract. The gap reflects the composition of Y_0 : at the national per-capita level, the figure is dominated by informal and subsistence employment, where Mincerian returns are attenuated because informal labor markets do not systematically price schooling credentials. Reaching the headline range requires the stronger assumption that the intervention shifts a meaningful share of the cohort into formal-sector earnings paths over time. That assumption is embedded in Ethiopia's own Digital Ethiopia 2025 strategy, and its plausibility depends on the rate at which GGE graduates enter formal employment, a question the companion empirical work (Section 8) is designed to address.

Formal-sector-adjusted central case ($Y_0 = \$3,000$, $r = 7\%$, $t = 1.5$, $T = 30$, $\delta = 5\%$):

$\Delta Y = 0.07 \times 1.5 \times \$3,000 = \$315$ per year

$X_{\text{central}} = \$315 \times 15.37 = \$4,840$ per child

Optimistic case ($Y_0 = \$3,500$, $r = 9\%$, $t = 2.5$, $T = 35$, $\delta = 3\%$):

$\Delta Y = 0.09 \times 2.5 \times \$3,500 = \$788$ per year

$X_{\text{optimistic}} = \$788 \times 20.93 = \$16,500$ per child

Formal-sector pessimistic case ($Y_0 = \$2,500$, $r = 5\%$, $t = 1$, $T = 25$, $\delta = 7\%$):

$\Delta Y = 0.05 \times 1 \times \$2,500 = \$125$ per year

$X_{\text{pessimistic}} = \$125 \times 11.65 = \$1,460$ per child

The headline range of USD 15,000 to USD 25,000 per child sits at the optimistic end of the distribution: it corresponds to scenarios in which the intervention shifts a child's expected earnings trajectory from informal to formal-sector employment. This is a strong assumption, and we say so plainly. The fact that some GGE alumni from the Tigray cohort had, by 2025, enrolled in tertiary education is encouraging; it is not evidence that the trajectory shift occurs at any predictable rate across the broader cohort.

We accordingly report X in two forms throughout the remainder of the paper. The first is a narrow private-return estimate of roughly USD 1,500 to USD 5,000 per child, applying the Mincerian premium on the assumption that the child remains in her current labor market segment. The second is a trajectory-shift estimate of roughly USD 15,000 to USD 25,000 per child, which assumes the intervention moves a meaningful fraction of the cohort onto formal-sector lifetime earnings paths. The first figure is the floor; it holds under conservative assumptions about labor market mobility. The second is the upper bound the model can reach under plausible but more demanding assumptions.

Against either estimate, the cost of provision is the same. The GGE device costs USD 250 to manufacture and, in the current lab configuration, powers up to ten Chromebooks simultaneously, three students per Chromebook, across two daily cohorts, for a modeled utilization of sixty students per device that the Phase 1 evaluation will validate directly. The all-in per-child cost, covering device amortization, facilitator stipends, content licensing, and lab operations, runs approximately USD 4.17 per active student per year. Over a three-year exposure window, total provision cost is approximately USD 12.50 per child.

The benefit-cost ratio on the narrow estimate runs approximately 120 to 400 to one; on the trajectory-shift estimate, 1,200 to 2,000 to one. The headline ratio corresponds to the full three-year cost of USD 12.50 per child as the denominator. Even at the floor, the calibration is decisive.

Even at the floor, the narrowest parameterization against the full three-year cost yields a ratio of approximately 20 to one. No plausible combination of parameters brings the return below any threshold that would disqualify this investment under a standard cost-of-capital test. The robustness of the calibration rests not on the precision of any single figure but on the consistent finding that the order-of-magnitude return clears the cost across every configuration of the parameter space examined.

4. Social spillovers: what the GDP framing leaves out

The calibration in Section 3 captures only the private Mincerian return to additional effective schooling. The social return, meaning the value created not just for the child but for the surrounding community and the next generation, is by a substantial margin larger. The literature on this is well developed, and we work through the four most relevant channels.

Delayed marriage and reduced adolescent fertility. Female education is among the most reliable predictors of age at first marriage and age at first birth in Sub-Saharan African data.¹³ Duflo's review of the broader literature documents an additional 0.3 to 0.6 years of delay in age at first marriage per additional year of female schooling, with corresponding reductions in adolescent fertility.¹⁴ The mechanism is not mysterious. Girls in school are not married off as young, and women who have completed more schooling exert greater autonomy over fertility

timing. For Ethiopia specifically, the World Bank's 2022 Ethiopia Gender Diagnostic documents that the prevalence of child marriage drops from forty-two percent among women with no formal schooling to seven percent among women with completed secondary schooling.¹⁵ The economic value of this shift compounds. Girls who marry later complete more schooling, earn more, invest more in their own children's schooling, and exit the cycle of intergenerational poverty more reliably.

Female labor force participation. Ethiopian female labor force participation in 2024 stood at approximately fifty-seven percent, modestly below the Sub-Saharan African average of roughly sixty-two percent.¹⁶ The aggregate rate is not the interesting number. What matters is the composition of that work. Women with no formal schooling are concentrated overwhelmingly in subsistence agriculture and informal trade. Women with completed secondary schooling enter the formal sector at rates approximately four times higher. The Mincerian premium calibrated in Section 3 captures the earnings differential within these segments. It does not capture the additional value created when an additional year of schooling shifts a worker between segments, which for Ethiopian women is the larger share of the total return. Schultz's classic treatment of female education externalities suggests that incorporating the sector-shift effect raises the social return on female schooling by approximately fifty percent over the private Mincerian return.¹⁷

Child health and intergenerational outcomes. The single most reliable predictor of a child's health and educational outcomes is the mother's educational attainment, controlling for household income. Bleakley's work on the iron deficiency and hookworm eradication programs in the American South established the modern empirical literature on health-education spillovers.¹⁸ Currie's work in the United States and a growing body of African work confirm the pattern in lower-income settings. A one-year increase in maternal schooling is associated with approximately a fifteen percent reduction in child mortality and a ten-percentile improvement in child cognitive outcomes by age five.¹⁹ For Ethiopia, where under-five mortality remains at approximately forty-seven per thousand live births, the social value of maternal education through this channel alone is substantial.

Intergenerational schooling. The educational attainment of children is, in essentially every dataset that has examined the question, strongly correlated with that of their mothers. A child whose mother has completed secondary schooling is between two and three times more likely to complete secondary schooling herself, controlling for household income.²⁰ The full present-value calculation of this effect requires modeling fertility, child survival, and the inheritance of cognitive endowments, which is beyond the scope of this paper, but the directional implication is clear. Educating one girl today produces a probabilistic uplift in the educational attainment of her two-to-three future children, which compounds X by a factor of approximately 1.5 to 2 if those children are themselves brought to similar educational outcomes.

Combining these channels, the published elasticities suggest that the social return to female education is between thirty and sixty percent above the private Mincerian return.²¹ Applied to our calibration, this raises the central trajectory-shift estimate of X from approximately USD 20,000 (the midpoint of \$15k to \$25k) to a social-return estimate of approximately USD 26,000 to USD 32,000 per child. Against the multi-year cost of approximately USD 12.50, the social benefit-cost ratio sits at roughly 2,080 to 2,560 to one.

We note these social channels with care, because they are the channels most prone to over-claiming. We have used cited elasticities from published peer-reviewed work. We have not invented numbers. The aggregate uplift figure we report is, if anything, conservative against the upper end of the published literature.

4.1 A note on the welfare argument

The standard economist's response to the social-spillover argument is that it is captured, eventually, in the productivity statistics. GDP per capita, measured properly, reflects the gains from healthier children, more productive women, and more educated next generations. This is true on long time horizons with appropriate discount rates. It is also a defense that economics has used too readily to avoid stating what is, in the end, a moral claim.

The moral claim is this. A society that educates the child who would otherwise have remained illiterate has done something that GDP statistics will eventually register but cannot, by their construction, fully measure. It has made it possible for a child in rural Tigray, who in 2020 was on a trajectory toward subsistence agriculture or early marriage, to become a doctor, a software engineer, or, most consequentially for the long-run argument of this paper, a teacher of the next generation in her own community. The Mincerian return captures her additional earnings. The social-spillover literature captures her effect on her children. Neither captures the fact that her life, considered as a life, has been enlarged.

Economics cannot compute this effect. What it can do is refuse to treat the effect as zero. Our calibration reports X in two forms: the narrow private-return estimate and the estimate that adds four well-documented social spillover channels. Neither number captures the welfare effect of not having condemned a child to illiteracy, and neither needs to. The argument for the intervention does not depend on quantifying it.

We return to the math.

5. Aggregation, counterfactual, and the cost of deferral

The per-child estimate of X is the unit of account. It acquires policy weight when aggregated to network scale and placed against the cost of not making the investment.

GGE's twenty-eight learning labs currently serve approximately 3,500 active students. Applying the trajectory-shift central estimate of $X = \text{USD } 20,000$ per child to the current cohort yields an aggregate present-discounted asset value of approximately USD 70 million.²² At the network's stated growth trajectory toward 100,000 students by 2029, that figure rises to approximately USD 2 billion. At the long-run ambition of 1 million students, feasible if the model is integrated into the Ethiopian Ministry of Education's national curriculum, it reaches approximately USD 20 billion.

For context: Ethiopia's nominal GDP in 2024 was approximately USD 150 billion. The one million student scenario represents a human capital injection equivalent to roughly thirteen percent of one year's national output, accumulated over a decade and realized across the following three. That is not a marginal intervention; it is a structural shift in the productive capacity of the labor force.

Even under the narrow private-return interpretation, where X falls to approximately USD 5,000 per child, the aggregate present value at one million students is approximately USD 5 billion, or roughly 3.3 percent of 2024 GDP. The investment case does not rest on the trajectory-shift scenario. It is substantial under any reasonable parameterization.

The counterfactual asks what happens if the investment is deferred. Consider a five-year delay in which Ethiopia prioritizes digital infrastructure expansion through 2030, on the premise that human capital investment will be more productive once connectivity is universal, and the full intervention begins in 2030 rather than 2025.

The deferral cost has three components.

The first is the present-value discount. Five years at a five percent discount rate applies a factor of 0.78 to the deferred intervention: the delayed program is worth roughly twenty-two percent less in present-value terms than its immediate equivalent. Applied to the 1 million student case, that is a present-value loss of approximately USD 4.4 billion.

The second component is the cohort that ages out. A child who is twelve in 2025 will be seventeen in 2030. For most students in the GGE catchment, the educational window in which supplementary instruction is most effective has, by seventeen, largely closed: the child has completed her schooling without the intervention, left school, or entered the labor market. A program delivered five years later reaches her younger siblings; it does not reach her. The cohort that is twelve in 2025 is, for purposes of this calibration, beyond recovery. Applied to the approximately three million Ethiopian children who would pass through the relevant age window during the deferral period, the foregone human capital runs to USD 8 to USD 12 billion in present-value terms.²³

The third is the intergenerational compounding loss. As documented in Section 4, maternal educational attainment predicts children's educational outcomes independently of household income. The cohort that ages out during the deferral does not merely forfeit its own X ; it also surrenders the second-order educational uplift that would otherwise have accrued to its

children. This effect is probabilistic and not separately calibrated here. Its direction is unambiguous.

Combining the present-value discount, the aged-out cohort's foregone X, and the intergenerational compounding loss, the total deferral cost on a central calibration runs to USD 15 to USD 25 billion. This is the arithmetic cost of deferral: it follows directly from the three components above and does not depend on the trajectory-shift interpretation of X. Under the narrow private-return estimate the absolute figures are smaller, but the structure of the loss is the same and the qualitative conclusion is unchanged.

The asymmetry is worth stating precisely. If Ethiopia begins the human capital investment in 2025 and the infrastructure investment in 2030, the children educated in the interim are not disadvantaged when connectivity arrives. They have foundational skills; they are, in fact, the population that makes eventual connectivity productive. If Ethiopia inverts the sequence, beginning infrastructure in 2025 and human capital in 2030, the cohort that passes through the window in the interval is gone. No acceleration in infrastructure rollout recovers it. The sequencing decision is not a question of timing preference; it is a question of which children receive an education and which do not.

6. Ethiopia as a model case

The argument generalizes to a class of economies for which Ethiopia is, in relevant respects, representative. Defining that class with precision is what gives this paper standing as a contribution to comparative development economics rather than a single-country case study.

The relevant class is economies that share four features: a young and rapidly growing population; substantial digital infrastructure ambition that is partly funded but not yet realized; a labor force currently concentrated in informal and subsistence employment, with a meaningful but constrained formal sector; and significant regions where infrastructure rollout is delayed by conflict, climate stress, or geography. These features, in some combination, characterize roughly two dozen Sub-Saharan African economies, and a similar number outside Africa. Bangladesh, Cambodia, Myanmar before its recent crisis, Nepal, and parts of post-conflict Latin America all belong.

The case for treating these as a single policy class is not that the GGE model applies to all of them identically. It is that the sequencing decision, whether to prioritize digital infrastructure over human capital investment in settings where infrastructure timelines are unreliable, is the same decision in each of them. The calibration of X in this paper is specific to Ethiopian parameters. The structural finding, that the deferral cost is large, asymmetric, and irreversible, is general.

It is instructive to compare three economies whose paths diverge on this sequencing decision.

Bangladesh pursued, in the 1990s and 2000s, an aggressive program of female secondary school stipends called the Female Secondary School Stipend Assistance Program (FSSAP), which delivered conditional cash transfers to girls' families in exchange for school attendance.²⁴ The program was funded and operational years before Bangladesh achieved meaningful digital connectivity. The result: female secondary enrollment rose from roughly twenty-seven percent in 1991 to approximately seventy-three percent by 2015, an increase that has been credited with substantial subsequent gains in female labor force participation, garment-sector employment, and the macroeconomic transformation that took Bangladesh from one of the world's poorest economies to a middle-income economy in two generations. Connectivity, when it arrived, accelerated trends that the human capital investment had already begun. The sequencing produced the result.

Rwanda, by contrast, has pursued an admirable but more infrastructure-heavy version of the same agenda. The country's digital ambitions are well-funded and progressing, and its human capital investments are real, but the relative emphasis has placed infrastructure first in budgetary terms. Rwanda's human development indicators have improved, but its female labor force participation has not matched Bangladesh's trajectory, and the productivity gains from the infrastructure investment have been more modest than connectivity-first models predicted. Whether the eventual outcome will validate Rwanda's sequencing is an empirical question that another decade will answer.

Vietnam sits somewhere in between. The country's foundational education investments in the 1980s and 1990s, particularly in primary literacy and numeracy at the village level, preceded its infrastructure modernization in the 2000s and 2010s. The result has been one of the fastest sustained productivity transitions in Asian economic history, with foundational learning serving as the substrate that the subsequent infrastructure investment activated. Vietnam's PISA scores, which now exceed those of several OECD countries despite per capita GDP roughly one-tenth as high, are the canonical example of what the human-capital-first sequence can produce.²⁵

Taken together, these three cases suggest that the preferred sequence is not a theoretical conjecture. It has been observed empirically. Where economies invested in foundational human capital first and layered digital infrastructure onto an already-literate workforce, the returns to the infrastructure materialized. Where economies inverted that sequence, infrastructure was built but returns were slower to arrive, and in some cases took forms the original investment case had not anticipated.

The relevance of this for Ethiopia's Digital Ethiopia 2025 strategy is direct. Digital Ethiopia is, in substance, the right ambition. The risk lies in sequencing. If the strategy's human capital pillar is funded and timed after its infrastructure pillar, on the implicit premise that infrastructure must come first, the country will pay the deferral cost calculated in Section 5. If the two are sequenced in the order Bangladesh and Vietnam demonstrate, human capital first

or alongside infrastructure, the strategy will produce returns commensurate with its stated goals.

For the comparable economies in the relevant class (Tanzania, Mozambique, Uganda, Senegal, the Sahel, Bangladesh's neighbors, and post-conflict economies in the Horn) the same logic applies. The IMF Article IV consultation process, the World Bank Country Partnership Framework process, and the bilateral donor strategies that shape education and infrastructure spending in these economies are all, currently, designed in ways that allow infrastructure-first sequencing to occur by default. The calibration in this paper suggests this default is expensive. The deferral cost is real, the asymmetry is structural, and the cohort that ages through the window in the meantime is not recoverable.

7. Policy implications

If the calibration holds, three implications follow for the institutions that shape education and infrastructure spending in lower-income countries.

The first concerns Ethiopia's own Digital Ethiopia 2025 strategy and its successor frameworks. The strategy's existing human capital pillar should be advanced in the budgetary sequence, not deferred until infrastructure milestones are met. The incremental cost of doing so is modest relative to the scale of the overall strategy. The cost of not doing so, calculated explicitly in Section 5, is on the order of USD 15 to USD 25 billion in foregone present-value human capital, together with an intergenerational compounding loss that no subsequent acceleration can recover. The correct macroeconomic posture treats human capital investment as the enabling condition for infrastructure returns, not as a downstream beneficiary of them.

The second is for the IMF Article IV consultation process and the analytical frameworks that surround it. Article IV consultations with economies in the relevant class, characterized by young populations, infrastructure ambition, informal-sector labor concentration, and uneven regional development, should incorporate explicit sequencing analysis of education and infrastructure investment. The current convention, in which education spending is reviewed as a discrete budgetary line and infrastructure spending as another, is not adequate to the macroeconomic significance of the sequencing decision. The Article IV process is the natural institutional location for this analysis because it is where the macroeconomic implications of competing investment paths are formally evaluated. The IMF's recent work on the macroeconomic returns to human capital in Sub-Saharan Africa, including Cherif and Hasanov on industrial policy, provides analytical machinery that extends naturally to the sequencing question.²⁶

The third is for the World Bank's Human Capital Project and the broader donor community that funds education and infrastructure programming in the relevant class of economies. The Human Capital Index is a measurement framework, not a programmatic one, and it does not

yet incorporate sequencing analysis. Adding it, in the form of an explicit measure of the deferral cost associated with delaying human capital investment until infrastructure thresholds are achieved, would represent a methodological contribution and a programmatic one. The Bank's lending operations in education and digital development are organized in separate global practices for sound institutional reasons, but the consequence is that sequencing decisions tend to be made by default rather than by analysis. A short, technical addition to the Human Capital Project's analytical apparatus could change this without requiring institutional restructuring.

A fourth implication, for the broader policy community, warrants brief attention. The "leapfrog" thesis, the claim that lower-income countries can bypass stages of development by adopting frontier technologies earlier, has shaped donor strategy for two decades. The thesis is well-founded in some applications: mobile money in Kenya, off-grid solar across East Africa, satellite-delivered education in selected pilot programs. It is also frequently misapplied. The version that holds that infrastructure investment alone permits lower-income economies to bypass the gradual accumulation of human capital that historically preceded productivity transitions is not supported by the evidence reviewed here. Bangladesh did not bypass that accumulation. Vietnam did not. The economies that have transformed most durably built foundational education first; the infrastructure investments that followed were activated by the human capital already in place. Returns from frontier technology adoption are conditioned on the absorptive capacity of the labor force, which foundational education creates and infrastructure alone cannot substitute for.

8. Limitations and the path to causal inference

A calibration paper has known limitations. We name them plainly.

The first is that this is, in the technical sense, an accounting exercise rather than a causal one. We have not estimated the impact of the GGE intervention on subsequent earnings. We have assumed, using cited elasticities from the published literature, that the intervention produces additional effective years of schooling, and we have priced those effective years at standard Mincerian returns. The accounting is rigorous within its assumptions. The assumptions themselves require empirical validation.

The path to validating them is empirical and identified. It requires a randomized controlled trial, or a quasi-experimental design with credible identification, that estimates the actual impact of the GGE lab model on measured learning outcomes and on subsequent labor market trajectories. The companion paper to this one, *Building the Last Mile Before the Pipe Arrives* (Bazzichelli, forthcoming), outlines a research agenda in which the Abdul Latif Jameel Poverty Action Lab (J-PAL) at MIT would conduct exactly such a trial in the GGE network

and in refugee settings.²⁷ The trial would generate the causal estimates that the calibration in this paper currently substitutes for with literature-derived elasticities.

The second limitation is that the calibration treats Ethiopia as a single labor market. In practice, Ethiopia has at least three distinct segments: formal urban, informal urban, and rural agricultural, each with different earnings distributions, different returns to schooling, and different probabilities of inter-segment mobility. The trajectory-shift interpretation of X depends on the probability that the intervention shifts a child across segments. We have not estimated that probability. Instead, we report X under both the narrow and trajectory-shift interpretations and leave the judgment to the reader. A more granular calibration, using Living Standards Measurement Study microdata and an explicit mobility model, would sharpen this considerably. We have not undertaken it here because the headline finding, that the order-of-magnitude return decisively exceeds the cost, does not depend on the additional precision.

The third limitation concerns the social spillover channels. The elasticities used in Section 4 are drawn from the published literature, but they are themselves estimates with confidence intervals, and they were produced in settings that differ from contemporary Ethiopia in non-trivial respects. The thirty-to-sixty-percent uplift we apply is a reasonable central case but carries its own uncertainty. A more rigorous treatment would assign confidence bounds to each spillover channel separately and propagate the uncertainty through to the aggregate estimate. We have not done so here both because the result would not differ qualitatively and because the normative argument in Section 4.1 stands independent of the precise spillover quantification.

CODA

The argument in this paper has been made in technical terms. The claim underneath it is not.

What distinguishes human capital from every other class of developmental investment is that it does not depreciate in the relevant sense. A road built in 2025 deteriorates. A fiber-optic line laid in 2025 becomes obsolete within fifteen years. A power station built in 2025 will, by mid-century, either be retrofitted or replaced. The infrastructure asset classes that absorb the largest share of development finance produce returns over fixed, bounded, and often shorter-than-planned horizons.

Human capital does not behave this way. A child educated in 2025 generates returns across her entire working life. She then generates second-order returns through her children's educational attainment, and third-order returns through the institutions her children go on to build and inhabit. The compounding is real, multi-generational, and irreversible in the direction that matters. No subsequent political shock, climate event, or technology disruption can disinvest her of the foundational skills she has acquired. A road can be bombed. A fiber-

optic line can be cut. The child who has learned to read has, for the rest of her life, learned to read.

There is something this paper's numbers cannot reach, and it is worth naming plainly at the close as at the opening. Education permanently enlarges a life: it expands what a person can think, do, and become, deepens her participation in civic life, and carries across generations in ways that no physical infrastructure replicates on its own. The calibration is a floor. The full value is larger than any accounting framework can hold.

This is the asymmetry the calibration has attempted to make visible. Investment in human capital is inexpensive and predictable in its cost; the return is high and it compounds. Deferring the investment is expensive, and the deferral cost cannot be recovered on any reasonable horizon. Neither claim depends on the precise parameterization; both follow from the underlying property of the asset.

The case for Ethiopia, and by extension for the comparable economies that share its structural features, is a case for investing in the one layer of national capacity the future cannot take back. The children who receive foundational education in Ethiopia in 2025 will be the workers, the engineers, the teachers, and the mothers of the Ethiopia of 2055. The children who do not will not be those things. No subsequent investment recovers them for those roles. The calibration has made that choice legible in dollars. The choice itself is between two Ethiopias.

We end where we began. X is a number, calibrated as carefully as the available data permits. The cost of leaving X unsolved is a generation. That cost is not one the calibration can properly footnote.

No such value can fully capture the expansion of human agency, or the dignity of a child considered in the fullness of her life.

METHODOLOGY

The quantitative projections in this paper used generative AI and computational algorithms in a limited, assistive capacity: to optimize the statistical sequencing of the Mincerian calibration, to assist with coding and syntax review, and to support the literature review that grounds the parameter choices in the published record. This use is consistent with the increasing role of AI in macroeconomic research and surveillance work at institutions including the IMF, where AI-assisted tools are now widely accepted and used to scale literature synthesis, parameter verification, and scenario analysis under human direction and with human verification. All data parameters, variable inputs, baseline earnings assumptions, structural policy recommendations, and the narrative drafting of this manuscript are the author's independent work. The platform and field data used for the calibration were collected by the author and the Girls Growing Education team in Addis Ababa between 2022 and 2025. Tool versions are documented in the author's working files.

Errors are the author's alone.

Dr. Samuele Rosa

Senior Economist | International Monetary Fund

NOTES

Footnote markers in the body text correspond to the numbered references below.

1. George Psacharopoulos and Harry Anthony Patrinos, "Returns to Investment in Education: A Decennial Review of the Global Literature," *Education Economics* 26, no. 5 (2018): 445-458. The Sub-Saharan range is drawn from the regional decomposition in Table 2.
2. Eric A. Hanushek and Ludger Woessmann, *The Knowledge Capital of Nations: Education and the Economics of Growth* (Cambridge, MA: MIT Press, 2015). See in particular the long-run growth specifications in Chapters 3 and 4.
3. T. Paul Schultz, "Returns to Women's Education," in *Women's Education in Developing Countries: Barriers, Benefits, and Policies*, ed. Elizabeth M. King and M. Anne Hill (Baltimore: Johns Hopkins University Press, 1993).
4. Girls Growing Education, "Internal Lab Analytics, 2025." On file with the corresponding author; available on request.
5. Psacharopoulos and Patrinos, 2018, op. cit.
6. Ibid., Tables 3 and 4.
7. T. Paul Schultz, "Evidence of Returns to Schooling in Africa from Household Surveys: Monitoring and Restructuring the Market for Education," *Journal of African Economies* 13, suppl. 2 (2004): ii95-ii148.
8. Måns Söderbom et al., "The Determinants of Survival among African Manufacturing Firms," *Economic Development and Cultural Change* 54, no. 3 (2006): 533-555.

9. Tassew Woldehanna, “Education, Experience and Earnings in Urban Ethiopia,” in *Education and Development in Ethiopia*, ed. T. Woldehanna (Addis Ababa: Ethiopian Economic Association, 2008); see also subsequent work using the Young Lives Ethiopia panel.
10. Hanushek and Woessmann, 2015, op. cit., Chapter 3.
11. Lant Pritchett, *The Rebirth of Education: Schooling Ain't Learning* (Washington, DC: Center for Global Development, 2013).
12. World Bank, *World Development Indicators*, 2024 release. Ethiopia GDP per capita in current US dollars, 2024.
13. Quentin Wodon et al., “Economic Impacts of Child Marriage: Global Synthesis Report,” World Bank and International Center for Research on Women, 2017.
14. Esther Duflo, “Women Empowerment and Economic Development,” *Journal of Economic Literature* 50, no. 4 (2012): 1051-1079.
15. World Bank, “Ethiopia Gender Diagnostic,” 2022.
16. World Bank / International Labour Organization modeled estimates, labor force participation rate, female (ages 15+) for Ethiopia, 2024 release. Sub-Saharan African regional average from the same series.
17. T. Paul Schultz, “Why Governments Should Invest More to Educate Girls,” *World Development* 30, no. 2 (2002): 207-225.
18. Hoyt Bleakley, “Disease and Development: Evidence from Hookworm Eradication in the American South,” *Quarterly Journal of Economics* 122, no. 1 (2007): 73-117.
19. Janet Currie and Enrico Moretti, “Mother’s Education and the Intergenerational Transmission of Human Capital: Evidence from College Openings,” *Quarterly Journal of Economics* 118, no. 4 (2003): 1495-1532; for African evidence, see Esther Duflo, “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review* 91, no. 4 (2001): 795-813, and subsequent literature.
20. World Bank, “World Development Report 2018: Learning to Realize Education’s Promise,” Washington, DC, 2018.
21. This range corresponds to the combined effect of the four channels above, taking midpoint elasticities. Higher estimates exist in the literature.
22. 3,500 active students $\times$ USD 20,000 per child = USD 70 million. The figure assumes the trajectory-shift interpretation of X applies at the midpoint of the central range.
23. Three million children $\times$ USD 4,000 per child (using the narrow private-return central estimate as the conservative case) yields USD 12 billion. Using the trajectory-shift estimate yields a substantially higher figure; we report the conservative number to avoid over-claiming.
24. Janneke Pieters and Sabina Marchetti, “Female Secondary School Stipend Programme in Bangladesh: A Critical Assessment,” UNESCO Background Paper, 2014; World Bank, “Bangladesh: Female Secondary School Assistance Project Implementation Completion Report,” various years.
25. OECD, PISA 2018 Results and subsequent rounds; see also Suhas D. Parandekar, Elisabeth Sedmik, and Akihiko Hosomi, “Why Are Students in Vietnam So Smart?” World Bank Education for Global Development Blog, June 2016.
26. Reda Cherif and Fuad Hasanov, “The Return of the Policy That Shall Not Be Named: Principles of Industrial Policy,” IMF Working Paper WP/19/74, 2019; see also subsequent IMF Sub-Saharan Africa Regional Economic Outlook publications.
27. Ginevra Bazzichelli, “Building the Last Mile Before the Pipe Arrives: Human Capital as the Real Prerequisite for AI Participation in Climate- and Conflict-Affected Regions,” working paper, Girls Growing Education, 2026. Companion paper to the present work.