Policy Research Working Paper 9588

Will Every Child Be Able to Read by 2030?

Defining Learning Poverty and Mapping the Dimensions of the Challenge

João Pedro Azevedo
Diana Goldemberg
Silvia Montoya
Reema Nayar
Halsey Rogers
Jaime Saavedra
Brian William Stacy

WORLD BANK GROUP
Education Global Practice
March 2021

Policy Research Working Paper 9588

Abstract

In October 2019, the World Bank and UNESCO Institute for Statistics proposed a new metric, Learning Poverty, designed to spotlight low levels of learning and track progress toward ensuring that all children acquire foundational skills. This paper provides the technical background for that indicator, and for its main findings—first, that even before COVID-19, 53 percent of all children in low- and middle-income countries could not read with comprehension by age 10, and second, that at pre-COVID-19 trends, the Learning Poverty rate was on track to fall only to 44 percent by 2030, far short of the universal literacy envisioned under the Sustainable Development Goals. The paper contributes to the literature in four ways. First, it formally describes the new synthetic Learning Poverty metric, which combines the dimensions of learning with schooling and thus reflects the learning of all children, and it presents, for the first time, standard errors associated with the proposed measure. Second, it documents how this indicator is calculated at the country, regional, and global levels, and discusses the robustness associated with different aggregation approaches. Third, it documents historical rates of progress and compares them with the rate of progress that would be required for countries to halve Learning Poverty by 2030, as envisioned under the learning target announced by the World Bank in 2019. Fourth, it provides heterogeneity analysis by gender, region, and other variables, and documents learning poverty's strong correlation with metrics of learning for other ages. These results show that the Learning Poverty indicator, together with improved measurement of learning, can be used as an evidence-based tool to promote progress toward all children reading by age 10—a prerequisite for achieving all the ambitious education aspirations included under Sustainable Development Goals 4.

This paper is a joint product of the World Bank Education Global Practice and the UNESCO Institute for Statistics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The authors may be contacted at jazevedo@worldbank.org.

Will Every Child Be Able to Read by 2030?

Defining Learning Poverty and Mapping the Dimensions of the Challenge[1]

João Pedro Azevedo, Diana Goldemberg, Silvia Montoya, Reema Nayar,
Halsey Rogers, Jaime Saavedra, and Brian William Stacy

JEL Classification Codes: I21, I24, I25; O15; E24
Keywords: Education; Learning; Schooling; Human Capital, Learning Assessments; SDGs

---

[1] *Many World Bank colleagues contributed to the ideas and data summarized in this paper. Together with some of the authors, Michael Crawford, Syedah Aroob Iqbal, Sergio Venegas Marin, Harry Anthony Patrinos, Shwetlena Sabarwal, and Jason Allen Weaver served as the Learning Target Working Group that conceptualized the indicator and the global target. Regional and global colleagues who helped construct and validate the indicators and its inputs included Kristoffer Gustav Bjarkefur, Marguerite Clarke, Natasha De Andrade Falcão, Ning Fu, Tihtina Zenebe Gebre, Koen Martijn Geven, Laura Gregory, Marcela Gutierrez Bernal, Maria Jose Vargas Mancera, Ryoko Tomita Wilcox, and Hongxi Zhao, in addition to Aroob and Sergio. Within the World Bank, the work was carried out under the direction of Jaime Saavedra and Omar Arias. The team is also grateful to Adolfo Imhof and the Foresight and innovation team of the UNESCO Institute for Statistics for their comments and strong collaboration in building the harmonized dataset; to Luis Crouch, Ambar Narayan, and Nobuo Yoshida for helpful early discussion; and to Simeon Djankov and Elizabeth Ninan, in addition to Ambar and Luis, for serving as peer reviewers. We are also grateful to those whose dedicated work produced all the learning data used in this paper—the national governments and the teams at the IEA (PIRLS and TIMSS), UNESCO (LLECE), Confemen (PASEC), and SACMEQ, as well as the OECD; and the staff from the EMIS departments of all Ministries of Education, without whom none of this work could be possible. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.*

1. Introduction

In October 2019, the World Bank and UNESCO Institute for Statistics launched a new metric, Learning Poverty, which highlighted that 53% of all children in low- and middle-income countries were not able to read an age-appropriate text with comprehension by age 10 (World Bank 2019). This paper provides the full technical background and main results for the learning poverty metric, as well as robustness tests, heterogeneity analysis, tests of external validity, and extensions.

Before proceeding to the technical details, it is important to explain the rationale for the metric and preview the main findings that it generates. In the realm of international education goals, ensuring that all children acquire basic reading skills should not seem like much of a stretch. The Sustainable Development Goals, which every UN member signed onto in 2015, embody much higher global aspirations for education. SDG 4 makes the following commitment: by 2030, the signatories will "Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all." The various targets under this goal cover the educational landscape, starting with universal access to quality ECD and preschool and extending to equal access to affordable university education. But the very first of these commitments is Target 4.1, which is to "ensure that all girls and boys complete free, equitable, and quality primary and secondary education leading to relevant and effective learning outcomes." In other words, the world has committed to achieve universal completion of *secondary* school for all youth — and with meaningful learning — by 2030.

But given the depth of the "learning crisis" in many low- and middle-income countries (Pritchett 2013, UNESCO 2017, World Bank 2018a), there are reasons to question whether this target is feasible and whether it will be a useful motivator to propel the required actions (see also UNESCO 2019). This was true even before the COVID-19 pandemic hit, and it is even more so now, as the school closures and global recession triggered by the pandemic interrupt student learning and reduce attachment to schooling (Azevedo 2020, Azevedo et al 2020, World Bank 2020).

The learning poverty metric is designed to spotlight one fundamental skill at the core of the SDG aspirations: the ability to read by age 10 with at least a minimum level of comprehension. Ensuring that all students read with comprehension is essential to achieving the ambitious SDG targets and to building human capital. Children need to learn to read so that they can read to learn. Those who do not become proficient in reading by the end of primary school often cannot catch up later, because the curriculum of every school system assumes that secondary-school students can learn through reading. Reading is a gateway to all types of academic learning. And intuitively, a target of "every child reading by age 10" seems attainable. In high-income countries, 90% of all children learn to read with comprehension before the end of primary school, and for the highest-performing countries, the figure reaches 97% or more. While it may take decades to build up an entire high-quality education system, teaching children to reach a minimum proficiency in reading should require much less time. Finally, reading proficiency can serve as a good proxy for (contemporaneous) foundational learning in other subjects, particularly at the level of the educational system. (See Annex 1.)

But more than aspiration and intuition are needed to guide action, which is why the World Bank and UIS developed the learning poverty measure. This paper contributes to the literature in four ways. First, it explains this new synthetic measure, which combines learning with schooling, thus capturing learning for all children and not only of those currently in school. Second, it shows how the indicator was generated using newly combined data to measure how far the world is from achieving the target of all children reading by age 10. Third, it documents the rate of progress that would be required for countries to halve learning poverty by 2030—the target year for the SDGs—and compares it to historical rates of progress. Fourth, it provides detailed analysis documenting the robustness and external validity of the learning poverty indicator, as well as the insights that can be gained by disaggregating by gender, region, and other variables. To develop these estimates, we combine data from 100 countries, accounting for 81% of children worldwide and 80% of children in low- and middle-income countries, using internationally comparable learning thresholds produced as part of the Global Alliance to Monitor Learning (GAML) led by UIS (UNESCO Institute for Statistics). The results, documented in detail in this paper and summarized briefly in World Bank (2019), are sobering:

- **More than half of children in low- and middle-income countries have not achieved minimum levels of proficiency in reading by age 10, or in most cases by the end of primary.** An estimated 53% of children at or near the end of primary school age are not yet able to read a short, age-appropriate story with comprehension. By labeling this deprivation as "learning poverty," we hoped to give focus to those children left behind and to accentuate just how important achieving at least a minimum proficiency in reading ability is as a vehicle to a productive, fulfilling life in the modern world.
- **At the rates of progress seen so far this century, the goal of ensuring that all children can read by 2030—in other words, reducing the rate of learning poverty to zero—is far out of reach.** While the share of children who are "learning-poor" has been declining, the pace of progress is far too slow to ensure that all children will be able to read by 2030. We estimate that under "business as usual"—that is, with progress at the rate we saw during 2000-17—44% of children in 2030 will still be unable to read at age 10. This indicator is an early warning that all the education-related SDGs are in jeopardy, and that grounding aspirations in reality requires a more plausible medium-term target. The high global learning poverty rate is also an early warning for countries with low Human Capital Index scores (Kraay 2018). As shown below, children who suffer from learning poverty typically end up with low levels of secondary-school learning as well; this limits their countries' ability to improve on the learning-adjusted years of schooling metric that is a major component of the Human Capital Index (Filmer et al 2018).
- **Even if countries reduce their learning poverty at the fastest rates we have seen in recent decades, the world will not come close to attaining the "every child reading" goal by 2030.** The simulations in this paper document that even if every country were to reduce learning poverty like the top performers over the 2000-15 period—meaning that they matched the rates achieved by countries at the 80th percentile of the regional distribution of gains—the global learning poverty rate can be reduced from 53% in 2015 only to 27% in 2030.[2] Stated differently, if every low- and middle-income country ramped up its efforts to address learning poverty and doubled or tripled its historical rate of progress, it would be possible to cut the global learning poverty rate in low- and middle- income countries by nearly half.

[2] For EAP and SAR, the two regions that lack sufficient data on annualized progress, we have used the global values to simulate rates of progress. Note also that in this simulation, countries that have achieved above the 80th percentile during 2000-15 are assumed to sustain their higher rates of progress.

- **All of these findings are based on data from *before* the COVID-19 pandemic—meaning that the situation is now even worse than indicated by these estimates.[3]** While the data sources used in this analysis are not yet available for 2020, nor will they be available for at least another year or two, there is no question that the levels of learning poverty are now higher, and the recent trends worse than reported here.

In summary, this analysis documents the magnitude of the learning crisis in a key dimension of foundational skills—basic literacy—and it shows that historical rates of progress were far too slow, even before COVID, to achieve meaningful progress toward global and national goals. The learning poverty indicator is simple and intuitive enough that it can be used as a focal indicator for global and national campaigns to change that trajectory. Indeed, it is already influencing the World Bank's operational engagement:[4] in October 2019, the World Bank announced a corporate commitment to support countries to "by 2030, reduce by at least half the share of children in low- and middle- income countries who cannot read by age 10." The intent of this learning target is to promote tangible progress toward the SDGs and improved Human Capital by focusing on medium-term learning goals and motivating immediate action to improve foundational skills.[5] Our analysis shows that this intermediate goal was already highly ambitious when announced, requiring the global rate of progress to increase to nearly triple its 2000-2017 rate. The effects of COVID will make it even harder to achieve, but they also underline how important it is to have a summary indicator like learning poverty to track progress in building more effective and equitable basic education systems in the wake of the crisis.[6]

In addition to providing the full technical analysis underlying the summary results presented in World Bank (2019), this paper presents the accompanying standard errors for the estimates, explores the robustness of its regional and global aggregation approach, external validity, and heterogeneity by variables of interest, among other features of the measure. The rest of this paper is organized as follows. Section 2 defines the learning poverty indicator, and describes the methodology used to construct it and Section 3 describes the data used. Section 4 presents the learning-poverty estimates for the population of children in low- and middle-income countries and other country groupings, and examines the robustness of this global measure. It also unpacks the measure into its different components and discusses the heterogeneity of the measure within countries. Section 5 presents simulations of the likely progress by 2030 under business-as-usual and high-case scenarios and compares them to the medium-term learning target adopted by the World Bank as a feasible stretch target—halving learning poverty by 2030. Section 6 concludes.

[3] Azevedo (2020) simulate that in a pessimistic scenario, learning poverty in low- and middle-income countries could increase from 53% to 63% due to COVID-19.
[4] World Bank (2019).
[5] This paper is focused on defining learning poverty and mapping trends and levels of the indicator, providing the rationale, data, and methodology underpinning the indicator and the learning target. A parallel policy note discusses the policy interventions that the World Bank is using to support country efforts toward the achievement of their national targets, and toward this global learning target (Crawford and others Forthcoming), while another companion paper provides additional analysis and proposes an inequality sensitive extension of the learning poverty measure (Azevedo 2020).
[6] See Azevedo and Montoya (2021) for a brief discussion on how the learning poverty measure can be of particular value to help countries focus their education policy response to COVID-19.

2. How we measure Learning Poverty: Definition and methodology

This section defines learning poverty and describes the methodology used to operationalize it.

Defining Learning Poverty

At the country level, as per World Bank (2019), we define learning poverty as the percentage of 10-year-olds who cannot read and understand a short passage of age-appropriate material—in other words, those who are below a "minimum proficiency" threshold for reading. Following Azevedo (2020), this measure can be defined as the union of two deprivations: 1) schooling deprivation and 2) learning deprivation. A child is considered schooling-deprived (SD) if he or she is of primary school age and out-of-school.[7] The dimension of learning deprivation (LD) applies only for children *in* school, and identifies those pupils who are below this minimum proficiency (BMP) level for reading, as measured in standard learning assessments. The final learning poverty measure combines the two dimensions in a single indicator.[8]

This "union approach" to measurement reflects the choice that, as presented in the SDGs, all age 10 children must be both in school *and* learning (see the more detailed discussion on each deprivation dimension below). This gives us the following formula for learning poverty:

$$LP = SD + [(1-SD) \times LD] \qquad (EQ.1)$$

where:

LP = Learning poverty

SD = the schooling deprivation dimension, which captures the share of children of primary-school age who are out of school; this dimension is reflected by the indicator of Out-of-School children or OoS. This dimension is linked to the indicator 4.1.4 from the SDG 4 thematic framework.

LD = the learning deprivation dimension, which captures the *s*hare of children at the end of primary who are below the minimum proficiency level (MPL) for reading, as defined by the Global Alliance to Monitor Learning (GAML) in the context of the SDG 4.1.1b monitoring, and observed by the indicator BMP (for "below minimum proficiency")

The choice of a union approach to aggregate the deprivation dimensions of this measure implies that all schooling-deprived children are regarded as being learning-deprived, or below the minimum proficiency level for reading.

[7] The indicator used to capture school deprivation is the complement to the UIS Adjusted Net Enrollment Rate, Primary. The adjusted net enrollment is defined as the number of pupils of the school-age group for primary education, enrolled either in primary or secondary education, expressed as a percentage of the total population in that age group. Some children of primary school age might enter primary school early and advance to secondary school before they reach the official upper age limit of primary education. The Net Enrollment Rate does not include those children, underestimating the number of children who receive a full course of primary education. To overcome this limitation, an adjusted net enrollment rate in primary education can be calculated.

[8] See Azevedo (2020) for a more formal discussion, including the main axiomatic properties of this measure.

Given this formulation, countries can improve their performance on this indicator in two ways: (1) by strengthening the quality of learning in their systems, and in particular by focusing on raising proficiency levels for children below the minimum proficiency threshold to at least this minimum level; and (2) by expanding coverage and bringing their primary-school-age out-of-school population into the system (as long as at least some of those children in school learn enough to exceed the minimum proficiency threshold).

The remainder of this section explains how this measure is implemented to produce our global estimate.

Schooling Deprivation: Identifying out-of-school children

The first element of learning poverty is schooling deprivation. As discussed above, this element reflects the belief that all primary-age children should be learning in schools of some type, a belief that every country has enshrined in law and that is enshrined in the SDGs. In addition to fulfilling a universal right and serving as a necessary condition for sustained learning, schooling offers many benefits beyond learning. As the COVID-19 school closures have shown, schools have numerous other functions that contribute to children's health and well-being[9]—such as promoting safety, nutrition[10], and socialization, and facilitating parents' labor market participation[11]—and at the macro level schooling can help build social cohesion,[12] democracy, and peace.[13] All those complementary functions mean that schooling has value over and above the measured cognitive learning that it leads to, and they justify including schooling deprivation in the concept of learning poverty.[14]

Beyond reflecting these social values and goals, including schooling deprivation as a dimension of learning poverty also creates better incentives for policymakers than would a measure based only on learning measured in school. It gives countries in which enrollment is not universal the ability to reduce learning poverty by improving access,[15] and penalizing countries that are only able to provide quality education to a smaller fraction of their school-age population.

Learning Deprivation: Identifying reading proficiency

When we speak of a child "attaining minimum reading proficiency," we mean that the child has the ability to read and understand a short passage of age-appropriate material, whether a simple story or non-fiction narrative of a few paragraphs.

[9] UNICEF 2020; WFP 2020.

[10] Adelman, Gilligan, and Lehrer 2008; Bhutta et al. 2013.

[11] Blau and Robins 1988; Blau and Currier 2006.

[12] Easterly, Ritzen, and Woolcock 2006.

[13] Khan 2016.

[14] For further discussion on the implications of combining an ordinal (binary) variable and a cardinal variable in a single multidimensional measure, see Azevedo (2020).

[15] This construction also penalizes countries that might try to improve their learning poverty rate by encouraging worse-performing students to drop out of the system. Ideally, we would want to go a step further and control for any potential selection bias among enrolled children that might take place when the (in-school) learning assessment is administrated, by assigning zero learning to all children who were in the original sample frame but did not take the assessment. Unfortunately, data to implement this adjustment is not systematically available across learning assessments. Recognizing this possible selection effect is particularly relevant for countries considering the use of this measure to track progress over time and improve accountability.

To operationalize this concept, the World Bank collaborated closely with the UNESCO Institute for Statistics (UIS), which has the mandate to lead the SDG monitoring process in education. The UIS leads the Global Alliance to Monitor Learning (GAML), which in 2019 agreed on the following definition of the minimum proficiency level in reading at the end of primary (MPL)[16]

> "*Students independently and fluently read simple, short narrative and expository texts. They locate explicitly-stated information. They interpret and give some explanations about the key ideas in these texts. They provide simple, personal opinions or judgements about the information, events and characters in a text*." (UIS and GAML 2019)

In addition to this nutshell statement, which is intended to be accessible to the nonexpert, the GAML has also proposed a common terminology to describe classifications in the context of the MPL. This is a critical first step toward linking cross-national and national learning assessments with a common benchmark.

Equating across assessments

The next step is to equate other international and national assessments to this benchmark. As an example, on the international Progress in International Reading Literacy Study (PIRLS) of 4th-graders, "proficiency" is equated to reaching at least the Low International Benchmark in reading—or a score of 400. According to the PIRLS 2016 documentation, achieving this score signals that "when reading predominantly simpler *Literary Texts*, students can: Locate and retrieve explicitly stated information, actions, or ideas; Make straightforward inferences about events and reasons for actions; Begin to interpret story events and central ideas" (IEA 2016). Similarly, for "predominantly simpler *Informational Texts*, students can: Locate and reproduce explicitly stated information from text and other formats (e.g., charts, diagrams); Begin to make straightforward inferences about explanations, actions, and descriptions."

PIRLS is the major global primary-age assessment focused on reading, and if all countries participated in it, the task of constructing global estimates of minimum proficiency would be trivial, as it would require aggregating results from a single cross-national assessment.[17] However, most countries participating in PIRLS are high-income, and only a small minority of low- and middle-income countries participate in the assessment. One of the main contributions of the GAML process is that it has overcome this data gap by benchmarking several major cross-national assessments—and increasingly national learning assessments as well—against the standard.

To incorporate these other assessments into the analysis, we need to harmonize their benchmarks of reading proficiency with the GAML definition. For each new assessment incorporated into the database, the harmonization process requires looking at the definitions of each level of proficiency and selecting the one that maps most clearly into this definition. For PIRLS, the minimum proficiency level is Level 2—which, as noted above, represents the Low International Benchmark—whereas for the PASEC regional

[16] The GAML is designed to improve learning outcomes by supporting national strategies for learning assessments and developing internationally comparable indicators and methodological tools to measure progress towards key targets of SDG 4. It was established by the UNESCO Institute for Statistics, which also hosts the Secretariat.

[17] PISA, the best-known international student assessment, tests the competencies of 15-year-olds. Since our interest is in assessing reading ability in late primary school, that test comes too late in the children's development to be the primary source for this exercise. However, later in this analysis, we will use PISA as a robustness check on our core findings.

assessment of West and Central Africa (to take one example), it is Level 4. This level then is used to calculate the reading proficiency rate for that country, which the share of students reaching at least that level of proficiency.

This process of equating proficiency levels on different assessments to the GAML definition is not straightforward. Even the long-running regional assessment initiatives like PASEC (West and Central Africa) and LLECE (Latin America and the Caribbean) use different definitions and a different number of levels than other assessments like PIRLS, and those might not even be the same over time. Their test-development methodologies and test administration procedures also vary. Moreover, because not all countries participate in global or regional assessments, for some major countries we rely on their interim reporting using their national assessments; equating these assessments is even more challenging. For these reasons, the World Bank and UIS did the mapping using a combination of descriptor matching and empirical triangulation (see Table 1 for the minimum proficiency cutoff value used for each assessment).

Navigating age differences

Among the differences across assessments, one important point concerns the age at which children are tested. The reference age for our exercise is age 10, for reasons discussed above. However, all learning assessments used in this analysis are sampled based on specific grades rather than age.[18] PIRLS and TIMSS are administered in Grade 4, meaning that the average student assessed is indeed 10 years old, but this is not the case for the regional assessments. (See Table 1 for the grade and average student age for each assessment.) PASEC and LLECE are administered in Grade 6, so the average age in those assessments is 12.8 and 12.4, respectively.[19] The national assessments are administered at different grades, so to incorporate those assessments, we chose for each country the grade between 4 and 6 (inclusive) for which relevant and reliable data were available. This is consistent with the SDG monitoring by UIS and GAML, which lists "End of Primary (or Grades 4 to 6)" as the relevant age category for the end-of-primary students (SDG 4.1.1b).

[18] The only learning assessment which samples a specific age group is PISA.

[19] The 6th-graders tested in the SACMEQ assessment are even older, at an average age of 13.5. There are a few exceptional cases in which PASEC, PIRLS, and TIMMS were implemented in a different grade from the standard grade at the request of the tested countries, but this is not recommended nor encouraged, as it undermines the comparability.

Table 1 Assessment data used in constructing the consolidated global dataset

Assessment	**Minimum proficiency level (MPL)[1]**	**Grade(s) assessed**	**Most recent year**	**Number of countries** (total)	**Number of countries** (low- & middle-income after 2011)	**Total student population represented** (low- & middle-income after 2011) *(millions)*	**Mean Age[2]**
PIRLS	Low International Benchmark (400 points)	4	2016	62	15	58	10.1
TIMSS[3]	Low International Benchmark (400 points)	4	2015	65	7	17	10.1
LLECE[4]	Level 3 (514 points)	6	2013	17	15	47	12.4
PASEC[5]	Level 4 (595 points)	5 and 6	2014	17	13	33	12.8
SACMEQ[6]	Level 5 (510 points)	6	2013	14	-	-	13.5
National assessments	*Varies by country*	4, 5, or 6	2017 *(Varies by country)*	15	12	281	*Varies by country*

Notes: (1) For all cross-national assessments other than TIMSS and LLECE, Minimum Proficiency Levels (MPLs) for regional and international assessments are taken from the revised UIS proposals prepared for consideration by GAML meeting in August 2019. For National Assessments: UIS and WB staff estimates; (2) Mean age is for the total population of students that took the test; (3) For TIMMS, we have used science scores, the subject and threshold level with the highest cross-country correlation with PIRLS; (4) Using the SERCE scale for both SERCE and TERCE rounds; (5) For the Democratic Republic of Congo, PASEC data are for 5th-graders in 2010; and (6) SACMEQ was used only for estimating changes in learning poverty, not levels.

Our estimates of learning poverty can therefore be considered as a lower bound on the share of children who cannot read proficiently by age 10, given that in some countries students will have had an extra year or two to learn to read after age 10 before they are assessed. We do not adjust for age, because our goal is to develop a global measure of the proficiency of children at age 10, rather than to focus on cross-country differences in achievement.

Another way of thinking about this is to distinguish between the construct and indicators we use to measure learning poverty:

- The construct of “all children reading by age 10” is an *ideal* that embodies normative statements about both learning and access. To achieve it, not only should all children be reading proficiently after 3 full years in primary education, but they should also have entered school at age 6 or 7.
- By contrast, the actual indicators used to measure learning poverty are based on grade rather than age, as noted above. Since the assessments are of 4th- through 6th-graders, the children tested will have had at least 3 to 5 years in school to reach what, according to the ideal, should

be an age-10 minimum proficiency, or even the entire primary-school-age segment for the out-of-school indicator.

Therefore, the learning poverty results presented below are likely a conservative estimate of the extent of the literacy challenge.

Reference year and reporting window

Estimating the current share of children in learning poverty requires deciding how to define "current." In practice, given that the out-of-school information is collected regularly, how current the indicator can be is largely determined by the availability of the below-minimum-proficiency (BMP) indicator that measures learning deprivation.

We have chosen to set the reference year for the current estimates as 2015, and to include results of any qualifying assessments administered within four years before or after that year (2011 to 2019).[20] In other words, the global estimate given below is labeled as 2015 but relies on assessments distributed over a nine-year window. This decision is driven by availability of learning data. The international and regional assessments in the database are carried out every 3 to 4 years, and even where assessments have been implemented recently, there is a lag of a couple of years before the data becomes available. The most recent year of the four global and regional assessments ranges from 2013 (for LLECE) to 2016 (PIRLS). This might suggest creating a narrower window around 2015, but some countries that are not covered by those assessments have good data available only for 2011-12, and a handful of other countries have data from as recent as 2017. This approach allows us to produce a global estimate based on data from 62 countries representing 80% of the population of low- and middle-income countries. Using the same temporal window, we have data for 100 countries and 81% of the population, including high-income countries.

Note that this range is intended as a moving window. For future rounds of estimates, we plan to use the same nine-year window around the new anchor year. PASEC, TIMSS, and LLECE were all implemented in 2019, and PIRLS will follow in 2021, providing a wealth of new data to draw on. New rounds of national assessments will be available for some of the countries that have not applied the international assessments, and their coverage might allow us to re-assess the ideal reporting window. Going forward, it will be critical to ensure the temporal comparability of learning assessments, and in particular the national assessments; without such comparability, a meaningful measure of temporal progress will be impossible.

After we present the results, we will briefly discuss and show the robustness of our global estimate to different reporting windows.

Regional Reporting and Global Aggregation

Because one of the main objectives of this paper is to provide a global picture of learning poverty, it is critical to discuss how the global aggregate is constructed. This would be a trivial task if our indicator were available for all countries for a specific reference year within a reporting window. However, this is

[20] As previously discussed, this approach follows similar criteria used in monitoring of global monetary poverty, which faces similar constraints on availability of recent data.

not the case. As discussed below, there are substantial data gaps, and those gaps are not uniformly distributed across regions or income levels (see Table 2 for more details).

To overcome this limitation, and following the practice used for other SDG indicators, we aggregate our global number as a population-weighted average of the regional learning poverty rates. Implicitly, this is equivalent to imputing the relevant regional weighted average to all countries that lack the necessary data.[21] Thus, for example, the learning poverty rate for Sub-Saharan Africa is a population-weighted average based on 46% of children in the region (see Table 6), computed using data for 17 countries. However, in the global aggregate, all 48 countries in Sub-Saharan Africa are represented, as the regional learning poverty rate is applied to all 123 million children in the 10 to 14 age cohort in the region.

Hence, the learning poverty rate for each region is the population-weighted aggregate for the countries from that region for which there are assessment and out-of-school data that meet the criteria explained above. Regional figures are reported only if at least 40% of the region's population of reference is covered by actual data.[22] This protocol is aligned with those of other global indicators under the SDGs, such as the International Poverty line, which also uses a 40% regional-population coverage threshold for reporting regional aggregates (World Bank 2018b). [23]

3. Data

This section describes in detail the data used to calculate learning poverty rates and the changes in the rate over time.

Learning assessment data

The consolidated learning database draws on assessment data from 5 major international and regional student assessments administered since 2000, as well as selected national assessments from the same period that have been benchmarked against the GAML descriptors. Below, we briefly explain each source, the rationale for using it, and the judgments involved in equating proficiency levels.[24] Table 1 (above) summarizes the cutoff level used to define proficiency, the grade level assessed, most recent year, the number of countries, and total population covered by assessment program. Further details on each national assessment are available in our GitHub repository (see Annex 10 for more details).

In the few cases (for countries outside of the Latin America and Caribbean, or LAC, region) in which a given country had administered multiple types of assessments, we applied a hierarchy of assessments in the order listed below. Specifically, if we had PIRLS data for Grade 4 for a country (in the relevant time

[21] This choice is fully aligned with other global aggregation methodologies used in the SDGs, such as the Global Monetary Poverty using the International Poverty line (World Bank 2018b).
[22] Within the conditions previously specified—that is, GAML comparability and the 8-year window centered on 2015.
[23] In our work, we do not try to account for any temporal differences caused by the fact that not all countries have learning data for the same years. This is in contrast with the approach in some other exercises; for example, the Global Poverty monitoring exercise uses macro numbers such as private consumption to interpolate national poverty numbers prior to the regional and global aggregation.
[24] The GAML created some initial mappings between three regional assessments (PASEC, LLECE, and SACMEQ) as part of the SDG monitoring process. These have been updated during GAML workshops in early April 2019, and the revised proficiency level descriptors were discussed in the GAML meeting in Yerevan in August 2019. These are included in Table 2. While we typically use these thresholds, we triangulate them with other data where possible.

period), we used those data to estimate reading proficiency. If not, we would proceed next to the relevant regional reading assessment (PASEC or SACMEQ); if the country had no regional assessment data, we would proceed to the TIMSS Science or Math assessment; and so on.[25] For countries in the LAC region, LLECE was used as the preferred assessment, for reasons discussed below. The hierarchy of assessments is as follows:

1) *Global Assessments – Reading*

    a) *PIRLS (global):* As noted above, PIRLS is the anchor assessment used for this database. Of the major cross-national learning assessments, it is the one that is administered to children at roughly the target age: it assesses a random sample of Grade 4 students in each country, and the average age of tested children is 10.1 years. Proficiency is defined as reaching the Low International Benchmark, which means a score above 400 points.

2) *Regional Assessments - Reading*

    a) *LLECE (Latin America and the Caribbean):* LLECE has implemented three rounds of regional assessments in Latin America and the Caribbean. The most recent round for which we have data, the third round, was carried out in 2013 and covered 13 countries. We have used data from both the second and third rounds, also known as SERCE (2006) and TERCE (2013). Given the dual objective of estimating a global learning poverty rate for 2015 and simulating the expected rate of progress until 2030, we have used TERCE results expressed in the SERCE scale, as described and reported by OREALC/UNESCO (2014). For all countries in the LAC region, we also used LLECE as the preferred source of learning assessment information even when PIRLS data were also available for the country, as was the case for Chile, Colombia, and Honduras. The reasons for this choice were: (1) to increase within-region comparability; (2) to increase the data available to determine trends, given that there are more historical observations available for the LLECE assessments (Chile and Honduras participated in only one PIRLS each, and Colombia participated in two rounds of PIRLS with a 10-year gap between them—2001 and 2011); and (3) to avoid some data comparability problems (as in Honduras, which implemented the PIRLS 4th-grade exam for 6th-grade students). Based on the proposed revisions to the minimum proficiency threshold for this assessment, an examination of the proficiency-level criteria, and triangulation with other data sources, we defined minimum reading proficiency as reaching Level 3 on TERCE (in the SERCE scale).[26]

    b) *PASEC (West and Central Africa):* PASEC has also carried out several rounds of data collection in Francophone African countries. The most recent round was carried out in 2013/2014. We used

---

[25] We do not use EGRA assessments, for three reasons. First, EGRAs are typically administered on a younger population or at lower grades than the age 10/grade 4 cohort. Second, even in the case of the two countries for which nationally representative 4th- and 5th-grade EGRAs are available, the proficiency level descriptors for those assessments cannot be mapped as of now to the GAML proficiency benchmark. Third, and more fundamentally, EGRAs are not designed for this purpose. As one leading EGRA advocate colorfully put it while commenting on an early draft of this paper, "It is a good general principle not to try to saw with a screwdriver! In the early grades, too much depends on the inherent opacity of the orthography in which languages are rendered" for those assessments to be used for this purpose.

[26] Use of Level 3 is consistent with the revised UIS proposals for consideration by GAML meeting in August 2019 (but deviates from the 2018 provisional GAML recommendation of mapping to Level 2). A comparison with PISA results (reported below) supports this decision to adopt Level 3. It also supports the decision to use Level 4 as the MPL for PASEC, as recommended by the revised proposal for GAML (and thus to deviate from the 2018 GAML recommendation).

data from that round of PASEC to provide estimates for 13 countries. We defined reading proficiency as reaching Level 4, in line with the proposals to GAML.

c) *SACMEQ (Southern and Eastern Africa):* SACMEQ has carried out several rounds of data collection for Eastern and Southern African countries. The latest round of the SACMEQ assessment (SACMEQ IV) was carried out in 2013. Due to concerns about the quality of the data for this round, we do not use this data to establish levels. We do, however, use it for estimating changes in proficiency over time, adopting Level 5 as the proficiency threshold, in line with the proposals to GAML.[27]

3) *Global Assessments – Non-Reading (Science or Math)*

a) TIMSS (global): Some countries that have not participated in PIRLS or the regional assessments have participated in the TIMSS international math and science assessment of 4th-graders. For these countries, we use the science scores as a proxy for reading scores, counting children as proficient if they exceeded the Low International Benchmark. We have two rationales for using this proxy. The first is conceptual: the ability to answer science questions requires reading proficiency, since most science questions are word problems. The second is empirical: across the countries for which we have data for both subject areas, proficiency on science is highly correlated with reading proficiency. Within the PISA assessment, the science-reading correlation is 0.98, and for countries that have participated in both TIMSS and PIRLS, the correlation between the two is 0.965 (Table 12). There were 15 countries in which TIMSS was used to construct the learning poverty indicator (see country table in Annex 7 for more details). The preferred subject was science, and the minimum proficiency set at the "low" benchmark. In 14 countries we were able to use TIMSS Science; TIMMS Math was used only for Jordan, which had no TIMSS Science score.

4) *National assessments (interim SDG reporting):* For some larger-population countries that have not participated in any of the assessments listed above, we used interim data from National Assessments, as per the UIS SDG 4.1.1 reporting protocol.[28] As with the regional assessments, this required deciding on the appropriate level of proficiency to map to the global benchmark. We made these judgments by: (1) drawing on the proficiency-level descriptors within each national assessment to select the level that most closely matched the global description of reading proficiency; and then (2) consulting UIS and World Bank country experts and triangulating with other data sources (such as PISA or citizen-led assessments) to determine whether any adjustment was necessary. These decisions are highly consequential, given that the 12 countries for which we rely on national assessments are responsible for 57% of the children in our database.[29]

---

[27] If instead we omit the SACMEQ data from the spells database, the estimated rate of improvement for the 2000-2017 period declines somewhat. As a result, the simulations lead to even more modest expected improvements in learning poverty by 2030—reinforcing the idea that the core findings we present in this paper may, if anything, understate the current and future dimensions of the learning poverty challenge.

[28] UIS (2019).

[29] These countries are Afghanistan, Bangladesh, Cambodia, China, Ethiopia, India, Kyrgyz Republic, Malaysia, Pakistan, Sri Lanka, Uganda, and Vietnam. Note, however, that even if we exclude these 12 countries, the estimated global learning poverty rate is quite similar to the full global estimate, at 54% for low- and middle-income countries and 46% when high-income countries are also included.

The resulting consolidated global dataset includes estimates of reading proficiency for 116 countries. Of these countries, 100 have data from the period starting in 2011;[30] 62 of these are low- or middle-income country clients of the World Bank. The total population represented by the estimates for all years in the dataset is 515 million children between age 10 and 14, which corresponds to 84% of the global cohort. Once we restrict the dataset to data from 2011 onward, the corresponding global figures are 496 million and 81% of the global cohort, and the figures for low- and middle-income countries are 437 million and 80%. (See Annex 3, Table 15 for more details.)

Measuring out-of-school children

As discussed, one additional input is data on out-of-school children. Our preferred measure of school participation is adjusted net enrollment rates for primary school, computed using administrative records such as National Educational Management Information Systems (EMIS). The primary data source for this information was UIS, with validation by World Bank and UIS country education specialists.[31] Given that this data is not always available for every year, we filled in missing data with information from the closest available year. If data are available for two years equally close to the missing year, the older value is systematically used.[32] If, despite this procedure, we still had missing values for adjusted net enrollment rates, we followed a hierarchy of alternate enrollment rates for our measure of school participation; Annex 4 describes this hierarchy and also provides a summary of decisions made for specific countries. In using the adjusted net enrollment rate, we account for some children of primary school age who might enter primary school early and advance to secondary school before they reach the official upper age limit of primary education.

[30] See the next section for a discussion of why 2011 is used as a cutoff.

[31] In some cases, the evidence was compelling enough to require alternative sources of data. One such case was Afghanistan, for which the official UIS data show an adjusted net enrollment rate of 28% in 1993 and a gross enrollment rate of over 100% for 2017, while the Afghanistan Living Conditions Survey (2016/2017) reports that enrollment is 50.4%. This last estimate which was deemed closer to the reality by country teams and was therefore used in our estimates.

[32] The earlier year tends to be closer to the latest population census round (that is, 2010) and is therefore less sensitive to any errors in population projections. Moreover, for countries that are expanding their primary school systems (which are more likely to be low-income countries with higher levels of learning poverty), the older number would give a more conservative estimate.

Population

We use population figures for the age 10-14 cohort, which matches the expected population of children at the end of primary. Given that population numbers are needed in these calculations only as weights for the regional and global aggregates, the use of 5 age cohorts is preferred to a single one.[33] We advise great caution on the extrapolation of the learning poverty rates into the absolute number of children living in learning poverty, since such extrapolation will be extremely sensitive to the choice of population definition. To ensure cross-country comparability, we use an international population source; see Annex 5 for a detailed description of the population data. In the next section we discuss the robustness of our global learning poverty estimates to different population definitions.

Data coverage, access, and quality

The learning poverty indicator is based on data covering four-fifths of children at the end of primary. In other words, 80 percent of children in low- and middle-income countries live in a country with at least one learning assessment at the end of primary, carried out in the past 9 years, that is of sufficient quality to be used for SDG monitoring.[34] This extensive coverage became possible only in recent years, with the progress in measuring learning in countries and the GAML's efforts to establish comparability, which has made possible the construction of a global indicator based on harmonized proficiency levels. It is worth noting that the 80% population coverage for learning poverty is much higher than the coverage of the global monetary poverty indicator when it was first launched in the 1990s.[35]

Despite this progress, major gaps remain in data coverage, concentrated in the countries where the learning crisis is most acute**.** While LAC and EAP have almost 90 percent coverage, less than half of children in Sub-Saharan Africa live in a country with a National Large-Scale Learning Assessment (NLSA) or a cross-national learning assessment of adequate quality to be used for this purpose (Table 2). Differences in coverage by income level are also striking. Virtually all children in high-income countries are in educational systems with such monitoring at 4th grade; the corresponding figure for low-income countries is less than 40 percent, and often those assessments take places as late as 6th grade. Recency of the data also differs: in high-income countries, 70% of these assessments took place in the last four years, but in low- and middle-income countries, the figure is only 35%. Data comparability, both within countries over time and across countries, also still poses a significant challenge. These gaps underscore the urgent need for action on improving data availability and quality (see Annex 10).

---

[33] We also replicate results using different population data, such as a single age cohort (age 10) or all children of primary school age (as defined by the legislation of each country); this leads to no qualitative change in the results.

[34] Quality is assessed in this context in terms of design, implementation, comparability, frequency, timeliness, documentation, and data access. Note that if we count the number of countries with adequate learning assessments for the learning poverty indicator, rather than using this population-weighted figure, then coverage is considerably lower.

[35] This level of data scarcity is not unprecedented. In 1981, when global monetary poverty was first reported at a global scale, only 55% of the global population was covered by household surveys, with several regions not even reaching the 40% reporting benchmark. The coverage has increased substantially since 2000, thanks to the impetus generated by the MDG process, and it now remains close to 90%. For more information see http://iresearch.worldbank.org/PovcalNet/home.aspx.

Table 2 Learning Poverty indicator: Population and country coverage by country groups

	Country Group	All Countries				Low- and Middle-Income Countries*			
		N countries with data	N countries total	Population with data (millions)	Population Coverage (%)	N countries with data	N countries total	Population with data (millions)	Population Coverage (%)
	Overall	100	217	496	81.1	62	144	437	79.7
Region	East Asia and Pacific	13	37	129	86.6	6	23	119	86.9
Region	Europe and Central Asia	33	58	42	84.0	12	23	20	74.0
Region	Latin American and Caribbean	16	42	47	86.8	16	30	47	88.4
Region	Middle East and North Africa	14	21	27	71.4	6	12	22	68.8
Region	North America	2	3	23	100.0	N/A	N/A	N/A	N/A
Region	South Asia	5	8	171	98.1	5	8	171	98.1
Region	Sub-Saharan Africa	17	48	57	46.1	17	48	57	46.1
Income Level	High income	44	79	63	97.7	6	10	4	99.7
Income Level	Upper middle income	27	60	162	91.8	27	58	162	92.2
Income Level	Lower middle income	16	47	219	75.8	16	46	219	76.0
Income Level	Low income	13	31	51	63.3	13	30	51	64.8
Lending	Part 1	38	73	60	93.0	N/A	N/A	N/A	N/A
Lending	IBRD	39	68	335	91.4	39	68	335	91.4
Lending	IDA / Blend	23	76	101	56.0	23	76	101	56.0

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Notes: Data includes only assessments since 2011 (See Table 17 in Annex 3 for an expanded range); Population coverage considering share of population ages 10-14 years old. Lending Categories: Part 1 countries do not borrow from the World Bank Group; International Bank for Reconstruction and Development (IBRD); International Development Association (IDA); and IDA-eligible based on per capita income levels and are also creditworthy for some IBRD borrowing (Blend). (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

Relationship to other global learning databases

This initiative is not the first effort to build a global database that harmonizes learning data from different assessments. We have already noted above the UIS-led effort that is underway to develop harmonized learning indicators as part of the SDG monitoring process; that effort, on which this paper builds, is based on mapping proficiency levels from different assessments against each other. The current efforts to harmonize learning indicators for SDG monitoring include item-based linking of assessments. Two important global efforts precede the initiative described in this paper: the World Bank's Harmonized Learning Outcomes (HLO) effort, which published harmonized data from many learning assessments at the primary and secondary level (Patrinos and Angrist 2018); and the UIS global database, which relied on "doubloon countries'' to align scores across different assessments (UIS, 2017c).

This initiative contributes to those previous efforts in the following ways:

- First, it creates an actionable indicator based on internationally agreed standards for the learning that should be taking place in primary school. By drawing on the GAML-based proficiency estimates, it makes use of a detailed descriptor-based equating of proficiency levels across assessments. By measuring proficiency—specifically reading proficiency in late primary school—it provides an actionable indicator that we hope countries can change relatively quickly with focused effort. It is therefore based on whatever reliable and relatively recent information is available for this purpose. This approach to linking learning data differs from and complements the numerical linking approach of the HLO-ratio (Patrinos and Angrist, 2018) and the Altinok "doubloons" (UIS, 2017c), which are based on the creation of "exchange rates" across assessments using countries that participate in multiple assessment programs in a given subject, schooling level, and testing round. The HLO, for example, measures the overall learning performance of each system, so it casts a wider net of age ranges and subjects and years.
- Second, it allows us to produce estimates now, rather than waiting for the outcomes of psychometric linked learning assessments using common items which reflect the Global Proficiency Framework. That process will produce valuable insights, but good provisional estimates of the baseline and rates of change are necessary to guide action now. We are already in 2021—more than one-third of the way through the SDG period of 2015-30—and the 10-year-old children of 2030 were born last year.
- Third, it expands the data available for these comparisons. By drawing on knowledge of World Bank and UIS staff working in countries with data gaps, as well as on the WB and UIS teams' understanding of the national assessments in these countries, it can incorporate new data from a number of assessments that are not yet included in the other databases.

Relationship of Learning Poverty with the SDG 4 monitoring framework

The learning poverty measure is also well aligned the SDG 4 monitoring framework. In particular with the SDG 4.1.1 on learning and the Out-of-School Children rate indicator in the SDG 4 thematic framework (4.1.4). In which,

- Indicator 4.1.1 is stated as: Proportion of children and young people (a) in grades 2/3; (b) at the end of primary; and (c) at the end of lower secondary achieving at least a minimum proficiency level in (i) reading and (ii) mathematics, by sex
- Indicator 4.1.4 is stated as: Out-of-school rate (1 year before primary, primary education, lower secondary education, upper secondary education)

The minimum proficiency level used by Indicator 4.1.1 b is cut-off value used to identify the student population in Learning Deprivation (LD), while 4.1.4 corresponds to the School Deprivation measure in equation 1.[36,37]

[36] In the case of countries that lack a 4.1.1b value, we have used the 4.1.1a reported value as long as the grade covered is either Grade 4 or 5. One example is the use of PIRLS, which many high-income countries use to report 4.1.1a; given that it covers 4th grade (age 10), we have considered it a valid indicator for the learning deprivation measure.

[37] The SDG 4.1.4 indicator is defined as the complement to the *total net enrollment rate [100-(Total number of students of the official age group for a given level of education who are enrolled in any level of education, expressed as a percentage of the corresponding population)]*. In the case of the learning poverty measure, given the focus on the primary school age children, the

4. Results: Levels and patterns of learning poverty

This section presents the learning poverty estimates for low- and middle-income countries and explores the robustness of those estimates.

Where we are now: Half of children in low- and middle-income countries are learning-poor

The headline number that emerges from this analysis is that 52.7% of all children in low- and middle-income countries are not able to read proficiently by age 10—or even at age 12, when many of them are tested. Based on the experience of rich countries, it should be possible to reduce this learning poverty rate to the single digits—say, 5% to 8%—just as absolute poverty is near zero in those countries. In the case of monetary poverty, the global rate had already been reduced to 11% before the pandemic, well on the way to a global target of poverty elimination by 2030.[38] Yet in the education sphere, one out of every two children in the developing world is not learning to read by late primary-school age (again, even before the pandemic). And the rate is much higher in some regions: in Sub-Saharan Africa, learning poverty is 87%, or more than six times as high as the 13% rate found in Europe and Central Asia's low- and middle-income countries. There is considerable variation within regions too, with learning poverty in regions such as ECA and EAP ranging from 2% in the country with lowest learning poverty to over 50% in countries with the highest rates, or a ratio over 30.

The global average for low- and middle-income countries is held down by the inclusion of the upper-middle-income countries, which average 29% learning poverty. But in lower-middle-income countries, 55% of children cannot read proficiently, and in low-income countries, the rate is over three-quarters at 90%. (See Table 3 for more details.)

Similar patterns are found in the data by World Bank lending status. Even in the IBRD countries in the database, about 40% of children are not reading proficiently by late primary. And in the other groups, a substantial majority of children do not acquire minimum proficiency; for example, the learning poverty rate is 80% for IDA/Blend countries.

---

preferred indicator is the complement to the *adjusted net enrollment rate [100-(Total number of students of the official primary school age group who are enrolled at primary or secondary education, expressed as a percentage of the corresponding population)]*. For a detailed rank of the preferred data sources of the schooling deprivation used in the LP measure please see Annex 3.

[38] At the baseline of the MDGs, in 1990, the global monetary poverty rate (excluding other High-Income Countries), using the international dollar per day, was 43%. Under the MDGs, the world agreed on the target to halve this number by 2015. At the same time, the global number for 1990 was 36% and it reached 10% by 2015. The last time global monetary poverty was higher than 50% was in 1981; in that year, the first year in which it was reported, the rate was estimated at 52% (with a survey coverage of 55% of the World population). For more details, please visit the Povcalnet website at http://iresearch.worldbank.org/PovcalNet/povDuplicateWB.aspx.

Table 3 Share of children who are learning-poor by late primary, by country group

	Country Group	All Countries: Learning Poverty (%)	All Countries: S.E. L.P. (%)	All Countries: Minimum L.P. (%)	All Countries: Maximum L.P. (%)	Low- and Middle-Income Countries*: Learning Poverty (%)	Low- and Middle-Income Countries*: S.E. L.P. (%)	Low- and Middle-Income Countries*: Minimum L.P. (%)	Low- and Middle-Income Countries*: Maximum L.P. (%)
	Overall	48.0	0.4	1.6	98.7	52.7	0.4	1.7	98.7
Region	East Asia and Pacific	19.8	0.8	1.7	51.1	21.2	0.9	1.7	51.1
Region	Europe and Central Asia	8.8	0.2	1.6	64.5	13.3	0.4	2.2	64.5
Region	Latin American and Caribbean	50.8	0.8	20.7	80.7	50.8	0.9	20.7	80.7
Region	Middle East and North Africa	58.7	0.6	11.7	94.7	63.3	0.9	35.7	94.7
Region	North America	7.6	0.5	4.3	7.9	N/A	N/A	N/A	N/A
Region	South Asia	58.2	0.9	14.8	93.4	58.2	0.9	14.8	93.4
Region	Sub-Saharan Africa	86.7	0.3	48.3	98.7	86.7	0.3	48.3	98.7
Income Level	High income	8.7	0.2	1.6	66.6	23.9	0.5	4.0	66.6
Income Level	Upper middle income	29.0	0.5	2.2	80.7	29.0	0.7	2.2	80.7
Income Level	Lower middle income	55.1	0.7	1.7	85.1	55.1	0.7	1.7	85.1
Income Level	Low income	89.8	0.3	78.2	98.7	89.8	0.3	78.2	98.7
Lending	Part 1	7.7	0.2	1.6	51.0	N/A	N/A	N/A	N/A
Lending	IBRD	40.0	0.6	1.7	80.7	40.0	0.6	1.7	80.7
Lending	IDA / Blend	79.5	0.3	51.1	98.7	79.5	0.3	51.1	98.7

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: For specific country numbers please see Table 20, in Annex 6; Standard errors calculated through bootstrapping, see Annex 9 for details.
(*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients. See Table 2 for the number of countries per classification.

Robustness and Validation

As discussed earlier, these results build on several assumptions and methodological choices. In this section we check the robustness of the results to variations in some of those assumptions and triangulate the results with other indicators to confirm the external validity of the measure.

Choice of reporting window

To build our global aggregate, we have to establish a valid reporting window. As previously discussed, our preferred window is a 9-year window centered on 2015. It is reassuring to see that the choice of reporting window does not seem to change the global estimates drastically, at least with the currently available data. Table 4 shows the sensitivity of the results to this choice, using the latest available data and reporting windows of different lengths (9, 7, and 5 years). Although the choice of reporting window does affect the number of countries, and to a greater extent the population coverage, it does not change the learning poverty very much. However, narrower reporting windows tend to lower the global learning poverty estimate, reflecting a composition bias: countries with more recent data have lower learning poverty rates.

Table 4 Sensitivity of results with respect to choice of reporting window

	All Countries					Low- and Middle-Income Countries*				
Window	Learning Poverty (%)	S.E. L.P. (%)	Population Coverage (%)	N countries	Avg. Year	Learning Poverty (%)	S.E. L.P. (%)	Population Coverage (%)	N countries	Avg. Year
Latest	49.0	0.3	84.2	116	2015	53.8	0.3	82.9	74	2014
9 years	48.0	0.4	81.1	100	2015	52.7	0.4	79.7	62	2015
7 years	47.0	0.4	72.6	90	2015	51.7	0.5	70.3	52	2015
5 years	43.5	0.4	56.4	60	2016	47.9	0.5	52.1	22	2016

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Standard errors calculated through bootstrapping, see Annex 9 for details. (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

Choice of population of reference

Another important choice is which population of reference to use. In this subsection we present the robustness of our global estimates to different choices of our population numbers. We show that our results are qualitatively similar in terms of the global learning poverty rate, while they can significantly differ in terms of the total number of learning-poor. Table 5 shows the both the global rate and the total number of learning-poor, using four population definitions, namely: (i) our preferred measure, which combines the five age cohorts from 10 to 14 years of age; (ii) a single age cohort (10 years of age); (iii) all children of primary school age, defined based on the *de jure* ages at which children are supposed to start and finish primary school (5-16); (iv) all children from age 9 to the official *de jure* age at which children are supposed to finish primary; and, (v) all children enrolled in primary school.

Results for the global rate are quite robust to these changes in population cohorts: learning poverty ranges from 53% to 55% for low- and middle-income countries, and from 48% to 51% for the whole world. By contrast, the absolute number of learning poor unsurprisingly differs greatly based on the population definition, ranging from 60M to 720M children globally (Table 5). It is important to keep in mind that there is no perfect choice of population of reference given the data limitations, and those differences may generate noise around the credibility of the estimates.

This exercise has implications for how learning poverty figures are communicated. First, the learning poverty rate can be used with confidence that it is not sensitive to the primary-age population of reference used. Second, given that absolute numbers are often seen as more powerful as communication tools, those citing the numbers of learning-poor will need to be careful to specify clearly the population of interest, and to caveat that such indicator is just for illustration purposes. A final important point from Table 5 is that no matter what definition is used, between 98% and 99% of all learning-poor children in the world live in low- and middle-income countries.

Table 5 Sensitivity of results to choice of population of reference

	All Countries					Low- and Middle-Income Countries*				
Population Definition	Learning Poverty (%)	S.E. L.P. (%)	Population (millions)	Population Coverage (%)	Learning Poor (millions)	Learning Poverty (%)	S.E. L.P. (%)	Population (millions)	Population Coverage (%)	Learning Poor (millions)
Ages 10-14	48.0	0.3	612	81.1	294	52.7	0.4	548	79.7	289
Age 10	48.3	0.4	125	80.7	60	53.0	0.4	112	79.3	59
Ages 5-16	48.1	0.3	1498	80.8	720	52.8	0.3	1343	79.4	708
Ages 9+	50.5	0.3	449	77.9	227	55.4	0.3	402	76.0	223
Primary	48.0	0.3	804	80.8	386	52.9	0.3	718	79.2	380

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Standard errors calculated through bootstrapping, see Annex 9 for details. (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

Consistency with other learning data

As a check on the estimated learning poverty rates, we can compare them with estimates of reading ability for other age groups from other sources. Several datasets include data on reading abilities of adolescents, adults, and younger children; in this subsection, we use them to explore whether reading proficiency levels are indeed as low as suggested by our learning poverty measure.

Here, we check the end-primary reading proficiency results against four other sources:

Skills of 15-year-olds (PISA): The first check uses data from PISA on the literacy of 15-year-olds. This comparison may be especially useful as an independent check on the relevance and consistency of the learning poverty data, for two reasons: first, PISA assesses competencies of students just 3 to 4 years past the end of primary, and second, it is available for many middle-income countries and (more recently, with the PISA-D program) for a few low-income countries.[39] On the PISA Reading test, proficiency levels range from Below Level 1b to Level 6. Level 2 is typically considered the minimum proficiency level for lower-secondary, and using it is equivalent to the approach we have used at the end of primary.[40] In other words, the share of students scoring *below* Level 2 is the analogue, for 15-year-olds, to the end-of-primary learning deprivation measure used to produce our learning poverty measure. Given that the average age for the end-primary assessments in our dataset ranges from 10 to 13, the assessments measure students only a few years apart, and therefore the end of primary and lower secondary learning deprivation measures should correspond closely. In this exercise, we have used PISA assessments around 3 to 5 years after the year of our learning poverty measure, in an attempt to track similar age cohorts. These measures are equivalent to the complement of the SDG 4.1.1b and SDG 4.1.1c indicators.

[39] Although it is theoretically possible that a country could perform poorly on 4th-grade reading proficiency but well on reading of 8th-graders or young adults—for example, if its secondary schools are managed better than its primary schools—this is not likely in practice. Systems that are high-performing at the lower grades tend to do well at higher grades too. While no single international literacy assessment assesses children in both 4th and 8th grade, the TIMSS Science assessment does, and the cross-country correlation for 2015 between 4th-grade and 8th-grade scores is quite high, with a correlation coefficient of 0.96.
[40] Level 2 is described as the "baseline level . . . at which readers begin to demonstrate the competencies that will enable them to participate effectively and productively in life as continuing students, workers and citizens" (OECD PISA for Development brief). In its country reports, the OECD signals the importance of this baseline by highlighting the share of a country's student population that is below Level 2. This is also the level that the GAML process recommended for the SDG 4.1.1c monitoring.

Because virtually all of the countries participating in PISA are either high- or middle-income, this test can be applied only to those countries and not the low-income countries in our dataset (notable exceptions are the PISA-D participant countries). Figure 1 compares the PISA Level 2 or PISA/Secondary learning deprivation measures against our end-primary learning deprivation and our learning poverty measures for 60 countries (panels a and b), 25 of which are among our middle- and low-income countries (panel c and d). The results are instructive:

- As expected, there is a strong relationship between end-of-primary and lower-secondary learning deprivations in reading. The Pearson's correlation of the end-of-primary learning deprivation (used in our learning poverty measure) with the PISA reading learning deprivation is 0.87, meaning that we can predict ¾ of the variation in a country's performance in lower-secondary just by knowing its performance in end-primary proficiency. And this correlation appears to reflect actual skills of the cohort: the correlation weakens if, as a placebo test, we compare instead PISA scores from rounds that took place much earlier than our learning poverty was collected. For example, if we use PISA measures collected from 6 to 11 years prior to the year of our learning poverty estimates, or even more than 11 years before the assessment of the end of primary, we find correlations of, respectively, 0.69 and even 0.14 (see Table 18 in Annex 6). If we focus only on middle- and low-income countries, the latter correlation is also quite low.
- But despite the strong relationship between primary and secondary learning deprivation rates, the primary rates are much lower than secondary rates in most cases. Visually, this means that most of the points in the graph are well above the 45-degree line in Figures 1a and 1c. For the average country in the sample, the gap is 10 percentage points, and for middle- and low-income countries it is 16 percentage points (see Table 18 in Annex 6).
- This difference holds if we use either our measure of learning poverty or just the share of end-of-primary learning-deprived students (as shown by the different panels of Figure 1 and by Table 19 in the Annex).

Figure 1 Rates of learning deprivation in reading: comparison of end-primary and lower-secondary measures (15-year-olds, PISA)

a. All countries, using learning deprivation as the end-primary measure (n=60)

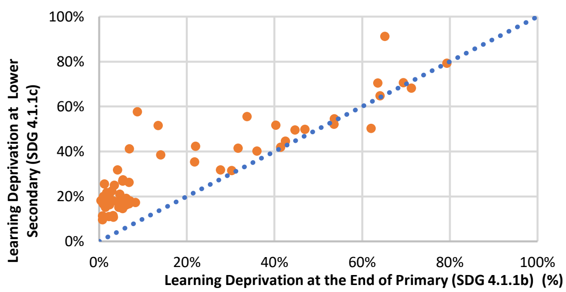

b. All countries, using Learning Poverty as the end-primary measure (n=60)

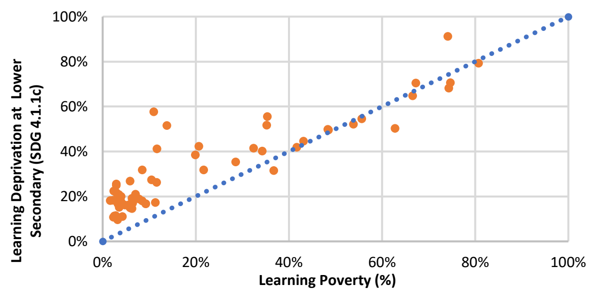

c. Only low- and middle-income countries, using learning deprivation measure as the end-primary measure (n=25)

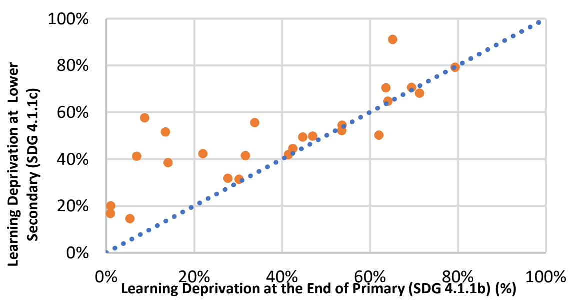

d. Only low- and middle-income countries, using Learning Poverty as the end-primary measure (n=25)

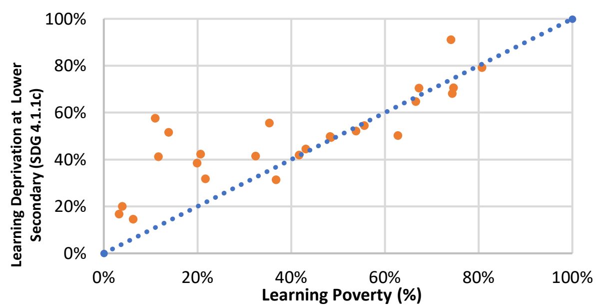

Source: Learning deprivation at lower-secondary (SDG 4.1.1c) is measured as the share of students scoring below Level 2 on the PISA reading test (OECD). PISA assessments used are those administered between 3 to 5 years after the learning assessment used for the Learning Poverty measure, as used in panels b and d (or the end-of-primary learning deprivation measure, panels a and c).

The fact that lower-secondary learning deprivation (as measured by PISA or the SDG 4.1.1c) is substantially higher than our end-primary measure of learning deprivation (SDG 4.1.1b) in these high- and middle-income countries is surprising, given selection into secondary schooling. In most middle-income countries, enrollment of 15-year-olds is far from universal. The students who remain in school long enough to be tested by PISA tend to be more advantaged (which usually means higher-scoring) than those who drop out. This selection effect should increase relative test scores, leading the lower-secondary learning deprivation rates to be *lower* than those for primary.

These results suggest that our estimated end-primary learning deprivation rates do not overstate the problem, at least for middle-income countries. Indeed, the comparisons with lower-secondary learning deprivation rates (PISA below Level 2) suggest that our learning poverty rates may *underestimate* the true values for many countries with low learning deprivation rates, by some 20 percentage points.

Adult skills (PIAAC): Another check of consistency is to compare with the literacy skills of young adults. Other factors such as on-the-job learning will also affect adults' literacy, but especially for the young cohort, the quality of learning during the school years is likely to be a major determinant. On the OECD's PIAAC survey of adult skills, Level 2 can be considered as the minimum proficiency. In Japan, less than 3% of the young-adult population (age 16-24) is non-proficient by this measure, and in the rest of PIAAC's top 10 countries in young-adult literacy, the number is less than 10%. By contrast, in Chile 39% of young adults are non-proficient on PIAAC; in Turkey, the figure is 37%; and in the Russian Federation, it is 14%. These young-adult nonproficiency rates from PIAAC are highly correlated with our learning poverty measure, with a correlation coefficient of 0.76. However, learning poverty is substantially lower, at 19% for Chile, 25% for Turkey, and just 3% for the Russian Federation. Across countries, learning poverty is a median of 6 percentage points lower than PIAAC nonproficiency. Like the PISA results, these results too suggest that, if anything, our estimates of primary age nonproficiency are conservative.

UIS estimates of "Children not learning": In addition to these measures of adolescent and adult skills, there is another recent estimate of the extent of the problem that is more directly comparable. The UIS estimated in 2017 that "more than 617 million children and adolescents are not achieving minimum proficiency levels (MPLs) in reading and mathematics" (UIS 2017a). Underlying those totals is UIS's estimate that 56% of primary school-age children do not attain the minimum proficiency in literacy.[41] This estimate was based on a more limited coverage of end-of-primary proficiency in low- and middle-income countries than the one we use here, as national learning assessments were not included, and it used a different harmonization methodology (very similar to that of the HLO) that appears to include inferring primary school learning from lower secondary school assessments in some cases.[42] Also, unlike our global learning poverty estimates, UIS's earlier global figure includes high-income countries; restricting the sample to low- and middle-income countries would further increase the Children Not Learning share.[43] Thus this comparison too suggests that the global learning poverty estimate of 53% is on the conservative side as an indicator of the scale of the learning challenge.

Early-grade learning (3rd and 4th Grade, LLECE and PASEC): Another check is the extent to which our end-of-primary learning measure is aligned with early primary grade results for each country. This is important as we would like to know if educational systems that are performing well at early grades are more likely to be performing well at the end of primary, and vice versa. This is analogous to the triangulation with lower secondary. Although only two regional assessment programs, LLECE and PASEC, enable us to make this comparison, this exercise still gives us a greater representation of low- and middle-income countries than the PISA and PIAAC comparisons do.

Given the choice of grade assessed and periodicity of these two assessment programs, we cannot compare scores for a single student cohort, so the results should be interpreted as the association of the quality of learning delivered in early grades and at the end of primary. We make the comparisons in terms of both the average reading score and the share of learning-deprived students. In the case of Latin America (Figure 2 panels a and c), a cross-country comparison of early-grade and end-of-primary

[41] Table 1 of UIS (2017a)

[42] UIS (2017b) reports that "to estimate meaningful regional aggregates, the proportion of students achieving minimum proficiency level at the end of primary education by subject were assumed to be equal to the proportion of students achieving the minimum proficiency level at the end of lower secondary education. This treatment was applied to countries with large regional weight such as China, Egypt and India" (p. 15).

[43] UIS (2017a) estimates the share of Children Not Learning in lower-middle-income and low-income countries at 76% and 91%, respectively (see Table 2). These estimates are substantially higher than our learning poverty estimates of 54% and 78% for the same country groups.

performance shows a very strong relationship, with approximately 90% of the variance being explained and a very low dispersion of the countries around the trend line. For Sub-Saharan Africa, the results are substantially different: while we find a positive relationship when we compare mean scores, the relationship is not very strong (Figure 2d), and we find virtually no relationship when we use the learning deprivation measure for end-primary (Figure 2b).

One channel that might explain the weaker correlation in Sub-Saharan Africa than in Latin America is the transition in language of instruction, which happens around age 10 in many African countries and which likely weakens students' literacy as tested in the second language. Using data from PASEC, we find results consistent with this hypothesis. First, the early-grade relationship is significantly stronger for math. Second, if we remove countries such as Burundi, in which there is a documented transition from Kirundi to French, the strength of this relationship would jump from an RHO of 0.03 to -0.53 in the case of learning deprivation (see Table 19 in the Annex). This suggests that in systems where there are transitions in language of instruction during primary school, interpretation of the learning poverty results should take this into account.

Figure 2 Learning Deprivation and Average scores in reading: Early Grade vs End of Primary

a. Latin America (LLECE) learning deprivation at the end of primary vs Early Grade average scores (n=31)

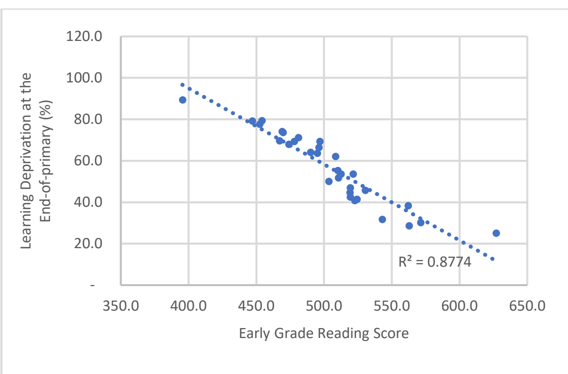

b. Sub-Saharan Africa (PASEC) learning deprivation at the end of primary vs Early Grade average scores (n=10)

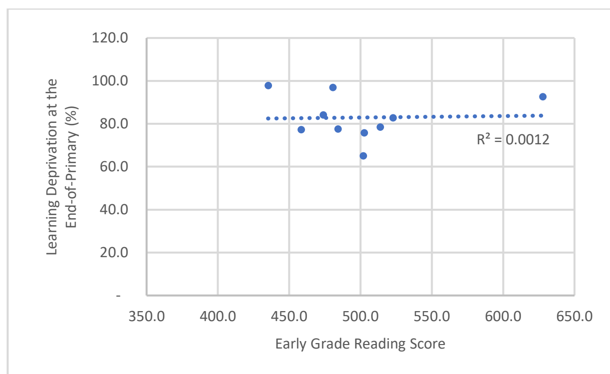

c. Latin America (LLECE) End-of-Primary vs Early Grade average scores (n=31)

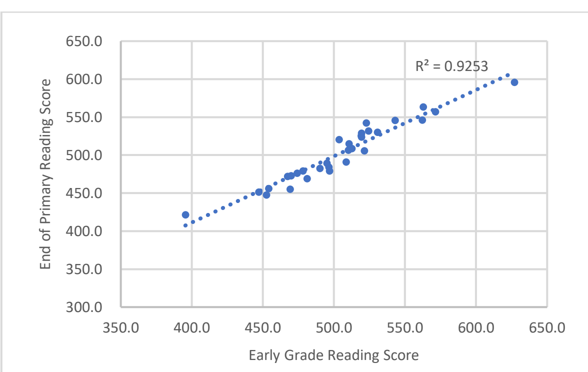

d. Sub-Saharan Africa (PASEC) End-of-Primary vs Early Grade average scores (n=10)

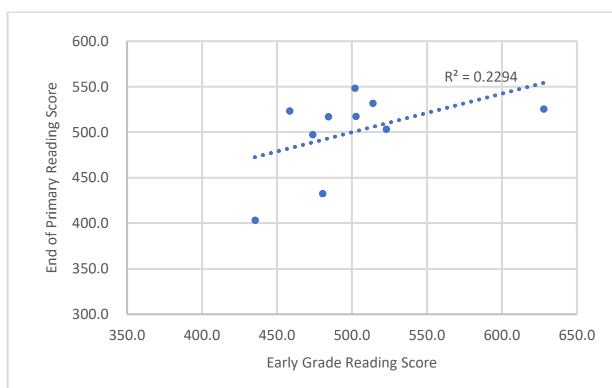

Source: Author's calculation. Note: for Early Grade we use the national average score in reading (grade 2 for PASEC and grade 3 for LLECE); for the end of primary we use both the learning deprivation share and the end-of-primary average reading score (grade 6 for both PASEC and LLECE). PASEC data from 2014 (n=10) and LLECE from 2006 (n=16) and 2013 (n=15).

Conclusion: These initial comparisons show that the learning poverty estimates for low- and middle-income countries correlate strongly with other estimates of reading ability, including those for other age groups (with some exceptions that may be driven by changes in language of instruction). The comparisons also suggest that, if anything, the learning poverty estimates are somewhat conservative as indicators of the scale of the learning crisis. For example, consider the implications of the PISA comparison, which shows that the PISA learning deprivation rates[44] are considerably higher than our primary learning deprivation and even learning poverty rates. This means that children who just clear our minimum proficiency threshold by the end of primary are not on track to meet PISA Level 2 by age 15—which in turn means that they will not even "begin to demonstrate the [reading] competencies that will enable them to participate effectively and productively in life as continuing students, workers and citizens." In other words, the actual level of learning poverty or share of learning-deprived students in low- and middle-income countries could be even higher than the 53% reported above. But whether the learning poverty rate is 53% or even higher, the high rate has major implications both for the prospects of reaching universal reading proficiency and for the way forward.

In what follows, we will build on these findings to present some evidence on the heterogeneity of the proposed learning poverty measure and will return to some of the points presented above.

Understanding the heterogeneity of Learning Poverty

In this subsection we discuss the heterogeneity of learning poverty across four important dimensions. First, we unpack the contribution of both schooling and learning to learning poverty, globally and for different groups of countries. Second, we introduce a learning gap measure. Since our learning indicator is the share of students below minimum proficiency, by definition it is not sensitive to changes in the distribution that take place *below* the threshold; below, we use the gap measure to show the importance of tracking those changes too. Third, we present learning poverty measures disaggregated by gender, which is a key dimension of inequality in many countries that may require special interventions. Fourth, we present evidence on spatial differences in learning poverty within a country, using Brazilian data for the illustration.

Unbundling Learning Poverty: The relative contribution of quantity and quality of schooling

learning poverty is multidimensional, as is it capture both learning and schooling by combining two separate yet conceptually related indicators. From a diagnostic and policy perspective, it is useful to unbundle this indicator and try to understand how much learning deprivation (LD) and schooling deprivation (SD) contribute to the overall level of learning poverty. The decomposition of our learning poverty measure suggests that weaknesses in learning (for enrolled children) account for 84% of the learning poverty in low- and middle-income countries (see Table 6). [45] This is not surprising, given that the out-of-school rate is about 8% in our data. However, there are important differences across country groupings. In the Sub-Saharan Africa, low-income, and IDA/Blend country groupings, we find that schooling gaps are still responsible for approximately 30% of learning poverty (see Table 6), consistent

[44] We define PISA learning deprivation using the same MPL used for SDG 4.1.1c: a PISA score of 407, or the PISA Low International proficiency level.

[45] For methodological details of the decomposition approach used see Azevedo (2020).

with the relatively high share of out-of-school children in these regions. What is more striking, however, is the relatively high contribution of the schooling component in the ECA region and high-income countries, where schooling contributes about a quarter of the learning poverty. In North America, the figure is 50%. One hypothesis is that the challenge of serving the "last mile" of children—those who live in remote locations or who face bigger barriers due to family background, migration status, or disability—becomes more important in relative terms as the quality of an educational system improves in absolute terms.

Table 6 Identifying the contributions of learning and schooling deprivation to learning poverty[46]

		All Countries					Low- and Middle-Income Countries*				
			Absolute Value		Relative Contribution			Absolute Value		Relative Contribution	
	Country Group	Learning Poverty (p.p.)	LD (p.p.)	SD (p.p.)	LD (%)	SD (%)	Learning Poverty (p.p.)	LD (p.p.)	SD (p.p.)	LD (%)	SD (%)
	Overall	48.0	40.0	8.0	83.4	16.6	52.7	44.0	8.7	83.5	16.5
Region	East Asia and Pacific	19.8	19.0	0.7	96.2	3.8	21.2	20.5	0.6	96.9	3.1
Region	Europe and Central Asia	8.8	6.6	2.2	75.1	24.9	13.3	9.6	3.7	72.1	27.9
Region	Latin American and Caribbean	50.8	47.2	3.6	92.9	7.1	50.8	47.2	3.6	92.9	7.1
Region	Middle East and North Africa	58.7	54.6	4.1	93.0	7.0	63.3	58.9	4.4	93.0	7.0
Region	North America	7.6	3.8	3.8	50.1	49.9	N/A	N/A	N/A	N/A	N/A
Region	South Asia	58.2	51.2	7.0	87.9	12.1	58.2	51.2	7.0	87.9	12.1
Region	Sub-Saharan Africa	86.7	62.3	24.3	71.9	28.1	86.7	62.3	24.3	71.9	28.1
Income Level	High income	9.1	6.7	2.4	73.6	26.4	22.3	16.6	5.7	74.5	25.5
Income Level	Upper middle income	30.3	28.5	1.8	94.0	6.0	30.0	28.1	1.9	93.8	6.2
Income Level	Lower middle income	55.8	50.0	5.8	89.6	10.4	55.9	50.1	5.8	89.6	10.4
Income Level	Low income	89.5	61.1	28.4	68.2	31.8	89.6	61.1	28.4	68.3	31.7
Lending	Part 1	8.1	5.9	2.2	73.3	26.7	N/A	N/A	N/A	N/A	N/A
Lending	IBRD	40.1	38.0	2.1	94.8	5.2	39.9	37.8	2.1	94.8	5.2
Lending	IDA / Blend	82.5	58.4	24.0	70.9	29.1	82.5	58.4	24.0	70.9	29.1

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Notes: Decomposition of learning poverty by its two dimensions: learning deprivation (LD) and schooling deprivation (SD) rates; for a methodological description, please see Azevedo (2020); (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

How important is the learning deprivation gap?

The distribution of children according to learning levels below the MPL matters. Here we introduce a complementary measure, learning deprivation gap, which is sensitive to changes in learning below the

[46] As per Equation 1, the absolute value of the schooling deprivation (OoS) contribution to learning poverty is the actual reported rate of out of school. For the learning deprivation measure, however, this is not the case, since what is reflected in the learning poverty is the learning deprivation weighted by the complement of the schooling deprivation, or out-of-school, population. Alternatively, this can be seen an adjustment to present both the LD results in terms of by the same population as the SD, namely the all children in the age-group of reference, and not just the student population. See Azevedo (2020) for a more detailed discussion on this point.

MPL.[47] While the underlying learning measure (the score on a student assessment) is a continuous variable, the learning deprivation indicator used to calculate learning poverty is the share of students Below Minimum Proficiency. As a result, the learning deprivation indicator is not sensitive to how far below this minimum proficiency level students might be. This omission can be especially relevant for countries or regions with high levels of inequality in the learning distribution among the learning-poor. In such countries, there could be substantial improvements in learning that fail to move students above the reference threshold, and thus fail to reduce learning poverty.

For this reason, we also explore a learning deprivation gap measure, defined as the population-standardized average distance to the MPL for the students below the MPL.[48] Figure 3 below illustrates that as learning deprivation increases, so does the dispersion of the learning deprivation gap across countries. Thus, two countries with similarly high learning deprivation rates can have very different learning deprivation gaps, and the policies required to get the same reduction in learning deprivation could differ considerably. For example, two PASEC countries (shown in orange in Figure 3) with over 90% of students below the minimum proficiency threshold turn out to have very different learning deprivation gaps of 13% and 33%, respectively, suggesting that the variance of proficiency *below* the threshold is quite different for these two countries. The country with the higher learning deprivation gap is likely to need greater flexibility in its learning policy packages to target learners at a wider range of learning levels, relative to the one with a smaller learning deprivation gap.

Figure 3 Learning deprivation and the learning deprivation gap by country

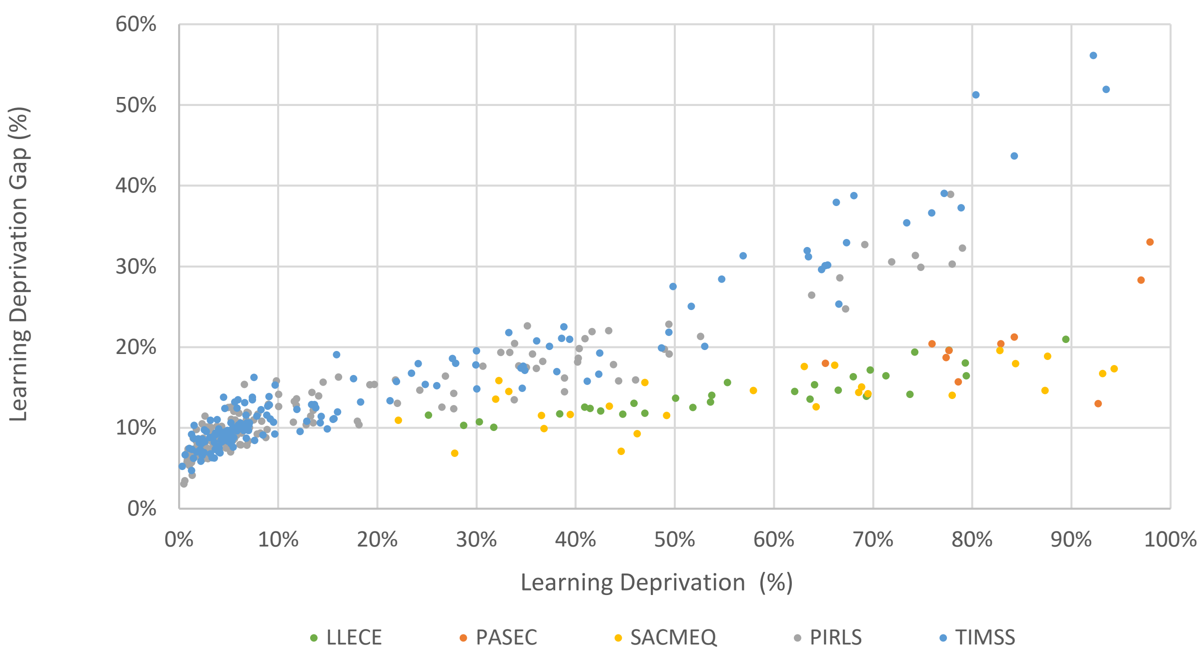

Source: Authors' calculations using the Global Learning Assessment Database; Note: each point represents one country-assessment observation (N = 399).
Note: This figure is for illustration purposes only, and comparisons of learning gaps should only be done within specific learning assessment programs.

[47] For a detailed discussion of the axiomatic properties of the learning poverty measure and an extension to a distributional sensitive class of measures, see Azevedo (2020).
[48] See Azevedo 2020 for the functional form.

How does learning poverty vary by gender?

Using all available cross-country assessments (as well as gender-disaggregated enrollment data from UIS), we have computed gender-specific learning poverty rates. Given data availability, we have only been able to compute this disaggregation for 93 countries, covering 59% of the global children and 55% of children in the developing world. Sub-Saharan Africa and East Asia are two regions with low coverage, largely because of problems in accessing the national learning assessment microdata, especially in China, Ethiopia and Uganda. In other words, our gender-disaggregated results have substantially lower coverage rates than our headline global number does (81% for all countries and 80% for low- and middle-income countries). This should be taken into consideration when interpreting these results, because gender patterns the missing countries might differ from what is shown here.

The results show that in virtually all countries for which we have data, girls have lower rates of learning poverty than boys do (Figure 4). Table 7 shows that girls are on average 4 percentage points less learning-poor than boys.[49] The difference is significantly smaller in Europe and Central Asia and in North America, and larger in the Middle East and East Asia and Pacific. The gender differential is significantly greater in middle-income countries: while in high-income and low-income countries differences are quite small, the gap reaches 5 percentage points in upper middle-income countries.

Table 7 Learning poverty by boys and girls, and country groups, for a subsample of countries

		All Countries					Low- and Middle-Income Countries*				
			Boys		Girls			Boys		Girls	
	Country Group	N countries w/ data	Avg LP (%)	SE LP (%)	Avg LP (%)	SE LP (%)	N countries w/ data	Avg LP (%)	SE LP (%)	Avg LP (%)	SE LP (%)
	Overall	93	50.1	0.5	46.3	0.4	53	57.7	0.5	53.6	0.5
Region	East Asia and Pacific	9	29.6	1.4	21.1	1.4	2	39.0	2.2	27.9	2.1
Region	Europe and Central Asia	37	10.0	0.3	8.2	0.2	13	14.4	0.5	12.7	0.4
Region	Latin American and Caribbean	19	52.7	0.9	48.6	1.0	18	53.0	1.0	48.9	1.2
Region	Middle East and North Africa	13	66.0	0.8	56.8	0.8	7	68.2	0.8	59.1	0.8
Region	North America	2	8.0	0.6	7.1	0.5	N/A	N/A	N/A	N/A	N/A
Region	South Asia	2	59.3	1.0	56.8	1.0	2	59.3	0.9	56.8	1.0
Region	Sub-Saharan Africa	11	86.9	0.5	83.8	0.5	11	86.9	0.5	83.8	0.5
Income Level	High income	45	8.4	0.3	6.6	0.2	6	25.7	0.7	22.0	0.6
Income Level	Upper middle income	28	44.5	0.6	39.4	0.7	27	44.6	0.6	39.5	0.7
Income Level	Lower middle income	13	58.3	0.8	54.3	0.8	13	58.3	0.8	54.3	0.8
Income Level	Low income	7	92.9	0.4	93.0	0.4	7	92.9	0.4	93.0	0.3
Lending	Part 1	40	7.5	0.3	5.8	0.2	N/A	N/A	N/A	N/A	N/A
Lending	IBRD	40	51.3	0.7	46.7	0.6	40	51.3	0.6	46.7	0.7
Lending	IDA / Blend	13	82.1	0.6	80.7	0.4	13	82.1	0.5	80.7	0.4

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Gender breakdowns calculated using all learning assessments since 2000 for which data was available. Estimates do not reflect all the learning assessments in the pooled dataset, because of a lack of gender-disaggregated data, and so the averages of the male and female

[49] If numeracy (rather than literacy) were used to calculate the learning poverty rate, the pattern might be expected to be different. Among adolescents, boys slightly outperform girls on PISA math scores, for example. Yet on the 4th-grade TIMSS math assessment, girls outperform boys in many countries, such as Saudi Arabia, Jordan, South Africa, and Indonesia, and in the median country there is no gender difference.

columns do not match the global averages reported earlier. Standard errors calculated through bootstrapping; see Annex 9 for details. Countries without gender-specific enrollment information and with a national enrollment higher than 98.5 are assumed to have gender parity on enrollment. (Nine countries are affected by this assumption: Austria, Canada, Costa Rica, Czech Republic, Germany, Georgia, Iran, Singapore, and Tunisia.)

Gender gaps are also higher for countries in the middle of the distribution of learning poverty. Figure 5 shows that the gender gaps are significantly higher for countries with a learning poverty rate between 30% and 70%. This finding is consistent with the patterns in other measures of horizontal inequality, such as spatial, rural/urban, and social-economic differences in learning poverty. It is also aligned with the work of Crouch, Rolleston, and Gustaffson 2020, who argue that tackling the learning crisis in high-learning-poverty countries also addresses learning inequity and inequality.

Figure 4 Learning poverty gender gap, by country

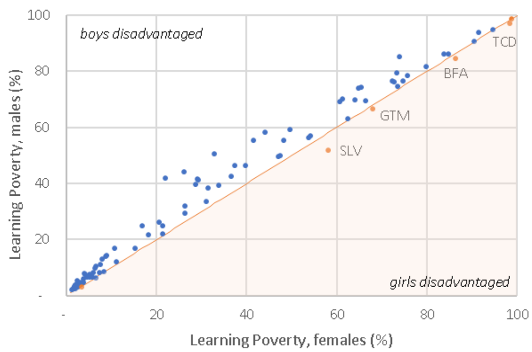

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers. Note: each point represents the latest available data for a country (N = 93)

Figure 5 Learning poverty gender gap by level of Learning Poverty

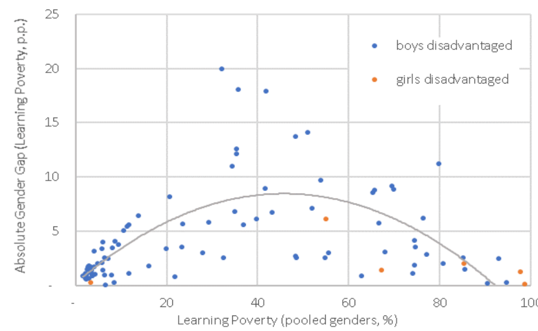

Source: Authors' calculations using the Global Learning Assessment Database; and UIS Enrollment Data. Note: each point represents the latest available data for a country (N = 93)

We can also examine the decomposition into the learning and schooling components by gender. At the global level, learning and schooling contributions to learning poverty are quite similar by gender, with 90% of the overall measure being accounted for by learning for both boys and girls, versus only 10% by schooling. Note that this is slightly higher than the 84% accounted for by learning above, likely because the countries for which gender disaggregation is not possible are countries with greater access challenges (for example, Sub-Saharan Africa). Here too, we see important differences across country groupings (Table 8). In particular, the contribution of the out-of-school component is greater for girls than for boys in Europe and Central Asia, the Middle East and Northern Africa, Sub-Saharan Africa, and North America, while the reverse is true in South Asia, East Asia, and Latin America. In virtually all cases, girls are doing systematically better than boys in terms of learning, both in absolute and relative terms, but the enrollment picture differs greatly across regions: in Sub-Saharan Africa and Middle East and Northern Africa, girls' lower enrollment remains a significant challenge both in relative and absolute terms, while in East Asia, Latin America, and South Asia the boys are the ones that tend to be left behind at primary school-age. In no other region does the access challenge remain as large as in Sub-Saharan Africa, especially for girls.

Table 8 Decomposition of learning poverty by learning and schooling, for boys and girls

		Boys, all countries					Girls, all countries				
			Absolute Value		Relative Contribution			Absolute Value		Relative Contribution	
	Country Group	Learning Poverty (p.p.)	LD (p.p.)	SD (p.p.)	LD (%)	SD (%)	Learning Poverty (p.p.)	LD (p.p.)	SD (p.p.)	LD (%)	SD (%)
	Overall	50.1	44.2	5.9	88.3	11.7	46.3	39.9	6.4	86.2	13.8
Region	East Asia and Pacific	29.6	26.0	3.6	88.0	12.0	21.1	20.3	0.8	96.2	3.8
Region	Europe and Central Asia	10.0	7.3	2.7	73.0	27.0	8.2	5.6	2.7	67.7	32.3
Region	Latin American and Caribbean	52.7	48.7	4.1	92.3	7.7	48.6	45.5	3.2	93.5	6.5
Region	Middle East and North Africa	66.0	62.6	3.5	94.7	5.3	56.8	51.7	5.1	91.0	9.0
Region	North America	8.0	4.4	3.6	55.1	44.9	7.1	3.1	3.9	44.2	55.8
Region	South Asia	59.3	53.6	5.7	90.4	9.6	56.8	50.6	6.3	89.0	11.0
Region	Sub-Saharan Africa	86.9	70.1	16.8	80.7	19.3	83.8	62.1	21.6	74.2	25.8
Income Level	High income	8.5	6.1	2.4	71.4	28.6	6.7	4.3	2.4	63.5	36.5
Income Level	Upper middle income	46.8	42.9	3.9	91.6	8.4	41.3	37.8	3.4	91.7	8.3
Income Level	Lower middle income	59.0	53.2	5.8	90.2	9.8	54.9	49.1	5.7	89.6	10.4
Income Level	Low income	92.7	71.6	21.1	77.3	22.7	92.8	62.3	30.6	67.1	32.9
Lending	Part 1	7.6	5.4	2.2	70.5	29.5	5.9	3.6	2.2	61.6	38.4
Lending	IBRD	51.9	48.4	3.5	93.2	6.8	47.0	44.7	2.3	95.2	4.8
Lending	IDA / Blend	83.9	64.5	19.4	76.9	23.1	82.7	54.7	28.1	66.1	33.9

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Notes: Shapley value decomposition by learning deprivation (LD) and schooling deprivation (SD) dimensions, for a methodological description please see Azevedo et al (2012); (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

How does learning poverty differ across regions within countries?

Another important dimension for disaggregation is the within-country spatial differences in learning poverty. Tracking these differences enables the spatial targeting of specific policy packages, and when temporal comparisons are available, it makes possible the identification of "what works" from within an educational system. This in turn creates opportunities for packaging, replicating, and scaling these best practices to reduce learning poverty. Using data from the Brazilian National Learning Assessment (NLA, *Prova Brasil*) and the Education Management information system, we were able to create a national learning poverty measure that captures both the out-of-school population and the share of 5th-grade students achieving the Brazilian minimum proficiency level (defined in the national learning assessment scale). Note that this is a national measure has not yet been "equated" to the GAML MPL threshold to facilitate international comparisons, although that could be a useful next step.

Because the Brazilian NLA is a census and is designed for temporal comparability of learning, we are able understand the spatial distribution of learning poverty in Brazil in two points in time, which allows an examination of how a national learning poverty has changed both temporally and spatially. One particular region of Brazil stands out: the state of Ceará, which is one of the poorest states in the North-East of Brazil but has among the lowest learning poverty rates (highlighted in Figure 6). While this success in reducing learning poverty cannot be attributed to any single policy or program, it is grounded on a series of reforms that include results-based financing and the provision of technical assistance to the municipal education boards (Loureiro and Cruz 2020, Loureiro et al. 2020). These results, which are well known in Brazil, have informed the design of programs to improve quality of education at the end of primary in other Brazilian states.[50]

Figure 6 Learning Poverty by Brazilian Municipality (national definition)

a. National Learning Poverty - 2013 *b. National Learning Poverty - 2017*

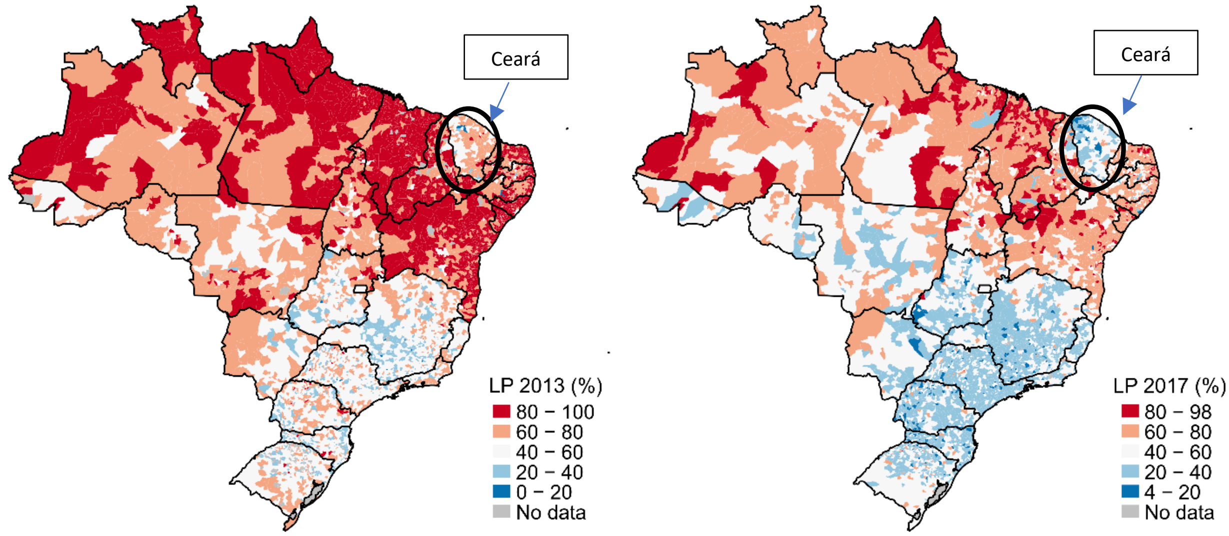

Source: Authors' calculations using IBGE and INEP/MEC data
Notes: The LP number for Brazil is calculated at the municipal level, using the microdata from Prova Brasil, INEP School Census and IBGE population estimative. The National MPL threshold used was 200 points in Portuguese.

5. Where we are headed: Simulations of future changes in learning poverty

This section moves from estimating current levels of learning poverty to trying to understand where learning poverty was headed before COVID-19 hit. Because the pandemic has altered that trajectory—although we lack the data to know by how much—the simulations presented here should be viewed as a best-case scenario.

The estimates above show that over half of children in low- and middle-income countries are not reaching minimum reading proficiency, and that the figure could be even higher. For low-income countries, the share is close to 80%. A crucial question, then, is how this learning poverty rate has changed in recent years. Is it declining rapidly enough to ensure that in 2030, all children will be proficient in reading by age 10, or at least by the end of primary school? The analysis in this section

[50] See Costa and Carnoy 2015, Loureiro and Cruz 2020, and Loureiro et al. 2020 for details. Programs influenced by the Ceará model include the PARC (Parceria para Alfabetização via Regime de Colaboração) program designed by Associação Bem Comum, Fundação Lemann and Instituto Natura.

shows that it is not: eliminating learning poverty by 2030 would require historically unprecedented progress.[51]

Progress in reducing learning poverty since 2000: Distribution of changes

To assess this question, we examine spells of within-country improvement in learning poverty within a given assessment, for the subsample of countries in which the same assessment was applied more than once between 2000 and 2017. This approach gives us 72 change spells for 50 low- and middle-income countries, distributed across the various assessments that can be compared over time. Figure 7 shows the distribution of spells for all countries in the database, arranged from lowest to highest rate of improvement.

Figure 7 Distribution of annualized changes in learning poverty for low- and middle-income countries, 2000-2017

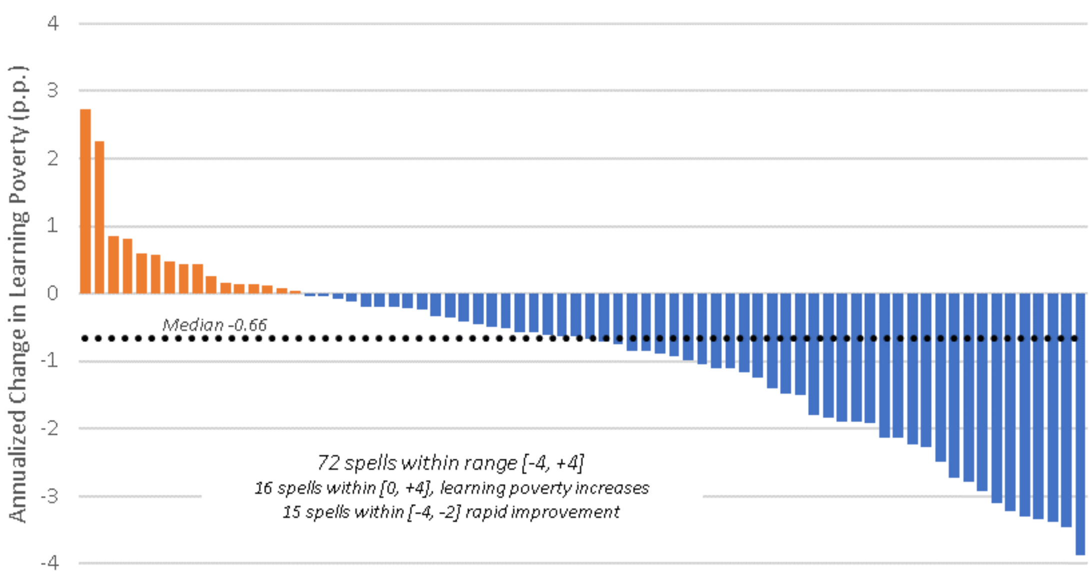

Source: Authors' calculations using the Global Learning Assessment Database; and UIS Enrollment Data.
Note: Each bar represents one spell, showing the annualized change in learning poverty for a given country between two points in time; includes only temporally comparable spells within the same assessment and grade, for low- and middle-income countries. See Table 5 for the breakdown of these 72 spells by assessment. Negative values represent a decrease in learning poverty—that is, an improvement in outcomes.

[51] For more detailed information on the spells database, please check Annex 8.

Three important points come out of this simple figure:

- ***The median annual reduction in learning poverty, across all spells, is less than 1 percentage point per year.*** With a global learning poverty rate of 53% in 2015, this suggests that unless improvement accelerates dramatically from recent historical patterns, the world will fall well short of eliminating learning poverty by 2030.[52]
- ***In around 20% of the recorded episodes of annualized change, learning poverty increases.*** Thus, although global rates of reading have been improving, there is no guarantee of progress in individual countries.
- ***Nevertheless, there have been some cases of rapid improvement.*** About 20% of the recorded spells show annualized reductions in learning poverty of 2 percentage points or more. Even if this tail of the distribution reflects some statistical noise, this indicates that it is possible to make rapid progress (in some cases, through a combination of better learning for enrolled students and increased enrollment).

Beyond what we learn from the averages and distribution, there can also be lessons from the individual country spells, especially in the tails. We need to be cautious with interpretation; these are just 3- to 4-year spells, so mean reversion can cause a spell of rapid reduction in learning poverty to be followed by a slowdown in the next period. Lessons drawn from sustained rapid improvements are therefore likely to be more reliable. For analysis of such cases, see Crawford and others (Forthcoming), which uses such rapid improvers to draw lessons on effective early-grade literacy interventions.

The data presented in this section have been for individual country spells, unweighted for population, but adjusted for multiple observations from the same country.

[52] As a reminder, because the latest learning data is from 2017, these estimates do not take the effects of the COVID-19 pandemic into account.

Global simulation of learning poverty to 2030: Business-as-usual vs high scenario

This section presents simulations of how our global learning poverty estimate can be expected to change between 2015 and 2030. We offer two scenarios: a business-as-usual scenario, based on the rates of change from 2000 to 2017; and a high scenario, which assumes that each country progresses at the most rapid rates achieved in its region over that period. The bottom line is that in neither scenario do the low- and middle-income countries achieve universal reading proficiency by 2030, even if we define "universal" loosely to mean 95% of the population, and even without taking the effects of COVID-19 into account.

These scenarios are based on the historical annualized reductions in learning poverty for each region and globally (Table 9).

Table 9 Annualized change in learning poverty (in percentage points) by country group, 2000-2017

Country Groups		Business as Usual (BaU)[(1)]	60th Percentile	70th Percentile	80th Percentile (High)	90th Percentile
	Overall	0.68	0.87	1.34	1.85	2.56
Region	East Asia and Pacific[(2)]	0.68	0.87	1.34	1.85	2.56
	Europe and Central Asia	0.68	0.86	1.06	1.25	2.24
	Latin American and Caribbean	0.62	0.64	1.29	1.81	2.02
	Middle East and North Africa	0.36	0.89	1.16	2.28	2.92
	South Asia[(2)]	0.68	0.87	1.34	1.85	2.56
	Sub-Saharan Africa	0.92	1.10	1.93	2.49	3.46
Income Level	High income	0.53	0.57	0.63	0.68	3.34
	Upper middle income	0.86	1.06	1.51	1.89	2.49
	Lower middle income	1.10	1.16	1.91	2.72	3.22
	Low income	-0.08	0.01	0.11	1.02	1.93
Initial Learning Poverty	0-25% Learning Poverty	0.58	0.68	0.86	1.25	3.10
	25-50% Learning Poverty	1.06	1.51	1.93	2.72	3.34
	50-75% Learning Poverty	0.62	0.86	1.10	1.81	2.14
	75-100% Learning Poverty	0.92	1.48	1.89	2.14	2.49

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Senate weights are used to ensure an equal contribution by each of the countries in the sample of spells. (1) Business as Usual (BaU) scenario reflects the country-specific growth rates; and for countries with no spells, the median of the group was used. (2) Global values were used due to an insufficient number of spells in the region.

The historical data also allow us to understand how much changes in schooling and learning affect the observed variations in learning poverty. Table 10 shows that shows that two-thirds of the average improvement of learning poverty can be attributed to improvements in learning, while one-third comes from reductions in schooling deprivation. However, the regional differences are striking. While in the Middle East all of the improvement can be attributed to schooling dimensions, in Latin America the reverse is true; in that region, improvements in learning more than compensated the worsening of the schooling indicator. In Sub-Saharan Africa we find a more balanced pattern, in which improvements in both schooling and learning have been almost equally important.

Table 10 Decomposition of the average reduction *in learning poverty by learning and schooling deprivations*

Region	Annualized Reduction in Learning Poverty (p.p.)	Annualized Reduction in Learning Deprivation (LD) (p.p.)	Annualized Reduction in Schooling Deprivation (SD) (p.p.)	Percentage of Annualized Reduction in Learning Poverty Explained by LD	Percentage of Annualized Reduction in Learning Poverty Explained by SD
Overall	0.93	0.59	0.35	63%	37%
East Asia and Pacific	2.72	2.11	0.62	77%	23%
Europe and Central Asia	0.74	0.32	0.42	43%	57%
Latin American and Caribbean	0.84	0.91	-0.07	109%	-9%
Middle East and North Africa	0.57	-0.16	0.73	-28%	128%
Sub-Saharan Africa	1.29	0.67	0.62	52%	48%

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Decomposition of change in learning poverty by the change in its two dimensions: learning deprivation (LD) and schooling deprivation (SD); for a methodological description, please see Azevedo et (2012). Includes all spells considered in the simulation (N=72). South-Asia has been omitted, as there was data on changes for only one country. Table 22 in the Annex reports the same mean values decomposed by this exercise.

Business-as-usual

For the business-as-usual scenario, we start from the current estimated level for each country and apply the country's historical rate of improvement (median value) to simulate the future learning poverty rate. For the many countries that have not applied the same reliable assessment at least twice—that is, the countries not represented in Figure 7 above—we lack good country-specific data on historical changes in learning poverty. In those cases, we use the median regional rate of change for the simulation.

Figure 8 shows the results of this simulation. Under business as usual, the mean population-weighted change in learning poverty is a reduction of 1 percentage point per year. At this pace, starting from a learning poverty rate of 53% in 2015, the rate will still be an estimated 44% in 2030—meaning that 44% of late-primary children in low- and middle-income countries will still not have reached minimum proficiency in reading.[53]

[53] This message is consistent with the concerns about slow progress raised by UNESCO (2019), a report for the 2019 UN High-Level Political Forum on the SDGs.

Figure 8 Learning poverty under two scenarios, 2015-30 (simulation)

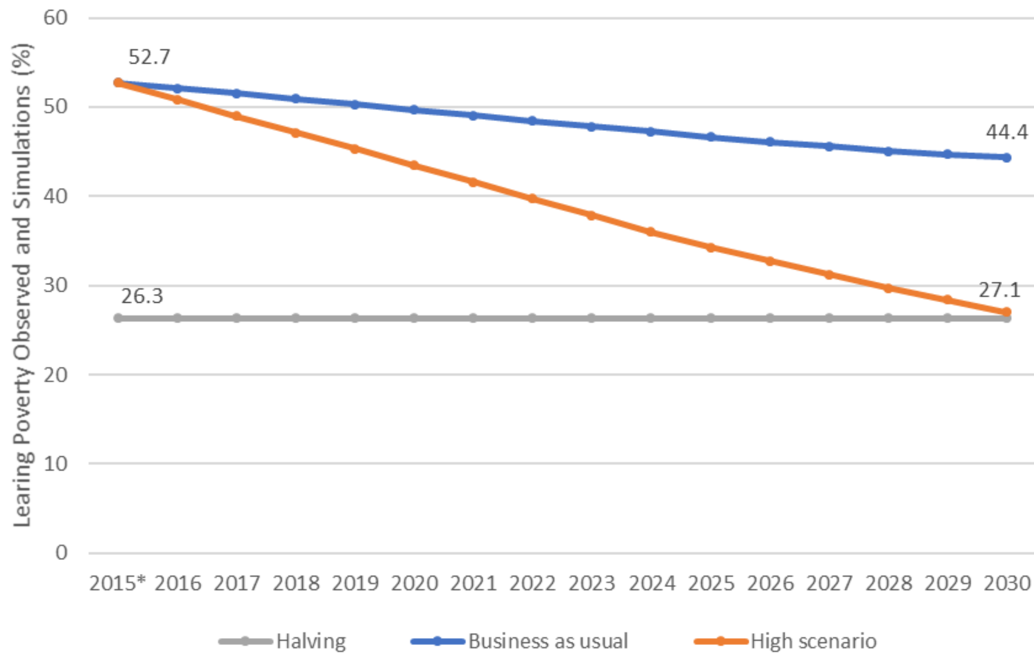

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.

High scenario

But perhaps business-as-usual is too pessimistic. For much of the 2000-2017 period, policymakers did not focus much on learning. To be sure, the substantial improvement in primary-school completion throughout this period should have reduced learning poverty. As described above, all children out of school count as learning poor in this statistic, so enrolling them should have lowered learning poverty, as long as at least some of them learned to read (as Table 10 suggests happened in practice). But policymakers' neglect of learning likely constrained the gains that were possible. Now that the international education community is focusing more on learning—a shift reflected in Sustainable Development Goal 4 and the learning indicators being used to track it—progress could well accelerate.

How much? To bound the possible improvements, the high scenario uses the 80th percentile of actual experience. Specifically, it assumes that every country can reduce learning poverty as quickly as the 80th-percentile country in its region did during the 2000-2017 period.[54] As Figure 8 shows, this allows much more rapid progress. Nevertheless, the learning poverty rate for 2030 will still be 27%—considerably higher than the rights-based target of zero, or even an alternate target set at the single-digit rates found in many wealthier countries.

---

[54] We also assume that if a country's historical rate of progress exceeds the regional 80th percentile, it will remain at that higher rate. An alternative assumption for this high scenario would be to set the rate of improvement for every country equal to the global 80th percentile. It turns out, however, that the regional 80th-percentile scenario yields faster global reductions in learning poverty. This is because the regional 80th percentile is higher in Sub-Saharan Africa than in other regions, and the regional scenario applies that more rapid rate to Africa's large child population.

Regional progress under the two scenarios

Table 11 reports both simulation scenarios by regions. The table indicates that between 2015 and 2030, learning poverty in Africa might drop 12 percentage points in the business-as-usual scenario, and 37 points in the high scenario. Although substantial, this learning poverty reduction is not enough to reduce Africa's share of the global learning-poor population, especially as demographic pressure is likely to strain education systems in Africa. Even under the high scenario, almost half of children in Africa and a third of children in South Asia will still be learning-poor in 2030. By contrast, two regions—East Asia and the Pacific, and Europe and Central Asia—will have single-digit learning poverty rates under either scenario.[55]

With these shifts, learning poverty might be even more unequally distributed between regions in 2030 than it was in 2015. At the baseline, learning poverty is concentrated in three main regions, with 82% of all learning-poor living in South Asia, Sub-Saharan Africa, and East Asia. In the business-as-usual scenario, concentration in the top three regions will increase to 88%, and the spatial pattern will also shift, with the Middle East taking the place of East Asia in that group. And in the high scenario, as learning poverty falls sharply, concentration in the top three regions will increase even more, to 92%. Latin America will take the place of East Asia among the regions with the most learning-poor, and Africa will be the only region that will see its share of the global learning-poor increase substantially, from 37% to 52%.

Table 11 Learning poverty rates in 2030 under two scenarios (simulation using spells by region)

	Population (millions)		Learning Poverty (%)			Region's Share of Global Learning-Poor (%)		
Region	2015	2030	2015	2030	2030	2015	2030	2030
			Base	BaU	High	Base	BaU	High
East Asia and Pacific	137	138	21.2	6.4	0.3	10.0	3.3	0.3
Europe and Central Asia	27	32	13.3	3.4	1.8	1.2	0.4	0.4
Latin American and Caribbean	53	51	50.8	39.6	23.7	9.4	7.5	7.4
Middle East and North Africa	33	44	63.3	56.2	28.9	7.1	9.2	7.8
South Asia	175	167	58.2	50.5	31.5	35.2	31.6	32.4
Sub-Saharan Africa	123	171	86.7	75.0	49.4	37.0	48.0	51.9
Overall	548	602	52.7	44.4	27.1	100.0	100.0	100.0

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.
Note: Business as Usual (BaU) scenario uses country-specific growth rates; for countries with no spells, the regional average was used. High scenario considers the rate of improvement equal to the regional 80th percentile. All figures consider only Part 2 countries.

How to interpret these simulations

The most important caveat is that these simulations do not represent a forecast. The simulations are based on historical rates of progress, using the best available data (with all its imperfections). They do not model the many factors influencing progress in each country—such as income growth, policy shifts,

[55] For these higher-performing regions, "single digits" is probably more useful as a guide to the future for these regions than the precise values in the table. Recall that these are results of a simulation exercise, not projections. In the case of ECA in particular, because the region has such a low initial value, the simulations to 2030 are in the range where simply extrapolating from past rates of improvement leads to unlikely outcomes. Almost no country has a learning poverty rate as low as 2 percent, so although these low- and middle-income ECA countries will likely achieve single digits, they won't likely reach the that low a rate.

conflict, migration, pandemics and any other shocks. Beyond this, the simulations do not take into account the impacts of the COVID-driven school closures and economic shock on schooling and learning, because the data are not yet available to construct learning poverty estimates for 2020-21. Nor do we know yet whether they will affect only the level or also the long-term trend rate of decline. But it is clear that the pandemic will make the task of eliminating learning poverty even more challenging.

What these simulations are intended to do, instead, is indicate the magnitude of the changes that were already necessary even before COVID, relative to what we've seen so far this century. The bottom line is that eliminating global learning poverty by 2030—or even some years after that—would require unprecedented rates of improvement.

Implications for target-setting

Meaningful action to improve foundational skills requires evidence-based targets to focus policies and operational engagement. The learning poverty indicator and this analysis can be used by policymakers and external actors like the World Bank to inform such targets, and indeed that has already begun to happen. Based on an early version of this analysis, the World Bank adopted the following target for its operational work with low- and middle-income countries: "By 2030, reduce by at least half the share of children who cannot read by age 10" (World Bank 2019). As shown above, this target can be attained (starting from the 2015 global level) if all countries accelerate their progress to at least the 80th percentile of the post-2000 distribution of spells of improvements. In other words, every country needs to perform like a country that has cut learning poverty rapidly enough to place it among the top 20% of improvers since 2000 (and at the same time, those that have exceeded that 80th-percentile rate need to sustain their previous rates of progress).[56]

It is important to note that while this target focuses on reducing the share of children with very low performance in low- and middle-income countries, for most of these countries this will also mean an improvement in their *average* performance. Evidence that in moving from very low mean performance to at least middling levels of mean performance, countries do so mostly by significantly cutting down the share of children at very low levels of performance. Policies to improve learning among lower-performing schools and pupils (the "bulging tail" of the distribution) are required to improve learning equitably and to reduce unfair inequality (Crouch and Rolleston 2017; Crouch, Rolleston, and Gustafsson 2020; Mullis et al 2016).

A final point is that the COVID-19 pandemic has made the challenge greater. Although we do not yet have any learning data for 2020 from these comparable assessments, simulations suggest the school closures and other shocks from the pandemic could increase the global learning poverty rate by up to 10 percentage points (Azevedo 2020). This shock could require revisiting learning targets and the policy packages necessary to reach them. At the same time, some of the interventions that will likely be most effective in reversing the learning losses caused by school closures could also lead the most rapid long-term reductions in learning poverty (Kaffenberger 2020, World Bank 2020).

[56] The 80th percentile rate is based on distribution of performance using the data available as of the writing of this paper. As more temporally comparable data on learning poverty becomes available it will be critical to update these estimates, and revise our expectations regarding the feasibility of halving learning poverty by 2030.

6. Conclusion

The learning poverty indicator presented in this paper is motivated by the urgent need to solve the learning crisis in low- and middle-income countries, a need that is made even more urgent by the COVID-19 shock. Rebuilding education systems, building countries' human capital, and advancing toward the Sustainable Development Goals all require ensuring that all children acquire foundational skills for life and work. The ability to read with comprehension is a fundamental skill that every education system around the world strives to impart by late in primary school—generally by age 10. The evidence presented in this paper adds to the considerable evidence that not only are many low- and middle-income countries failing to meet the higher-level education goals, but they are also struggling to impart these more basic foundational skills.

To highlight this problem and focus efforts to solve it, this paper has described in detail the learning poverty indicator launched by UIS and the World Bank in 2019 and has quantified the extent of the problem globally. Drawing on a dataset, newly assembled for this purpose, that covers 80% of children in low- and middle-income countries, this paper has documented that about 53% of all children in low- and middle-income countries cannot read age-appropriate material by age 10. In other words, more than half of the children in the world suffer from learning poverty.

We have also presented the most comprehensive rigorous estimates of historical progress in reducing learning poverty in low- and middle-income countries. Our estimates reveal that at the 2000-2017 rate of improvement, the learning poverty rate would fall only to 44% by 2030. An early version of this paper therefore inspired a new medium-term target to guide the World Bank's work in low- and middle-income countries: cut learning poverty by at least half by 2030. Starting at the 53% rate estimated for 2015, this target would be attainable if all countries match the 80th percentile of historical performance in their regions, which meant more than doubling the rate of progress globally. Further analysis presented here has shown that the learning poverty indicator, although simple in design, is robust and is highly correlated with metrics of learning covering other age groups. In addition, we have presented disaggregated analysis of Learning poverty along gender and regional dimensions.

In the wake of COVID-19, the global learning poverty rate is doubtless even higher than it was in 2015, and as a result the acceleration in learning and schooling improvement needed to halve it will be even greater. We hope that countries will be able to use this indicator to galvanize action and track progress as they recover from the pandemic.

References

Adelman, S., Gilligan, D.O. & Lehrer, K. 2008. How effective are food for education programmes? A critical assessment of the evidence from developing countries. Food Policy Review 9. Washington, DC, International Food Policy Research Institute.

Angrist, Noam, Simeon Djankov, Pinelopi Koujianou Goldberg, and Harry Anthony Patrinos. 2019. "Measuring Human Capital." World Bank Policy Research Working Paper No. 8742.

Azevedo, João Pedro, Viviane Sanfelice and Minh C. Nguyen. 2012. "Shapley Decomposition by Components of a Welfare Aggregate," MPRA Paper 85584, University Library of Munich, Germany.

Azevedo, João Pedro. 2020. Learning Poverty: Measures and Simulations. Policy Research Working Paper No. 9446. World Bank, Washington, DC. http://hdl.handle.net/10986/34654

Azevedo, João Pedro and Silvia Montoya. 2021. How the SDG 4.1.1 Framework and Learning Poverty Can Help Countries Focus Their Education Policy Response to COVID-19. Posted on February 17, 2021 by Data for Sustainable Development: UNESCO Institute for Statistics Blog. https://sdg.uis.unesco.org/2021/02/17/how-the-sdg-4-1-1-framework-and-learning-poverty-can-help-countries-focus-their-education-policy-response-to-covid-19/

Bhutta, Z.A., Das, J.K., Rizvi, A., Gaffey, M.F., Walker, N., Horton, S., Webb, P., Lartey, A., & Black, R.E. 2013. Evidence-based interventions for improvement of maternal and child nutrition: what can be done and at what cost? The Lancet, 382(9890): 452–477.

Blau, David M., Philip K. Robins. 1988. Child care costs and family labor supply. Review of Economics and Statistics 70(3):374–81.

Blau, David, Janet Currie. 2006. Pre-school, day care, and after-school care: who's minding the kids? In Handbook of the Economics of Education Vol. 2, chapter 20, 1163–278, Elsevier.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *Quarterly Journal of Economics* 129 (4): 1553–1623.

Crawford, Michael, and others. Forthcoming. "Attaining the Learning Target: Policies to Promote Literacy for All Children." Washington, DC: World Bank.

Costa, Leandro Oliveira, and Martin Carnoy. 2015. "The Effectiveness of an Early-Grade Literacy Intervention on the Cognitive Achievement of Brazilian Students." Educational Evaluation and Policy Analysis 37 (4): 567–90. https://doi.org/10.3102/0162373715571437.

Crouch, Luis, Caine Rolleston, and Martin Gustafsson. 2020. "Eliminating global learning poverty: The importance of equalities and equity." *International Journal of Educational Development*.

Easterly, William, Jozef Ritzen, and Michael Woolcock. 2006. "Social Cohesion, Institutions, And Growth," Economics and Politics, Wiley Blackwell, vol. 18(2), pages 103-120, July.

Filmer, Deon, Amer Hasan, and Lant Pritchett. 2006. "A Millennium Learning Goal: Measuring Real Progress in Education." Center for Global Development Working Paper 97.

Filmer, Deon, Halsey Rogers, Noam Angrist, and Shwetlena Sabarwal. 2018. "Learning-Adjusted Years of Schooling (LAYS): Defining a New Macro Measure of Education." World Bank Policy Research Working Paper No. 8591.

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

Hanushek, Eric, and Ludger Woessmann. 2015. *Universal Basic Skills: What Countries Stand to Gain.* Organisation for Economic Co-operation and Development.

IEA. 2016. "PIRLS 2016 International Results in Reading: About PIRLS 2016." TIMSS and PIRLS International Studies Center, Boston College. http://timssandpirls.bc.edu/pirls2016/international-results/wp-content/uploads/structure/PIRLS/0.-about-pirls-2016/P16-About-PIRLS-2016.pdf

Kaffenberger, Michelle, and Lant Pritchett. 2017. "The Impact of Education versus the Impact of Schooling: Schooling, Reading Ability, and Financial Behavior in 10 Countries." Background paper, World Bank, Washington, DC.

Khan, Zebun Nisa. 2016. Role of Education in Building Social Cohesion. International Journal of Secondary Education. Vol. 4, No. 2, pp. 23-26. doi: 10.11648/j.ijsedu.20160402.12

Kraay, Aart C. 2018. "Methodology for a World Bank Human Capital Index." World Bank Policy Research Working Paper No. 8593.

Jerrim, John, Luis Alejandro Lopez-Agudo, Oscar D. Marcenaro-Gutierrez, and Nikki Shure. 2017. "What happens when econometrics and psychometrics collide? An example using the PISA data," Economics of Education Review, Elsevier, vol. 61, pages 51-58.

Kaffenberger, Michelle. 2020. "Modeling the long-run learning impact of the COVID-19 learning shock: Actions to (more than) mitigate loss." *RISE Insight Series* 17.

Loureiro, Andre, and Louisee Cruz. 2020. "Achieving World-Class Education in Adverse Socioeconomic Conditions: The Case of Sobral in Brazil." Washington, DC: World Bank.

Loureiro, Andre, Louisee Cruz, Ildo Lautharte, and David K. Evans. 2020. "The State of Ceara in Brazil is a Role Model for Reducing Learning Poverty." Washington, DC: World Bank.

Lyon, G. Reid, and Vinita Chhabra. 2004. "The Science of Reading Research." *Educational Leadership* 61:6 (March), 12-17.

Mullis, I. V. S., Martin, M. O., Foy, P., and Hooper, M. 2017. *PIRLS 2016 International Results in Reading*.

Mullis, I. V. S., Martin, M. O., Foy, P., & Hooper, M. 2016. TIMSS 2015 International Results in Mathematics. Retrieved from Boston College, TIMSS & PIRLS International Study Center: http://timssandpirls.bc.edu/timss2015/international-results/)

OECD. 2017. *PISA 2015 Assessments and Analytical Framework: Science, Reading, Mathematic, Financial Literacy and Collaborative Problem Solving.*

OREALC/UNESCO (2014) Comparison of results of the Segundo y Tercer Estudio comparativo y explicativo: SERCE y TERCE, 2006-2013. Santiago.

Oye, Mari, Lant Pritchett, and Justin Sandefur. 2016. “Girls’ Schooling Is Good, Girls’ Schooling with Learning Is Better.” Education Commission, Center for Global Development, Washington, DC.

Patrinos, Harry Anthony, and Noam Angrist. 2018. “Global Dataset on Education Quality: A Review and Update (2000-2017).” World Bank Policy Research Working Paper No. 8592.

Pritchett, Lant. 2013. *The Rebirth of Education: Schooling Ain’t Learning.* Washington, DC: Center for Global Development. Baltimore: Brookings Institution Press.

Shaywitz, Sally E. 2003. *Overcoming dyslexia: A new and complete science-based program for reading problems at any level*. New York: Knopf.

Skibbe, Lori E., Janelle J. Montroy, Ryan P. Bowles, and Frederick J. Morrison. 2019. "Self-regulation and the Development of Literacy and Language Achievement from Preschool through Second Grade." *Early Childhood Research Quarterly* 46(1), 240-251.

UNESCO. 2019. *How Fast Can Levels of Proficiency Improve? Examining Historical Trends to Inform SDG 4.1.1 Scenarios*. Information Paper No. 62 December 2019

UNICEF (United Nations Children’s Fund). 2020. “Framework for Opening Schools.” New York: UNICEF. https://www.unicef.org/media/68366/file/Framework-for-reopening-schools-2020.pdf.

UIS. 2017a. “More Than One-Half of Children and Adolescents Are Not Learning Worldwide.” UNESCO Institute for Statistics Fact Sheet No. 46 (September).

UIS. 2017b. “Counting the Number of Children Not Learning: Methodology for a Global Composite Indicator for Education.” UNESCO Institute for Statistics Information Paper No. 47 (October).

UIS. 2017c. *Mind the Gap: Proposal for a Standardised Measure for SDG 4– Education 2030 Agenda*. Information Paper No. 46 October 2017

UIS. 2019. Protocol for Reporting on SDG Global Indicator 4.1.1. GAML6/WD/2. http://gaml.uis.unesco.org/wp-content/uploads/sites/2/2019/05/GAML6-WD-2-Protocol-for-reporting-4.1.1_v1.pdf

UIS and GAML. 2019. “Minimum Proficiency Levels (MPLs): Outcomes of the Consensus-Building Meeting.” UNESCO Institute of Statistics and the Global Alliance to Monitor Learning. http://gaml.uis.unesco.org/wp-content/uploads/sites/2/2019/02/MPLs_revised_doc_20190204.docx

Valerio, Alexandria, María Laura Sánchez Puerta, Namrata Raman Tognatta, and Sebastián Monroy-Taborda. 2016. “Are There Skills Payoffs in Low- and Middle-Income Countries? Empirical Evidence Using Step Data.” Policy Research Working Paper 7879, World Bank, Washington, DC.

Vries, Willem de. 2002. Dimensions of Statistical Quality: A discussion note about the quality initiatives of some international organizations. Inter-agency Meeting on Coordination of Statistical Activities SA/2002/6/Add. New York, 17-19 September 2002. https://unstats.un.org/unsd/accsub/2002docs/sa-02-6add1.pdf

WFP (World Food Programme). 2020. Joint Advocacy Brief: Stepping up Effective School Health and Nutrition. Rome: WFP. https://www.wfp.org/publications/joint-advocacy-brief-stepping-effective-school-health-and-nutrition.

World Bank. 2018a. *World Development Report 2018: Learning to Realize Education's Promise.* Washington, DC: World Bank. https://openknowledge.worldbank.org/handle/10986/28340 License: CC BY 3.0 IGO.

World Bank. 2018b. *Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle.* Washington, DC: World Bank. License: Creative Commons Attribution CC BY 3.0 IGO

World Bank. 2019. *Ending Learning Poverty: What Will It Take?* World Bank, Washington, DC. World Bank. License: Creative Commons Attribution CC BY 3.0 IGO http://hdl.handle.net/10986/32553

World Bank. 2020. *The COVID-19 Pandemic: Shocks to Education and Policy Responses*.

Annexes

Annex 1: Validity of reading as a proxy for other foundational skills

One argument for focusing on foundational literacy via the learning poverty indicator is that it can serve as an easily understood proxy for success in building other foundational skills, particularly at the level of the educational system. If it proves to be a good proxy, then low learning poverty rates can be a marker of the health of the education system—in the same way that low rates of child stunting are a marker of healthy early childhood development. The theory behind this is that systems that ensure that all children can read are likely to succeed in helping them learn other subjects as well. The data bear this out: across countries and schools, proficiency rates in reading are highly correlated with proficiency in other subjects in the same grades (Table 12). For example, the correlation between a country's reading score on the Progress in International Reading Literacy Study (PIRLS) assessment and its TIMSS math score is 0.92, and the cross-subject correlations within other assessments are strong too, ranging from 0.87 to 0.97 at the municipality and school levels. On the OECD PISA assessments of 15-year-olds, the cross-country correlation between reading scores on math or science remains high, even after controlling for socio-economic status. At the student level, this correlation unsurprisingly weakens significantly, as shown in the last column of Table 12; still, it remains strongly positive in all of the tests summarized here, at between 0.561 and 0.895. Beyond these correlations with other cognitive scores, evidence shows that higher levels and earlier development of language and reading skills is associated with early development of a child's self-regulation, a fundamental socioemotional skill (Skibbe et al., 2019).

Table 12 Correlation of reading scores with math and science scores by assessment and level of aggregation

	Subject	Country	Municipality	School	Student
PIRLS-TIMSS (4th grade)	Math	0.923			
	Science	0.965			
LLECE	Math	0.904		0.867	0.561
(6th grade)	Science	0.942		0.869	0.587
PISA-D	Math	0.939		0.936	0.818
(15-year-olds)	Science	0.974		0.958	0.868
PISA	Math	0.949		0.939	0.851
(15-year-olds)	Science	0.978		0.971	0.895
BRAZIL (5th grade)	Math		0.962	0.943	0.725
(9th grade)	Math		0.937	0.907	0.672

Source: Authors' calculations using the Global Learning Assessment Database.
Note: Correlations computed by Grade. The top row shows the country-level correlation between scores on the Latin American Laboratory for Assessment of the Quality of Education (LLECE), Program for International Student Assessment (PISA), and PISA for Development (PISA-D), the figure shows the correlation between reading and math/science scores within the given assessment. Correlations between scores on Progress in International Reading Literacy Study (PIRLS) literacy and Trends in International Mathematics and Science Study (TIMSS) math/science assessments are at the country level. In the case of Brazil, results for Prova Brasil were used. The table shows the correlation between reading and math/science scores within the given assessment.

Table 13 Relationship between reading and math proficiency (PISA)

	Model 1	Model 2	Model 3	Model 4	Model 1	Model 2	Model 3	Model 4
No weights								
Math	0.854***	0.833***	0.829***	0.838***				
	(0.027)	(0.038)	(0.038)	(0.036)				
Science					0.963***	0.934***	0.935***	0.936***
					(0.021)	(0.029)	(0.029)	(0.029)
R^2-Adj	0.891	0.910	0.913	0.914	0.937	0.945	0.948	0.947
N	417	417	417	412	420	420	420	415
Weighted								
Math	0.891***	0.893***	0.886***	0.892***				
	(0.029)	(0.046)	(0.044)	(0.045)				
Science					0.985***	0.973***	0.974***	0.975***
					(0.025)	(0.031)	(0.030)	(0.030)
R^2-Adj	0.886	0.908	0.911	0.912	0.945	0.955	0.957	0.957
N	417	417	417	412	420	420	420	415

Source: Authors’ calculations using the Global Learning Assessment Database (PISA) and WDI data.
Notes: Model 1: no controls; Model 2: region and income level; Model 3: Model 2 + temporal fixed effect; and Model 4: Model 3 + GDP per capita. * p<0.05, ** p<0.01, *** p<0.001; for details on these calculations, please see https://github.com/worldbank/learningpoverty

Annex 2 PIRLS and TIMMS temporal comparability

Table 14 Non-comparable spells in PIRLS and TIMMS

Country	Assessment	N	Non comparable spells
Israel	PIRLS	4	2001-2011, 2001-2016, 2006-2011, 2006-2016
Morocco	PIRLS	4	2001-2011, 2001-2016, 2006-2011, 2006-2016
Poland	PIRLS	2	2006-2016, 2011-2016
Qatar	PIRLS	2	2006-2011, 2006-2016
South Africa	PIRLS	3	2006-2011, 2006-2016, 2011-2016
Armenia	TIMSS	4	2003-2007, 2007-2011, 2007-2015, 2011-2015
Kazakhstan	TIMSS	2	2007-2011, 2007-2015
Kuwait	TIMSS	2	2007-2011, 2007-2015
Morocco	TIMSS	4	2003-2011, 2003-2015, 2007-2011, 2007-2015
Poland	TIMSS	1	2011-2015
Qatar	TIMSS	2	2007-2011, 2007-2015
Yemen	TIMSS	2	2003-2007
Total		32	

Source: PIRLS and TIMMS documentation, by the International Association for the Evaluation of Educational Achievement (IEA).
Note: Temporal comparability is impaired in these cases due to sampling and/or implementation differences.

Annex 3. Learning assessment population coverage

Table 15 Population and country coverage by country groups, latest available learning assessment

	Country Group	All Countries: N countries with data	All Countries: N countries total	All Countries: Population with data (millions)	All Countries: Population Coverage (%)	Low- and Middle-Income Countries*: N countries with data	Low- and Middle-Income Countries*: N countries total	Low- and Middle-Income Countries*: Population with data (millions)	Low- and Middle-Income Countries*: Population Coverage (%)
	Overall	116	217	515	84.2	74	144	454	82.9
Region	East Asia and Pacific	15	37	140	93.8	8	23	130	94.7
Region	Europe and Central Asia	39	58	45	89.7	15	23	22	82.5
Region	Latin American and Caribbean	19	42	48	89.2	18	30	48	89.6
Region	Middle East and North Africa	15	21	29	79.2	7	12	25	77.8
Region	North America	2	3	23	100.0	N/A	N/A	N/A	N/A
Region	South Asia	5	8	171	98.1	5	8	171	98.1
Region	Sub-Saharan Africa	21	48	58	46.8	21	48	58	46.8
Income Level	High income	47	79	64	98.6	6	10	4	99.7
Income Level	Upper middle income	33	60	166	94.1	32	58	165	94.1
Income Level	Lower middle income	23	47	233	80.7	23	46	233	80.8
Income Level	Low income	13	31	51	63.3	13	30	51	64.8
Lending	Part 1	42	73	61	94.9	N/A	N/A	N/A	N/A
Lending	IBRD	47	68	352	95.9	47	68	352	95.9
Lending	IDA / Blend	27	76	102	56.5	27	76	102	56.5

Source: Authors' calculations using the Global Learning Assessment Database; and UN Population numbers.

Note: Data includes assessments since 2000 (See Table 3 for assessments within the preferred reporting window only); Population coverage considering share of 10- to 14-year-olds in countries with assessment data; Regions: East Asia and Pacific (EAP), Europe and Central Asia (ECA), Latin American and Caribbean (LAC), Middle East and North Africa (MNA), North America (NAC), South Asia (SAR), and Sub-Saharan Africa (SSA); Lending Categories: International Development Association (IDA); International Bank for Reconstruction and Development (IBRD); and IDA-eligible based on per capita income levels and are also creditworthy for some IBRD borrowing (Blend). (*) Low- and Middle-Income countries refers to World Bank client countries and include high-income IBRD clients.

Annex 4. Out-of-school children (enrollment)

Enrollment Data Selection

Enrollment is the share of children in a specific age group that is attending school. We use Adjusted Net Enrollment Rate, defined as the total number of students in the official primary school age group who are enrolled in primary or secondary education, expressed as a percentage of the corresponding population.

General rules

Initial Source: World Bank (WB Open Data), UIS (UNESCO Institute of Statistics), and other sources suggested by World Bank regional or country teams.

To fill in missing values: Step function used, and data filled in with the value of the closest year. If there are data available for two years equally close to the year to fill, the older value is used.

For some countries, despite filling the missing values, the values for adjusted net enrollment rates are still missing. To obtain enrollment rates, we follow this hierarchy:

1. Adjusted net enrollment rate (ANER)
2. Total net enrollment rate (TNER)
3. Gross enrollment rate (GER)
4. Gross enrollment rate capped at 100%: if the gross enrollment rate is higher than 100%, it is adjusted to be 100%.
5. Household Survey

Table 16 Source of Enrollment Data

Type of enrollment indicator	Freq.	Percent
Adjusted net enrollment rate (ANER)	108	93%
Country Team Validation	4	3%
GER (capped at 100%)	2	2%
Gross Enrollment Rate (GER)	1	1%
Household survey	1	1%
Total	116	100%

When we implement this hierarchy, gross net enrollment rate is used for Austria, Czech Republic and Slovak Republic. For Austria and Czech Republic, the enrollment values appear as 100 percent due to the cap on gross enrollment rates. For Afghanistan, we used national household survey data. For China, Côte d'Ivoire, Costa Rica, and Moldova, we used national sources based on the validation discussions with the relevant World Bank country teams.

Annex 5. Population numbers and projections

The population data used in our estimates can be found at HealthStats' Population Estimates and Projection database from the World Bank,[57] under the indicator codes SP.POP.1014.MA and SP.POP.1014.FE. These aggregates use a 5-year age group population calculated to be consistent with World Bank's total population data but follow the gender and age distribution the United Nations Populations Division (UNPD). About 30% of World Bank's total population data come from country sources, and about 70% of them come from UNPD.

The population projections from 2015 to 2030 follow the UNPD estimates and projections (2017 Revision), prepared by the Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat. The 2017 Revision is the twenty-fifth round of global population estimates and projections produced by the Population Division since 1951. The detailed description of the methodology for preparing country estimates, including the assumptions that were used to project fertility, mortality and international migration up to the year 2100, can be found on the website of the Population Division.[58] This paper uses the 10-14 age group UN Population Division's "medium-variant projection."[59]

The single-year age population numbers (for example, Age population, age 10, male, interpolated; Age population, age 10, female, interpolated; Age population, age 14, male, interpolated; Age population, age 14, female, interpolated) were taken directly from UNPD data (not as percentages, rather absolute numbers).

Table 17 Population age 10-14 by region and income classifications (Year = 2015)

	2015 Population age 10-14 (in millions)				
Region \ Income Level	High income	Upper-middle income	Low-middle income	Low income	Total
East Asia and Pacific	10	88	49	2	149
Europe and Central Asia	25	19	5	1	51
Latin America and Caribbean	2	47	3	1	54
Middle East and North Africa	4	15	13	5	37
North America	23	0	0	0	23
South Asia	0	2	165	8	175
Sub-Saharan Africa	0	6	53	64	123
Global	65	177	289	81	612

Source: Authors' calculations using UN Population numbers and World Bank income level classification FY2020.

[57] Available at https://databank.worldbank.org/data/source/health-nutrition-and-population-statistics:-population-estimates-and-projections#

[58] www.unpopulation.org. See also https://esa.un.org/unpd/wpp/Graphs/Probabilistic/

[59] For details, see https://population.un.org/wpp/Publications/Files/WPP2017_Methodology.pdf

Annex 6. Validation with other learning datasets

Table 18 Weighted average and correlation of PISA and learning poverty country averages according to country groupings and moving windows of PISA data (weighted)

	Temporal windows of PISA values in respect to LP reference year	Indicator (in percent)			Rho		N
		Lower-secondary LD (SDG 4.1.1c)	End-of-primary LD (SDG 4.1.1b)	LP	SDG 4.1.1c & SDG 4.1.1b	SDG 4.1.1c & LP	
All countries	3 to 5 years	52.1	42.2	44.3	0.87	0.90	123
	0 to 4 years	45.4	33.6	35.6	0.88	0.90	228
	-3 to -5 years	34.1	20.3	22.2	0.79	0.80	107
	-6 to -11 years	32.6	15.0	16.9	0.69	0.71	110
	-11 or lower	15.3	3.4	4.6	0.14	0.33	28
Low- and Middle-income countries	3 to 5 years	22.5	7.0	8.4	0.79	0.76	76
	0 to 4 years	23.0	8.3	9.6	0.89	0.86	160
	-3 to -5 years	19.8	5.0	6.6	0.76	0.76	79
	-6 to -11 years	20.0	5.9	7.2	0.79	0.81	85
	-11 or lower	14.9	3.5	4.5	0.27	0.42	26

Source: Authors' calculations using the Global Learning Assessment Database and UIS Enrollment Data.

Note: All measures are country-weighted means; a negative window implies that the PISA data was collected prior to the Learning Poverty (LP) and Learning Deprivation (LD) data.

Table 19 Correlation of early grade and end-of-primary scores

Assessment	Indicator	Early Grade[1]	End of Primary[2]	Rho[3]	N
LLECE	Mean Score	504.3	502.3	0.96	31
	LD, end of primary	504.3	56.7	-0.94	31
	LD-Gap, end of primary	504.3	45.0	-0.85	31
PASEC	Mean Score	500.0	500.0	0.48	10
	LD, end of primary	500.0	82.9	0.03	10
	LD-Gap, end of primary	500.0	53.1	-0.72	10
PASEC (except Burundi)	Mean Score	485.8	497.2	0.63	9
	LD, end of primary	485.8	81.8	-0.53	9
	LD-Gap, end of primary	485.8	54.8	-0.67	9

Source: Authors' calculations using the Global Learning Assessment Database and UIS Enrollment Data.

Note: (1) Early Grade learning measures expressed in mean scores; (2) End-of-Primary learning measures expressed in mean score, Learning Deprivation (LD), and Learning Deprivation Gap (LD-Gap); and (3) country-level pairwise correlation coefficient between Early Grade and End-of-Primary learning indicators. The expected correlation between mean scores is positive, given that both indicators have the same directional interpretation; the expected correlation between a mean score and a deprivation measure is negative, given the different directional interpretation of the respective measures.

Annex 7. Learning poverty levels and data sources, by country

Table 20 Country numbers

Region Code	Country Name	Schooling deprivation (SD, %)	Learning Deprivation (LD, %)	Learning Poverty (LP, %)	Assessment	Assessment Year
EAP	Australia	3.2	5.5	8.6	PIRLS	2016
EAP	Cambodia	2.6	49.8	51.1	NLA	2013
EAP	China	0.0	18.2	18.2	NLA	2016
EAP	Hong Kong SAR, China	1.9	1.4	3.2	PIRLS	2016
EAP	Indonesia	2.4	33.8	35.4	PIRLS	2011
EAP	Japan	1.2	1.0	2.2	TIMSS	2015
EAP	Korea, Rep.	2.7	0.3	3.0	TIMSS	2015
EAP	Macao SAR, China	1.3	2.4	3.7	PIRLS	2016
EAP	Malaysia	1.4	11.7	12.9	NLA	2017
EAP	New Zealand	1.5	10.0	11.4	PIRLS	2016
EAP	Singapore	0.1	2.7	2.8	PIRLS	2016
EAP	Thailand	2.0	21.9	23.5	TIMSS	2011
EAP	Vietnam	0.6	1.1	1.7	NLA	2011
ECA	Armenia	7.2	30.0	35.0	TIMSS	2015
ECA	Austria	0.0	2.4	2.4	PIRLS	2016
ECA	Azerbaijan	5.0	19.2	23.3	PIRLS	2016
ECA	Belgium	1.3	5.1	6.4	PIRLS	2016
ECA	Bulgaria	6.8	5.2	11.7	PIRLS	2016
ECA	Croatia	3.0	1.0	4.0	PIRLS	2011
ECA	Cyprus	2.2	14.3	16.2	TIMSS	2015
ECA	Czech Republic	0.0	3.0	3.0	PIRLS	2016
ECA	Denmark	1.0	2.6	3.6	PIRLS	2016
ECA	Finland	0.9	1.7	2.6	PIRLS	2016
ECA	France	0.9	6.3	7.1	PIRLS	2016
ECA	Georgia	0.4	13.5	13.8	PIRLS	2016
ECA	Germany	0.2	5.5	5.7	PIRLS	2016
ECA	Hungary	3.1	2.9	5.9	PIRLS	2016
ECA	Ireland	0.0	2.3	2.3	PIRLS	2016
ECA	Italy	1.4	2.1	3.5	PIRLS	2016
ECA	Kazakhstan	0.3	1.9	2.2	PIRLS	2016
ECA	Kyrgyz Republic	1.9	63.8	64.5	NLA	2014
ECA	Latvia	3.2	0.8	4.0	PIRLS	2016
ECA	Lithuania	0.3	2.7	3.0	PIRLS	2016
ECA	Netherlands	0.3	1.3	1.6	PIRLS	2016
ECA	Norway	0.2	5.8	6.0	PIRLS	2016
ECA	Poland	4.4	2.0	6.3	PIRLS	2016
ECA	Portugal	3.6	3.0	6.5	PIRLS	2016
ECA	Romania	6.9	14.1	20.0	PIRLS	2011
ECA	Russian Federation	2.4	0.9	3.3	PIRLS	2016

ECA	Serbia	0.8	7.4	8.1	TIMSS	2015
ECA	Slovak Republic	2.1	6.6	8.5	PIRLS	2016
ECA	Slovenia	2.2	3.7	5.8	PIRLS	2016
ECA	Spain	1.5	3.4	4.9	PIRLS	2016
ECA	Sweden	0.4	1.9	2.3	PIRLS	2016
ECA	Turkey	5.0	17.6	21.7	TIMSS	2015
ECA	United Kingdom	0.2	3.2	3.4	PIRLS	2016
LAC	Argentina	0.6	53.6	53.9	LLECE	2013
LAC	Brazil	2.7	46.9	48.4	LLECE	2013
LAC	Chile	9.3	30.3	36.8	LLECE	2013
LAC	Colombia	6.9	44.7	48.6	LLECE	2013
LAC	Costa Rica	1.1	31.7	32.5	LLECE	2013
LAC	Dominican Republic	6.6	79.4	80.7	LLECE	2013
LAC	Ecuador	1.9	62.1	62.8	LLECE	2013
LAC	Guatemala	10.1	63.6	67.3	LLECE	2013
LAC	Honduras	17.1	69.4	74.7	LLECE	2013
LAC	Mexico	1.2	42.5	43.2	LLECE	2013
LAC	Nicaragua	1.6	69.3	69.8	LLECE	2013
LAC	Panama	7.1	64.1	66.6	LLECE	2013
LAC	Paraguay	10.8	71.3	74.4	LLECE	2013
LAC	Peru	4.2	53.7	55.7	LLECE	2013
LAC	Trinidad and Tobago	1.3	19.7	20.7	PIRLS	2016
LAC	Uruguay	0.5	41.4	41.7	LLECE	2013
MNA	Bahrain	2.1	30.6	32.1	PIRLS	2016
MNA	Egypt, Arab Rep.	1.4	69.2	69.6	PIRLS	2016
MNA	Iran, Islamic Rep.	0.9	35.1	35.7	PIRLS	2016
MNA	Israel	2.9	9.0	11.7	PIRLS	2016
MNA	Jordan	4.0	50.0	52.0	TIMSS	2015
MNA	Kuwait	3.3	49.4	51.0	PIRLS	2016
MNA	Malta	2.4	26.8	28.6	PIRLS	2016
MNA	Morocco	5.4	63.8	65.8	PIRLS	2016
MNA	Oman	1.5	40.9	41.8	PIRLS	2016
MNA	Qatar	2.2	33.8	35.3	PIRLS	2016
MNA	Saudi Arabia	2.5	36.7	38.3	PIRLS	2016
MNA	Tunisia	0.4	65.1	65.3	TIMSS	2011
MNA	United Arab Emirates	2.8	32.4	34.3	PIRLS	2016
MNA	Yemen, Rep.	18.9	93.5	94.7	TIMSS	2011
NAC	Canada	0.0	4.3	4.3	PIRLS	2016
NAC	United States	4.1	3.9	7.9	PIRLS	2016
SAR	Afghanistan	49.6	87.0	93.4	NLA	2013
SAR	Bangladesh	4.9	55.0	57.2	NLA	2015
SAR	India	2.3	53.7	54.8	NLA	2017
SAR	Pakistan	27.3	65.0	74.5	NLA	2014
SAR	Sri Lanka	0.9	14.0	14.8	NLA	2015

SSA	Benin	3.6	77.3	78.2	PASEC	2014
SSA	Botswana	7.2	44.3	48.3	PIRLS	2011
SSA	Burkina Faso	31.7	78.6	85.4	PASEC	2014
SSA	Burundi	2.7	92.7	92.9	PASEC	2014
SSA	Cameroon	5.2	75.9	77.2	PASEC	2014
SSA	Chad	21.1	97.0	97.7	PASEC	2014
SSA	Congo, Dem. Rep.	63.2	62.0	86.0	PASEC	2010
SSA	Congo, Rep.	12.8	82.9	85.1	PASEC	2014
SSA	Cote d'Ivoire	21.1	77.6	82.3	PASEC	2014
SSA	Ethiopia	14.0	88.7	90.3	NLA	2015
SSA	Madagascar	21.9	95.8	96.7	PASEC	2015
SSA	Mali	33.0	85.7	90.5	PASEC	2012
SSA	Niger	38.9	97.9	98.7	PASEC	2014
SSA	Senegal	25.7	65.2	74.1	PASEC	2014
SSA	South Africa	8.4	77.9	79.8	PIRLS	2016
SSA	Togo	8.5	84.2	85.6	PASEC	2014
SSA	Uganda	9.0	81.1	82.8	NLA	2014

Source: Authors' calculations using the Global Learning Assessment Database; and UIS Enrollment Data.
Note: Countries ordered alphabetically within regions; Only the latest assessment data since 2011 is reported for each country, with the exception of Congo, Dem. Rep. which is from 2010.

Annex 8. Understanding the spells database

An important requirement for the simulations is to have good measurements of how reading proficiency improved over the 2000-17 period. Calculating this rate of progress is even more challenging than estimating levels, because of the lack of data and thresholds that are comparable over time. For any individual country, there may be multiple estimates of proficiency rates available for the past 17 years, but simply calculating progress using those estimates would be misleading. Often, those estimates are based on data from different assessments—PIRLS and PASEC, for example. Although the equating process described above aims to harmonize the proficiency levels across assessments to allow calculation of proficiency levels, this process is too imprecise to be used for estimates of change over time. Given that reductions in learning poverty are typically much smaller than the levels (for example, 1% annual improvement on a baseline of 50%), the noise introduced by mixing assessments is likely to swamp the signal (the actual change in proficiency).

Drawing on our global database of learning poverty, which includes over 900 subject-specific national measures of learning poverty, we construct a database of rates of change for each country only *within* grade, subject, and assessments. For example, we calculate the improvement from PIRLS 2011 to PIRLS 2016, but do not estimate the change from PIRLS 2011 to LLECE 2013. We further restricted this database to include only grades 4, 5, and 6 and only assessments of reading and science.[60] It is important to note that 32 TIMMS and PIRLS spells have been removed due to lack of comparability over time.[61] In addition, because PASEC is not comparable over time (although it is comparable across countries within a cycle), we do not use the data from PASEC. Although we do not use SACMEQ for estimating global proficiency levels, given quality concerns about the latest round of SACMEQ, we do use it in estimating changes over time, since earlier SACMEQ rounds are of sufficient quality. Lastly, we restrict our sample to spells within the range of -4 to +4 annualized change in learning poverty,[62] to use measures that are robust to the presence of outliers in our simulation.

With these restrictions applied, our dataset yields 207 spells of change for 91 countries between 2000 and 2017. Table 21 shows the summary statistics of the spell database by assessment. For this paper, we focus on the subset of observations from low- and middle-income countries. These additional filters restrict our database further to a final number of 72 spells for 50 countries.

---

[60] Jordan was the only country in which TIMSS math scores were used as a proxy for reading, since neither reading nor TIMSS science scores are available.

[61] Compiled from multiple PIRLS and TIMSS assessment documentation. Changes in the assessments meant that within particular rounds the assessments were comparable across countries in the round, but not between rounds.

[62] The spells outside the considered range and their respective values are SACMEQ 2007-2013 grade 6 for Zambia (4.3), Kenya (4.6), Botswana (5.1), South Africa (5.9), Uganda (6.0), Malawi (6.3), Namibia (7.4) and Lesotho (7.5); TIMSS 2007-2011 grade 4 for Georgia (4.3); and TIMSS 2011-2015 grade 4 for Morocco (4.5).

Table 21 Summary statistics of the annualized changes in learning poverty and initial condition by assessment

Assessment		Annualized Change (p.p.)					Initial conditions				
		mean	p50	min	max	N	mean	p50	min	max	N
All Countries	LLECE	0.77	0.62	-0.43	2.14	14	61.3	61.9	29.5	91.1	14
	NLA	-0.48	-0.48	-0.48	-0.48	1	57.2	57.2	57.2	57.2	1
	PIRLS	0.26	0.09	-3.96	3.34	95	14.0	7.1	0.8	80.4	95
	SACMEQ	1.48	1.10	-0.59	3.88	17	62.0	59.7	35.3	94.5	17
	TIMSS	0.37	0.29	-3.35	3.10	80	16.0	8.7	1.6	94.1	80
	Total	0.43	0.21	-3.96	3.88	207	22.1	9.1	0.8	94.5	207
Low- and Middle-Income Countries*	LLECE	0.77	0.62	-0.43	2.14	14	61.3	61.9	29.5	91.1	14
	NLA	-0.48	-0.48	-0.48	-0.48	1	57.2	57.2	57.2	57.2	1
	PIRLS	0.81	0.57	-2.25	3.34	23	28.2	23.0	3.7	80.4	23
	SACMEQ	1.48	1.10	-0.59	3.88	17	62.0	59.7	35.3	94.5	17
	TIMSS	0.71	0.76	-2.73	3.10	17	35.9	27.7	4.8	94.1	17
	Total	0.92	0.66	-2.73	3.88	72	44.8	41.1	3.7	94.5	72

Source: Authors' calculations using the Global Learning Assessment Database and UIS Enrollment Data.
Note: "Initial conditions" refers to the level of learning poverty at the beginning of each spell. (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients.

Table 22 shows the summary statistics of the spell database by region (excluding Part 1 countries), both unweighted and weighted. In the latter, senate weights are used to ensure an equal contribution by each of the countries in the sample of spells, for some countries have multiple spells. A few results stand out. First, there is not enough coverage of sufficiently comparable cross-national assessment in two regions, EAP and SAR. This is striking, given the population density of these regions. Second, for most regions, the distribution of spells is relatively well behaved. It is also important to highlight the range of initial conditions that our spell database covers: poverty rates at the start of the spells range from very low (4%) to very high (91%). In designing the simulations, it will be important to account for both of these features of the data.

Table 22 Summary statistics of the annualized changes in learning poverty and initial condition by region, low- and middle-income countries (weighted and unweighted)

	Assessment	Annualized Change (p.p.) mean	p50	min	max	N	Initial conditions mean	p50	min	max	N
Unweighted	East Asia and Pacific					1					1
	Europe and Central Asia	0.62	0.57	-0.86	3.10	23	16.9	13.4	3.7	43.9	23
	Latin American and Caribbean	0.88	0.62	-0.43	3.34	18	54.6	52.8	22.2	91.1	18
	Middle East and North Africa	0.73	1.03	-2.73	3.30	12	58.1	60.6	24.4	94.1	12
	South Asia					1					1
	Sub-Saharan Africa	1.48	1.10	-0.59	3.88	17	62.0	59.7	35.3	94.5	17
	Total	0.92	0.66	-2.73	3.88	72	44.8	41.1	3.7	94.5	72
Weighted	East Asia and Pacific					1					1
	Europe and Central Asia	0.74	0.68	-0.86	3.10	23	20.1	17.3	3.7	43.9	23
	Latin American and Caribbean	0.84	0.62	-0.43	3.34	18	58.2	55.5	22.2	91.1	18
	Middle East and North Africa	0.57	0.36	-2.73	3.30	12	70.2	73.5	24.4	94.1	12
	South Asia					1					1
	Sub-Saharan Africa	1.29	0.92	-0.59	3.88	17	65.9	66.7	35.3	94.5	17
	Total	0.93	0.64	-2.73	3.88	72	50.9	52.0	3.7	94.5	72

Source: Authors' calculations using the Global Learning Assessment Database and UIS Enrollment Data
Note: "Initial conditions" refers to the level of learning poverty at the beginning of each spell. In the cases of EAP and SAR (which have only 1 spell each), their single spells were used in the global distribution of spells, but they were not used as representative of that particular region. The weighted panel uses senate weights to ensure an equal contribution by each of the countries in the sample of spells.

Simulation robustness checks

The main simulation results, presented in Table 12, used spells by region as a basis for projected growth. As a robustness check, we present analogous simulation results using spells by income level (Table 23) and by initial learning poverty rate (Table 24).

Table 23 Learning poverty rates in 2030 under two scenarios (alternative simulation using spells by income level)

Region	Population (millions) 2015	2030	Learning Poverty (%) 2015 Base	2030 BaU	2030 High	Share of Global Learning-Poor (%) 2015 Base	2030 BaU	2030 High
East Asia and Pacific	137	138	21.2	4.4	0.2	10.0	2.3	0.1
Europe and Central Asia	27	32	13.3	3.2	0.8	1.2	0.4	0.1
Latin American and Caribbean	53	51	50.8	39.4	22.8	9.4	7.4	6.7
Middle East and North Africa	33	44	63.3	51.2	29.4	7.1	8.3	7.5
South Asia	175	167	58.2	45.4	19.5	35.2	28.0	18.9
Sub-Saharan Africa	123	171	86.7	85.1	66.9	37.0	53.7	66.6
Overall	548	602	52.7	45.0	28.5	100.0	100.0	100.0

Note: Business as Usual (BaU) scenario uses country-specific growth rates; for countries with no spells, the income level average was used. High scenario considers the rate of improvement equal to the group's 80th percentile. All figures consider only Part 2 countries.

Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.

Table 24 Learning poverty rates in 2030 under two scenarios (alternative simulation using spells by initial learning poverty)

	Population (millions)		Learning Poverty (%)			Share of Learning Poor (%)		
Region	2015	2030	2015	2030	2030	2015	2030	2030
			Base	BaU	High	Base	BaU	High
East Asia and Pacific	137	138	21.2	7.5	2.4	10.0	3.9	1.9
Europe and Central Asia	27	32	13.3	3.4	1.8	1.2	0.4	0.3
Latin American and Caribbean	53	51	50.8	39.6	19.7	9.4	7.4	5.8
Middle East and North Africa	33	44	63.3	54.4	33.4	7.1	8.8	8.6
South Asia	175	167	58.2	51.2	26.4	35.2	31.8	25.8
Sub-Saharan Africa	123	171	86.7	75.2	57.4	37.0	47.7	57.5
Overall	548	602	52.7	44.8	28.3	100.0	100.0	100.0

Note: Business as Usual (BaU) scenario uses country-specific growth rates; for countries with no spells, the average of countries within the same initial learning poverty level was used. High scenario considers the rate of improvement equal to the group's 80th percentile. All figures consider only Part 2 countries in the World Bank classification.
Source: Authors' calculations using the Global Learning Assessment Database; UIS Enrollment Data; and UN Population numbers.

Annex 9. Calculation of standard errors

At the country level, learning poverty is a weighted sum of two indicators: the share of children in school who are below a minimum proficiency threshold (learning deprivation) and the share of out-of-school children (schooling deprivation). The first indicator is estimated using sample-based learning assessments, while the latter is estimated in virtually all cases using administrative records from national Education Management Information System and the population census. Both measures have an associated error term, but they are of a very different nature.

Our learning data is sample-based, and as such its error term reflects the sampling error and the psychometric procedure used to estimate the latent learning variable (for more details, see Jerrim et al 2017 on inference of learning data, taking into consideration both its complex survey design and the use of plausible values). By contrast, our out-of-school measure is a population measure estimated using administrative records and census data, and therefore does not have a sampling error. Both measures, however, are also affected by non-sampling error, such as questionnaire or measurement error, implementation challenges, and behavior effects. Unfortunately, we have no way to capture this non-sampling error; hence we use bootstrap for error propagation of the sampling error associated with our learning measure.

Using our student-level assessment microdata, we form an indicator for whether each student is above the minimum proficiency level defined above and estimate the mean to produce the proportion proficient in each country, along with the standard error of that mean estimate for each country-year-assessment combination. Applying the Central Limit Theorem, which can be justified because the assessment databases typically contain several hundred student observations, our estimator of the proportion above minimum proficiency in each country follows an asymptotically Normal distribution. To produce standard errors for our final numbers, which are based on these country-year-assessment-level proficiency numbers, we take 100 bootstrap random draws of our country-year-assessment-level proficiency database; for these draws, each individual observation in our database is drawn from the Normal distribution whose mean is equal to our estimate of the proportion minimally proficient and whose variance is the squared standard error of that estimated proportion. Then our final global and regional numbers are calculated in each of these 100 bootstrap simulated databases, and our standard error is the standard deviation of our estimate across those 100 bootstrap datasets.

As discussed, our enrollment numbers, which also feed into our learning poverty measure, are population measures based on administrative records without an associated sampling error. We acknowledge that this indicator might suffer from non-sampling errors, which can lead to inaccurate counts of students enrolled in schools. Non-sampling errors affect both of our measures, as they have a direct impact on the out-of-school measure, and an indirect effect in our learning estimates as they impact the sample frame. These types of misreporting errors are difficult to account for, and we are unable to incorporate this type of error into our standard error calculations. We acknowledge this as a limitation.

Special Cases for Standard Error Calculations

In some cases, where we are using country-year observations based on national assessments, we do not have a standard error associated with this observation. When no standard error is available, we use a value of 1.2%, which is approximately the median standard error across all country-year-assessment combinations used.

Annex 10. The need to improve learning data for learning poverty estimates

In statistics, quality used to be associated primarily with accuracy. It is now recognized that there are other important dimensions in data quality. Even if data are accurate, they do not have sufficient quality if they are produced too late to be useful, or cannot be easily accessed, or conflict with other credible data. Therefore, quality is increasingly approached as a multi-dimensional concept.[63]

In the main text, Table 2 and the accompanying section describe the substantial data quality gaps. Filling these gaps is an important step toward improving the overall quality of learning poverty statistics, and require tackling significant problems in terms of data availability (existence), relevance, accessibility, timeliness, and comparability both across countries and over time.

A first critical aspect is the lack of availability or existence of learning data and low- and middle-income countries with a sufficient periodicity, or even temporal regularity. Several low-income countries have never measured learning (or not measured it in the last 10 years), as their participation in cross-national programs is in some cases irregular, or even non-existent. The rigidities imposed by cross-national assessments also do not help. Going forward, it will be critical to continue to promote the production of new learning assessments and to increase the flexibility of existing programs.

To be policy-relevant, learning data must meet users' needs. Relevance requires the identification of users and their expectations. It is often the case that users of learning data are too narrowly defined, and as a consequence, many users and policy questions are not recognized and end up not being served.

Another critical data quality aspect is accessibility. Data have the most value when they are easily accessible by users, are available in the form users desire, and are adequately documented (with "metadata" appropriate for the type of user). Currently, 12 countries are represented in the learning-poverty dataset through their respective National Learning Assessments (NLAs). Many other countries have national assessments, and the number will likely increase as countries strengthen their national systems and take greater ownership of the learning assessment process. But unlike data from cross-national assessments, the micro- and meta-data from these national assessments is often not accessible; instead, at best only aggregate results are reported, often without sufficient breakdowns and lack of sufficient documentation. This will need to change if we are to include as many countries as possible in the learning poverty estimates.

A fourth critical aspect, beyond availability, relevance and accessibility, is timeliness**.** Currently most cross-national learning assessments take two to three years to go from data collection to dissemination. By contrast, in the case of surveys of monetary poverty and labor-force surveys, several national offices of statistics are now able to collect and disseminate data within just 3 to 6 months. Agile, timely production of learning data should therefore be a priority.

A fifth issue is grade comparability across countries and more importantly within countries. While most rich countries assess their learners at Grade 4, low- and middle-income countries tend to be less consistent and evaluate more of their learners at a later grade, such as 5 or 6 (see Table 25). Moreover, some countries participate in different cross-national assessments; this gives them information on

[63] The International Organization for Standardization (ISO 8402) defines quality as follows: “Quality is the totality of features or characteristics of a product or service that bear on its ability to satisfy stated or implied needs of customers.” For a summary of how different international organizations define dimensions of statistical quality, please see Vries (2002).

learning outcomes at different grades at each point in time, but it prevents them from using these for comparisons and for tracking progress to inform policies and programs. For example, Chile and Colombia participated in LLECE (which assesses learners in Grade 6) in 2013 and in PIRLS (Grade 4) in 2016, but these assessments cannot be used to provide information on changes in learning between 2013 and 2016. In 2011, Honduras applied the PIRLS assessment, but did it by sampling Grade 6 students. As consequence, these results are noncomparable to those of other PIRLS countries, which applied the assessment in Grade 4. And because the PIRLS questions were different, results were also inconsistent with Honduras' LLECE Grade 6 assessment from 2006 and 2013.

Table 25 Assessment comparability in terms of grade assessed, by income level

	Number of Countries by Grade			Population Distribution by Grade (%)		
Grade	High-Income Countries	Low- and Middle-Income Countries	All Countries	High-Income Countries	Low- and Middle-Income Countries	All Countries
4	37	21	58	94	38	54
5	4	1	5	6	13	11
6	0	28	28	0	49	35
Total	41	50	91	100	100	100

Source: Authors' calculations using the Global Learning Assessment Database; and UN Population numbers.
Notes: Due to lack of comparability over time, based on documentation by International Association for the Evaluation of Educational Achievement (IEA), 32 TIMMS and PIRLS spells have been removed. PASEC is not comparable over time, although it is comparable across countries within a cycle. For LLECE, the TERCE-SERCE scale is used to ensure comparability; otherwise, those spells would also have to be removed.

The lack of temporal comparability of assessments and scales over time within many countries makes it impossible for those countries to monitor progress on student learning. Even between rounds of a given assessment, there is often a challenge of comparability. Using the cross-national assessment data from the past 20 years, we are able to track 219 episodes of measured changes in learning for end-of-primary students. But most of these are for high-income countries; for low- and middle-income countries, the number is only 72 episodes, after excluding outliers (see Table 26). This limitation reflects the sparsity of the data, as assessment cycles follow five- to seven-year intervals or have been entirely irregular. Moreover, some cross-national assessment programs make significant changes in their scales between rounds or have even adopted instruments designed only for cross-national comparison within rounds. This challenge is even worse in the context of National Learning Assessments, which are designed to reflect national curricula that are often subject to change.

Table 26 Temporal comparability within assessments

	All Countries		Low- and Middle-Income Countries*	
Assessment	all spells	without outliers	all spells	without outliers
LLECE	14	14	14	14
NLA	1	1	1	1
PASEC	0	0	0	0
PIRLS	95	95	23	23
SACMEQ	25	17	25	17
TIMSS	83	80	19	17
Total	218	207	82	72

Source: Authors' calculations using the Global Learning Assessment Database.
Notes: Due to lack of comparability over time, based on existing documentation by International Association for the Evaluation of Educational Achievement (IEA), 32 TIMMS and PIRLS spells have been removed. PASEC is not comparable over time, although it is comparable across countries within a cycle. Spells indicating decrease in learning poverty greater than 4 percentage points per year or increase greater than 4 percentage points per year were considered outliers. For LLECE, the TERCE-SERCE scale is used to ensure comparability; otherwise, those spells would also have to be removed. (*) Low- and Middle-Income countries refers to Part 2 countries, which are eligible to borrow from the World Bank Group and include high-income IBRD clients. See Table 3 for the number of countries per classification.

While this discussion has focused on the end-of-primary reading assessments required for our learning poverty measure, these data issues are often magnified in other cross-national and national assessments that assess different grades and subjects. Data on enrollment and administrative records also need improvement. In some cases, there are substantial discrepancies between administrative records and household surveys, and these differences are poorly understood.

The picture that emerges from this brief analysis is one of a highly fragmented system of education data, with significant variance across regions over time and within countries, in terms of coverage, comparability, and frequency—corroborating the findings from UIS 2019, which points to similar weaknesses. It will be important to make sure that learning data and administrative records from line ministries such as education are included in the minimum data package of any state-building or capacity-strengthening initiative, with particular effort to strengthen country data systems in Sub-Saharan African and fragile and conflict-affected states. To fill these data gaps and make better use of available data will require innovative implementation modalities and better coordination among development partners.

Annex 11. Scientific reproducibility and GitHub repository

To ensure the full replication and documentation of this work, all code required to reproduce the numbers of this paper are stored in two public World Bank GitHub repositories. The Global Learning Assessment Database (GLAD) Repo assembles the microdata harmonization of all learning assessments used in this paper, and the Learning Poverty Repo documents the construction of the learning poverty indicator, the spells database, the simulations, and all tables used in this paper.

Figure 9 GitHub – Learning Poverty

https://github.com/worldbank/LearningPoverty

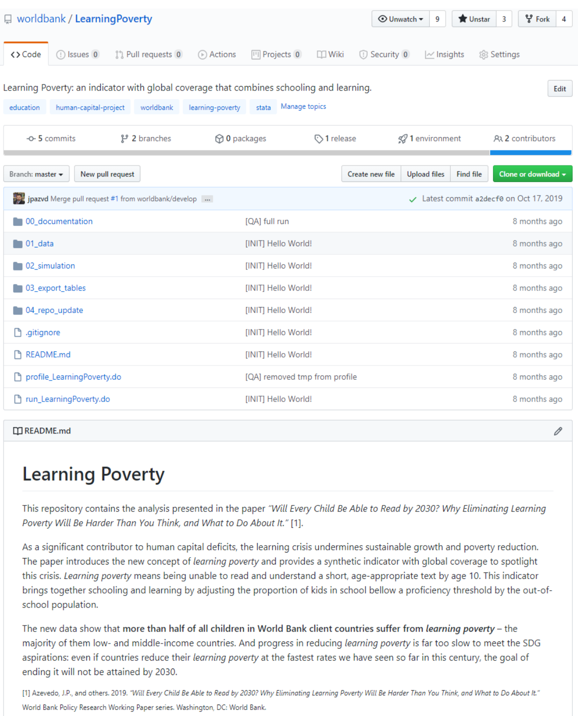

Figure 10 GitHub – Global Learning Assessment Database (GLAD)

https://github.com/worldbank/GLAD