

Building the Last Mile Before the Pipe Arrives

*Human Capital as the Real Prerequisite for AI Participation in
Climate and Conflict Affected Regions*

Ginevra Bazzichelli

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

Acknowledgments: I thank the Girls Growing Education field team in Ethiopia, without whose lab operations, site visits, and continued work through the Tigray conflict and its aftermath this paper would not exist. Any errors are mine.

Author's Note

I wrote this paper first, during the summer and autumn of 2025, after two years of building the Girls Growing Education (GGE) lab network in Ethiopia. I did not visit the Tigray deployment sites myself. The Tigray observations described in this paper were relayed to me by a GGE on-the-ground partner in Ethiopia who conducted the site visits during and after the 2020 to 2022 conflict, and who has been running lab operations in the region since. What the observations turned out to be, when I sat with them, was evidence that the entire sequencing framework the development finance community currently operates under has the order of operations wrong. This paper is my attempt to make that argument in the form the audience I am addressing responds to, and to name the specific research partnerships that would either confirm the argument or refute it. A companion paper, Solving for X, followed. That paper constructs the per-child calibration that this one references but does not attempt, and it substitutes literature-derived elasticities for the causal estimates that only the trials proposed here can eventually produce. Any errors in the argument that follows are mine; the credit for the underlying work in

the labs, and for the field observations in Tigray that ground the argument, belongs to the field team in Ethiopia.

Abstract

The dominant framing of the global digital divide treats it as an infrastructure problem to be solved in sequence: lay fiber and cellular networks, then connect schools, then teach digital skills, then participate in the artificial intelligence economy. That framing has shaped roughly two decades of bilateral aid, multilateral lending, and foundation strategy. In the regions where the divide is most severe, it is also the wrong sequence.

This paper argues that in climate-stressed and conflict-affected populations (northern Ethiopia after the Tigray war, the Sahel, refugee-hosting areas of East Africa, and the low-income cohort of the sub-Saharan youth population more broadly) the binding constraint on eventual AI participation is not the arrival of the pipe. It is the foundational human capital of the child who will eventually plug into it. Where infrastructure is absent, delayed, or actively destroyed, the child's window for acquiring foundational literacy, numeracy, English, and digital habits does not pause. It closes.

The argument draws on Girls Growing Education's (GGE) field deployments across 28 learning labs in Ethiopia, including in Tigray during and after the 2020 to 2022 conflict, and on a comparative reading of adjacent efforts, notably Arizona State University's SolarSPELL. We describe what was deployed, what was observed, and what the deployments imply about the design of learning infrastructure that does not wait for connectivity to arrive.

We propose a research agenda, not a product roadmap, for testing whether the inversion holds. Stanford HAI, MIT's J-PAL, and Cornell Tech are each positioned to answer one of the three open questions the argument raises. We name the question we think each is best positioned

to take on, and we sketch what a study designed around that question would need to measure. The paper closes with policy implications for the multilateral institutions that finance both education and digital infrastructure but rarely sequence the two in the order the field evidence suggests they should.

Keywords: digital divide, human capital, AI readiness, refugee education, offline learning, conflict-affected populations, Ethiopia, Tigray

1. The sequence we have been told to believe

Walk into almost any donor meeting on digital development since roughly 2015 and you will hear a version of the same story. It runs like this. The world is dividing into countries that participate in the artificial intelligence economy and countries that do not. Participation requires connectivity. Connectivity requires infrastructure. Once the infrastructure is in, content can be delivered. Once content is delivered, skills accumulate. Once skills accumulate, populations participate. The policy implication follows: build the pipes first.

The story is told with conviction because parts of it are true. Connectivity matters. Skills matter. The productivity gains from AI, over the long run, will be captured largely by populations that can use it. The problem is not with any single link in the causal chain. It is with the assumption that the chain has to be built in that order, and that the timeline over which it will be built is one the children currently in school can afford to wait for.

Consider what the sequence requires. It requires that infrastructure investment, once committed, actually reaches the population it was intended to reach. It requires that reaching happen on a timeline compatible with the childhood window during which foundational skills can be acquired. And it requires that the populations most in need of the sequence, meaning

children in the poorest, most climate-stressed, and most conflict-affected regions, are the ones the sequence eventually gets to.

In Ethiopia in 2020, all of those upstream conditions collapsed at once. The Tigray war began in November of that year. By the time the Pretoria peace agreement was signed in November 2022, UNICEF and Save the Children would document the destruction or damage of more than 7,000 schools across northern Ethiopia. The Ethiopia Education in Emergencies dashboard would later put the figure above 9,000 once Amhara and Afar were included. Roughly 2.4 million children in Tigray alone were out of school. Eighty-eight percent of schools in the region had been damaged or were occupied as shelters, military positions, or aid distribution points.¹

What does the connectivity-first sequence look like in those conditions? It looks like nothing at all. The fiber that was being laid before the war is now buried under rubble or has been ripped up for the copper. The cellular towers that were powering 4G expansion across the region have been targeted as military communications infrastructure. The schools that would have housed the computer labs that would have used the connections do not exist. The teachers who would have been trained have been displaced, killed, or are themselves teaching from tents in informal settlements. A donor meeting in Geneva or Washington that begins with the phrase "once connectivity is restored" is, in this context, a donor meeting that postpones the question of what to do for the children who are currently fourteen years old.

This paper is about those children. It is also, more broadly, about the much larger group of children (refugee, internally displaced, climate-migrated, or simply born in the wrong part of a country) whose access to eventual AI participation depends on decisions being made in the next

five to ten years by institutions whose default operating assumption is still that the pipe comes first.

The argument of this paper is that the sequence is wrong because it confuses what is necessary with what is causally upstream. Connectivity is necessary for AI participation. It is not what causes AI participation. What causes AI participation is a foundational cognitive layer, built during childhood, that gives the eventual user the capacity to make use of what arrives on the pipe. That layer can be built in advance of the pipe. It can be built during conflict. It can be built in refugee settlements. And critically, the window for building it will close on the cohort that is currently in school long before the pipe reaches them.

2. The infrastructure-first orthodoxy and why it has the timeline wrong

It is worth being precise about what the orthodoxy actually claims, because the strongest version of it is more defensible than the caricature.

The GSMA's annual Mobile Economy and State of Mobile Internet Connectivity reports are the empirical core.² They document, with care, the correlation between mobile broadband penetration and GDP growth. The World Bank's own Digital Progress and Trends report, published in 2024, updates the analysis and extends it to affordability, coverage, and quality dimensions.³ In the academic literature, the most-cited study on the causal effect of connectivity is Hjort and Poulsen's American Economic Review paper on the arrival of fast internet in Africa, which found employment and firm-level productivity gains along the East African submarine cable landings.⁴

The empirical correlations are real. The strongest version of the connectivity-first claim is that, in stable middle-income countries with functioning education systems, marginal investment in broadband infrastructure produces marginal gains in human capital outcomes. There is good

evidence for that. Hjort and Poulsen's work on the staggered arrival of submarine cables in Africa found significant employment effects in the cities the cables reached.⁴ Studies of the China Broadband Universal Service Reform, of Brazil's *Plano Nacional de Banda Larga*, of Kenya's fibre rollout, all point in the same direction.

The problem is the generalization from those settings to the settings where the divide is most severe. The implicit syllogism (connectivity drives GDP; GDP is what you want; therefore prioritize connectivity) has three flaws when transposed to the low-income, climate-stressed, conflict-affected populations that are the subject of this paper.

First, it conflates the marginal effect in functioning systems with the marginal effect in collapsed ones. The cable arriving in Nairobi reaches a city with electricity, schools, formal-sector employers, a banking system, and a population literate in at least one global language. The same cable arriving in Gambella or Kakuma reaches none of those things. The complementary inputs that turn connectivity into welfare are absent. The connection is a connection to nothing usable.

Second, it assumes that the cable will arrive on a timeline relevant to the children currently in school. In Ethiopia, ICT contributes roughly two percent of GDP, against a Sub-Saharan African average of approximately four percent.⁵ Closing that gap through infrastructure investment alone, even under optimistic assumptions about capital deployment and political stability, is a project measured in decades. The children who are nine years old today, and whose foundational learning window closes at fifteen or sixteen, do not have that timeline. The population most in need of the argument the orthodoxy makes is precisely the population the orthodoxy's timeline cannot reach in time.

Third, and most fundamentally, the orthodoxy reverses the actual causal direction in the AI era. The 2024 Stanford HAI AI Index Report and the extensive literature on AI complementarity establish that the productivity return to AI tools accrues disproportionately to workers who bring foundational reading, mathematical, and reasoning skills to the interaction.⁶ AI amplifies human capital. It does not substitute for it. A worker without foundational skills does not become productive when given AI. She becomes worse than a worker with foundational skills and no AI at all.

The orthodoxy, in other words, has the policy implication backward. If the goal is AI readiness, the binding constraint is not the pipe. It is the prerequisite literacy, numeracy, and digital fluency that determines whether the pipe, once it arrives, has anything to connect.

3. The human-capital-first inversion: theoretical grounding

The argument for inverting the sequence does not require new theory. It requires taking seriously theory that has been on the shelf for forty years.

The cleanest statement is Hanushek and Woessmann's work on what they call the knowledge capital of nations.⁷ Their finding, replicated across every major cross-country dataset, is that measured cognitive skills (not years of schooling, not enrollment rates, but skills as assessed by international instruments like PISA and TIMSS) are what predict long-run country-level growth. Years of schooling matter only insofar as they produce skills. Where they do not, they do not.

Lant Pritchett's framing of the same point, that schooling ain't learning, has been more influential in development policy because it names a specific failure mode: children being credentialed without being educated.⁸ In low-income countries, an alarming share of primary school completers cannot read a paragraph or do basic arithmetic. The response of the World

Bank's Learning Poverty framework has been to name this failure and to argue that the goal of primary education is not attendance or completion but foundational skill acquisition. The J-PAL evidence base, particularly the Banerjee-Duflo-Kremer strand on Teaching at the Right Level (TaRL), has shown that this failure is remediable, at scale, with the right pedagogical design.⁹

What follows from these two bodies of work, taken together, is straightforward. Foundational cognitive skills (literacy, numeracy, English, the habits of self-directed learning) are the layer that determines whether everything downstream, including eventual AI participation, produces a return. That layer is built during a specific developmental window, roughly ages six through fifteen. It can be built with astonishingly modest resources when the pedagogy is right. It cannot be reliably built after the window closes.

The connectivity-first orthodoxy implicitly assumes that connectivity causes learning by giving children access to better content. In stable, well-schooled, upper-middle-income settings, that assumption is defensible. In the settings this paper is about, it is not. In those settings, connectivity gives access to a great deal, most of which is not educational content, and to educational content that assumes a level of foundational literacy the child does not yet have. Connectivity without foundational capacity produces consumption of digital media, not accumulation of human capital.

This is not an argument against connectivity. It is an argument about sequence and constraint. Build the human capital layer first, using tools that do not depend on connectivity, and connectivity, when it arrives, operates on a population equipped to make productive use of it. Build connectivity first, and it operates on a population that cannot. The order is not decorative. It is the difference between the pipe delivering a return on investment and the pipe delivering entertainment.

4. The GGE Ethiopia case: what was deployed, what was observed

The empirical anchor for this paper is GGE's deployment in Ethiopia between 2021 and 2025, and in particular the subset of that deployment that took place in Tigray during and after the war. This is not an evaluation paper. We do not present randomized-controlled-trial evidence, because we do not yet have it. What we present is the observational record of a specific design hypothesis tested under specific conditions, and the outcomes that the field generated.

The hypothesis was that a single piece of hardware, preloaded with high-quality educational content, generating its own Wi-Fi network for up to ten devices, and requiring no internet connection to function, could deliver structured foundational learning in environments where the connectivity-first orthodoxy said learning could not happen. The Tigray war provided the unfortunate opportunity to test that hypothesis under exactly the conditions the orthodoxy would predict were impossible.

The hardware is what GGE refers to as the EdTech device. It comes preinstalled with the Khan Academy library and additional curated K–12 resources across mathematics, English, the sciences, and digital literacy. It requires no internet connection. Up to ten students connect from any combination of desktop, laptop, tablet, or mobile phone to the network the device itself generates. The device is compact, low-power, and can be carried by a single adult. We will return to its technical capabilities and their implications in Section 5; here, the relevant point is that it was designed from the beginning to function in the absence of the infrastructure the orthodoxy assumes.

By the end of 2025, GGE had established 28 learning labs across Ethiopia, reaching more than 22,400 students. Of these, approximately 3,500 were active users in the sense captured by GGE's internal engagement definition (students who returned to the device at least weekly and

completed at least one full learning activity per session). The labs were located across Addis Ababa, Amhara, Oromia, and Tigray, with the Tigray deployments beginning in the immediate post-war recovery period.

The Tigray deployment is the case worth examining most closely, because it stress-tests the hypothesis. During and immediately after the war, formal schooling in northern Ethiopia was disrupted at scale. UNICEF Ethiopia and Save the Children documented school closures, damage to school infrastructure, teacher displacement, and prolonged interruptions to formal instruction affecting several million children.¹ Connectivity in the region was intermittent at best. The infrastructure-first orthodoxy would predict that in such conditions, digital learning is impossible. Any device that requires a network cannot function. Any content strategy that requires cloud services cannot deliver. Any pedagogical model that requires trained in-person teachers cannot operate.

The GGE response was orthogonal. Because the device was portable, solar-capable, and carried its own content library, learning did not depend on any of the systems that had collapsed. Because the device generated its own local network, students could work in groups of up to ten without any external connectivity. Because the pedagogical model was structured around a locally trained adult facilitator rather than a certified teacher, the shortage of teachers was not a hard blocker. The response required nothing from the fiber that was not being laid, or from the schools that were not yet rebuilt, or from the electric grid that was intermittent.

What we observed, across the Tigray deployment and the broader Ethiopia program, falls into three categories.

First, attendance held. The attendance logs from labs in Tigray during the post-war recovery period showed weekly participation rates consistent with the labs in regions of Ethiopia

that had not experienced conflict. The conventional development assumption is that conflict-affected populations have disrupted attendance for years after a peace agreement, because the schools, the teachers, and the routines are not yet in place. In the GGE labs, the routines were portable, and they reformed faster than the schools did.

Second, the labs functioned with the staffing and energy budget we had projected. The facilitator model (one locally trained adult per lab, supervising rather than instructing) worked as designed. Solar charging worked. The device's Wi-Fi hotspot mode reliably served the intended cohort size. No lab failed for technical reasons attributable to the environment. In the Ethiopian context, this is not a trivial result. The prior expectation in the humanitarian education literature is that programs of this kind fail on hardware, energy, or maintenance within the first year.

Third, and most importantly: some of the students who used the EdTech device during the war years have gone on to enroll in tertiary education. This is not a randomized impact estimate. It is not, by itself, evidence that the device caused the enrollment. It is, however, an outcome that the connectivity-first sequence does not produce in this population at all, and it is one we did not have any other plausible mechanism for producing. The children in question were, in 2020, in conditions that had reduced their realistic education trajectories to subsistence-sector labor or early marriage. They are, in 2025, in college. The mechanism, as best we can reconstruct it, was that the device let them keep learning during the period when no other learning was available, and that the continuity of learning kept their academic trajectories intact through a window that would otherwise have closed permanently.

We claim no more for this observation than it can carry. It is consistent with the hypothesis. It is not a test of the hypothesis. The research agenda we propose in Section 6 is

specifically about how to convert observations of this kind into evidence that would survive review.

What we also did not observe is worth recording. We did not observe spontaneous mastery in the absence of facilitation. The labs in which facilitators were poorly trained, or where scheduling was chaotic, produced attendance without progression. This confirms one of the strongest findings in the education literature: the technology does not teach. The lab teaches. The device is what makes the lab possible where connectivity, teachers, and electricity are unreliable, but it is not itself the mechanism of learning.

5. Generalizing to refugee and conflict-affected populations: prior art and design distinctions

If the GGE deployment in Tigray works, the next question is whether it generalizes: to refugee camps, to other conflict-affected populations, to climate-migration corridors, to the broader class of settings where the connectivity-first orthodoxy fails. The comparative literature and the operational track record of adjacent efforts are useful here.

The closest comparator is Arizona State University's SolarSPELL project, founded by Laura Hosman and operating since 2014. SolarSPELL is a solar-powered, offline-capable, WiFi-emitting digital library that has been deployed in Peace Corps sites, refugee camps, and rural schools in more than ten countries.¹¹ The design philosophy is congruent with GGE's in the dimensions that matter most: no dependency on external connectivity, solar-first energy, ruggedized hardware, and a curation model that assumes the operating environment is austere. SolarSPELL is the field-tested existence proof that offline, solar-capable digital learning in fragile settings is not conceptually speculative.

The design philosophy GGE shares with SolarSPELL is the one this paper has been arguing for: offline-capable, solar-compatible, infrastructure-independent, deployable to populations that connectivity will not reach in time. On this central point, the two projects are aligned, and the existence of SolarSPELL strengthens rather than weakens the case for the inversion this paper proposes. Two independent teams, working separately on different continents, converged on the same architectural conclusion. That is the pattern that suggests the conclusion is correct.

Where the two diverge is in the design philosophy on what comes after content delivery.

SolarSPELL is, by its own description, a library. It optimizes for breadth of content, low cost per unit, and self-service access. A learner walks up to the device, finds material, reads or watches it. The model assumes the learner can navigate the library and extract value from it. In contexts where the learner is a teacher who already knows what they are looking for, this works well. ASU's evaluations in the South Sudanese refugee teacher population in Ethiopia are encouraging on exactly this dimension.

GGE optimizes for something different. The unit of deployment is not a library but a lab, meaning a structured learning environment with facilitator supervision, scheduled attendance, curriculum sequencing, and observed student progression. The content is not exhaustive; it is curated to the specific pedagogical goal of foundational skill acquisition, primarily through the Khan Academy library at K to 12 levels. The device is designed to serve up to ten students simultaneously through its integrated Wi-Fi hub, which makes the lab format economically viable at very low per-student cost. SolarSPELL and GGE are not competitors. They are complements. SolarSPELL is the library. GGE is the school-in-a-box.

There is also a technical difference that matters for the question of where this paper's argument goes next. The GGE EdTech device is designed with an on-device inference capacity that is not part of the SolarSPELL specification. In its current deployment, that capacity is used for content recommendation, adaptive sequencing, and offline speech interaction. In its near-term roadmap, the device is designed to host locally-hosted small language models that can provide tutoring interactions without any cloud dependency at all. This is technically nontrivial and it is the point at which the offline learning device stops being a container for content and starts being an AI-augmented learning environment that does not require the pipe.

This is the point at which the paper's larger argument closes its loop. If the connectivity-first orthodoxy is wrong, and the human-capital-first inversion is right, then the design category the field needs is neither the connected classroom nor the offline library. It is the offline AI-augmented learning lab. That category exists. It is being deployed. And it is being deployed in exactly the environments the orthodoxy has told the field to wait on.

The candidate sites for testing the lab model in refugee contexts are well-mapped. The Gambella regional state in western Ethiopia hosts the second-largest refugee population in Ethiopia, primarily from South Sudan, in camps administered under the Ethiopian government's Comprehensive Refugee Response Framework commitments. The Kakuma refugee complex in northern Kenya, operated in partnership between UNHCR and the Kenyan government, has been the site of prior J-PAL work on refugee education access.¹² The Cox's Bazar settlements in Bangladesh, hosting Rohingya refugees, remain among the largest and most educationally underserved refugee populations globally.

None of these is a hypothetical site. All of them have the conditions the GGE model was designed for: stable but constrained population, limited or absent formal school infrastructure,

unreliable connectivity, restricted electricity, and a large cohort of children whose foundational learning window is currently open and will not remain open indefinitely.

6. A research agenda: what Stanford HAI, MIT J-PAL, and Cornell Tech could each test

The argument in this paper is, at this stage, a position paper backed by observational field experience. Converting it into research-grade evidence requires studies of a specific kind, designed and conducted by partners with specific track records. We name three, and we describe what each is positioned to test and roughly how a study would be structured. We do not attempt to specify the trials in full here; we specify the research questions each partner is best placed to answer.

Stanford HAI is the right partner for the question this paper opens but does not close: what does AI readiness actually look like as a measurable individual-level construct, and does foundational cognitive skill acquisition produce it? The Stanford AI Index research team has, more than any other group we are aware of, developed the frameworks that turn AI-related constructs into measurable indices, and the institutional access to run them at scale.

The specific question Stanford HAI is positioned to answer is whether the foundational cognitive skills the GGE labs build, operationalized as measured performance on standard reading, mathematics, and English instruments, translate into measurable AI-tool productivity gains once the student encounters AI-augmented tasks. A study designed around that question would need three elements. First, a paired-cohort design comparing GGE-lab-exposed students with matched non-exposed students on identical AI-augmented tasks. Second, a task battery that isolates the productivity delta AI produces, rather than absolute performance. Third, a

measurement instrument for AI readiness itself, treated as a first-order object of study rather than a derivative of general education. Stanford HAI's existing work on the AI Index and on human-AI complementarity provides the methodological base for constructing such an instrument.

MIT J-PAL is the right partner for a different question, and a more difficult one: does the lab model actually produce the outcome improvements the observational evidence in Section 4 suggests it does, at a magnitude and cost that would justify scaling? This is a causal identification question, not a descriptive one, and J-PAL's methodological standing in development economics is unmatched for this specific class of question.

The specific design would be a multi-arm cluster randomized controlled trial across refugee or conflict-affected settlements (Gambella is the most natural first site, with Kakuma as an obvious second) with random assignment at the settlement level. The primary outcomes would be foundational reading and numeracy gains measured against validated instruments, with secondary outcomes on attendance, completion of structured learning sequences, and, over a longer horizon, tertiary education enrollment. The trial would include treatment arms varying facilitator training intensity, content curation, and scheduling structure, so that the design isolates which components of the lab model are causally responsible for the outcomes. Cost per learning-adjusted year would be reported alongside effect sizes, so the trial produces the input development finance actually needs when comparing alternative uses of the marginal dollar. A five-year time horizon would let the trial capture both proximal skill outcomes and the first observable downstream indicators of tertiary enrollment. J-PAL's Refugee, Displaced Persons and Migration Initiative has the operating footprint for this class of trial and has published on refugee education access in adjacent designs.¹²

Cornell Tech is the right partner for the third question, which is the one this paper has touched but not developed: what is the right architecture for an AI-augmented learning device that operates entirely offline in low-power, harsh-environment conditions? This is an engineering research question with policy implications, and Cornell Tech's Jacobs Institute has both the hardware research capacity and the applied AI infrastructure to answer it.

The specific question is what an edge-deployable, AI-augmented learning device (running small language models for tutoring, adaptive assessment, and speech interaction) can actually do at a cost and power budget compatible with the GGE deployment envelope. The research would address three concrete engineering constraints: model compression sufficient to fit useful tutoring capability on a device with the memory and compute budget of a low-cost tablet or single-board computer; power draw compatible with solar charging in Sub-Saharan operating conditions; and inference latency acceptable for real-time student interaction without cloud fallback. The output would be an engineering reference design and a set of benchmarks against which subsequent hardware iterations could be evaluated. This is not a hypothetical research direction; it is what GGE's current device roadmap requires, and it is the design problem the field will need to solve for the offline-AI category to develop beyond a single implementation.

There is also a fourth question, which we name here without proposing a partner for, because we think GGE should answer it internally in the first instance and then invite academic replication. That question is whether the cost per active student per year, currently in the range of USD 20 to USD 50 across the Ethiopian deployment, is a floor or a ceiling under expected

scaling. GGE's operating data over the next three years should produce a credible answer, which will then be relevant to the sequencing debate the rest of this paper is about.

What ties all three of these research relationships together is the framing this paper has built up to. The argument is not that any one of these trials, on its own, settles the question of sequencing. The argument is that the three questions the trials address (does foundational human capital predict AI-tool productivity; does the lab model actually produce foundational human capital under identified conditions; can the model be carried forward with AI-augmented tooling at the offline edge) collectively constitute the evidence base a serious sequencing argument requires. Absent them, the debate about pipe-first versus lab-first will continue to be conducted with reference to correlations from settings where the correlations were easy and to a syllogism whose weakest link, the assumption of a timeline, is the one nobody wants to name.

7. Policy implications

If the argument of this paper holds, three things follow for the institutions that currently fund education and connectivity in low-income, fragile, and refugee contexts.

The first is that the order of operations in funding portfolios should be revisited. A donor allocating one hundred million dollars across education and connectivity in a fragile-state portfolio should not be advised, on the current evidence, to prioritize connectivity. The child in Gambella is not waiting for the fiber. She is waiting for the lab. The lab costs a small fraction of what the fiber costs and can be operational within a fiscal quarter. The fiber cannot.

The second implication is about refugee and conflict education specifically. The international refugee architecture (UNHCR, the Global Compact on Refugees, the Comprehensive Refugee Response Framework, and the bilateral and foundation funders that operate inside it) treats digital learning as an emerging modality with promise but unclear

operating specifications. On the evidence in Section 4, the specifications are becoming clear. What is required is not another consultation. What is required is procurement at scale of the class of device the GGE deployment used, staffing of the lab model that surrounds it, and evaluation partnerships of the class Section 6 describes.

The third implication is for the universities themselves, and for the academic networks that surround them. The research agenda this paper proposes is not glamorous. It is trial work in settlements that are logistically difficult to reach, with populations that face real security and welfare concerns, on questions where negative findings are entirely possible and would themselves be valuable. The universities we have named have the standing and the resources to conduct this class of research. The question is whether the incentive structures inside academia will reward it in proportion to its policy consequence. On the current margin, we think the answer is not yet, and we think this is worth naming.

8. What we are not claiming

A paper that argues for a contrarian position has to be precise about what it is not arguing.

We are not arguing that connectivity is unimportant. It is important, and we will be among the most enthusiastic users of it when it arrives.

We are not arguing that the GGE model is the only model that works. SolarSPELL, the various Khan Academy offline distributions, and adjacent efforts by organizations including Rumie, One Laptop per Child in its later iterations, and the offline learning branches of Google.org and the Gates Foundation have demonstrated in their own ways that offline learning in fragile settings is a category, not a single product. What the GGE case adds is the lab format,

the facilitator model, and the near-term integration of on-device AI. Other models will emerge and should be tested against the same evidence standard.

We are not arguing that AI will save Ethiopian education. The opposite is closer to the truth: Ethiopian education will, with the right human-capital foundations, be ready to absorb whatever AI turns out to be, when it arrives. Foundations come first. AI second. Anyone telling refugee populations that AI tutors will substitute for teachers is selling something.

We are not, finally, arguing from a complete evidence base. The Tigray observations described in Section 4 are real, but they are not yet RCT-grade. The research agenda we have laid out in Section 6 is the work that would convert observation into evidence. We are publishing this paper before that conversion is complete because the policy decisions it concerns are being made now, and waiting for the evidence would be its own kind of failure.

9. Coda: the children who are fourteen now

The argument of this paper has been technical and methodological because the audience it is addressed to (donors, university researchers, and the multilateral institutions that shape both) responds to technical arguments made in the venues they read. The substantive claim underneath the technical argument is not a technical one at all. It is a claim about which children the sequencing decision reaches in time.

The children who were nine years old when the Tigray war began are fourteen now. The children who were in first grade in Cox's Bazar when the Rohingya crisis began are about to age out of the informal education they have received. The children in the Gambella settlements who were born after their families crossed the South Sudanese border have, in many cases, never seen a working school. None of them is going to wait while the institutional debate about connectivity-first development is resolved.

What we have observed in Ethiopia, including in Tigray, including during the war, is that a comparatively modest piece of hardware combined with a facilitator and a curriculum can keep the foundational learning window open for children whose environments have made every other approach fail. The technology is not what teaches them. What teaches them is that something is there, structured, dependable, and not waiting on a pipe.

The case for inverting the sequence is, in the end, a case about whose timeline counts. The orthodoxy's timeline is the donor's. The inversion's timeline is the child's. We submit this paper on the assumption that the child's timeline should win.

Notes

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

1. UNICEF Ethiopia, "Humanitarian Situation Reports," 2021–2023, available at <https://www.unicef.org/ethiopia>; Save the Children, "More than 2.3 million children out of school in northern Ethiopia despite peace agreement," September 2023, available at <https://www.savethechildren.net>; Education Cannot Wait, *Ethiopia Country Situation Analysis*, 2022. The figure of 9,382 damaged or destroyed schools was reported by the Ethiopian Ministry of Education in August 2022. See also the Education in Emergencies Hub Ethiopia country brief, November 2025.
2. GSMA, *The State of Mobile Internet Connectivity Report 2024* and the annual *Mobile Economy Sub-Saharan Africa* series, available at <https://www.gsma.com/r/somic>.
3. World Bank Group, *Digital Progress and Trends Report 2024*, Washington, DC: World Bank, 2024.
4. Jonas Hjort and Jonas Poulsen, "The Arrival of Fast Internet and Employment in Africa," *American Economic Review* 109, no. 3 (2019): 1032–1079.
5. GSMA, *Mobile Economy Sub-Saharan Africa 2024*; World Bank, Ethiopia data portal, <https://data.worldbank.org/country/ethiopia>.
6. Nestor Maslej et al., *The AI Index 2024 Annual Report*, Stanford Institute for Human-Centered AI, Stanford University, April 2024, and the 2025 update of the same report. See in particular the regional concentration analysis in Chapter 4.

7. Eric A. Hanushek and Ludger Woessmann, *The Knowledge Capital of Nations: Education and the Economics of Growth* (Cambridge, MA: MIT Press, 2015); see also their summary in "Economic Growth in Developing Countries: The Role of Human Capital," Hoover Institution working paper, 2020.
8. Lant Pritchett, *The Rebirth of Education: Schooling Ain't Learning* (Washington, DC: Center for Global Development, 2013); see also Karthik Muralidharan et al. on India ASER data.
9. Abhijit Banerjee, Rukmini Banerji, James Berry et al., "From Proof of Concept to Scalable Policies: Challenges and Solutions, with an Application," *Journal of Economic Perspectives* 31, no. 4 (2017): 73–102; J-PAL evidence summary, "Teaching at the Right Level to accelerate learning," 2023, available at <https://www.povertyactionlab.org/evidence-effect/teaching-at-the-right-level>.
10. Julián Cristia et al., "Technology and Child Development: Evidence from the One Laptop per Child Program," *American Economic Journal: Applied Economics* 9, no. 3 (2017): 295–320; J-PAL Education sector summary on computer-assisted learning, 2023.
11. Laura Hosman et al., "The SolarSPELL Initiative: Digital Libraries for Offline Educational Access," various publications, Arizona State University, 2018–2024; see project documentation at <https://solarspell.asu.edu>.
12. J-PAL and Community Jameel, "Expanding access to higher education for refugees," 2018, available at <https://www.povertyactionlab.org/blog/5-7-18>.
13. J-PAL, "New J-PAL research and policy initiative to test and scale AI innovations to fight poverty," 2024.
14. Ginevra Bazzichelli, "Solving for X: Human Capital, the Marginal Child, and the Macroeconomic Case for Education Before Infrastructure in Ethiopia," working paper, *Girls Growing Education*, 2025. Companion paper to the present work; constructs the per-child valuation framework this paper references.