

Measuring Child Labor: the Who's, the Where's, the When's, and the Why's

Guilherme Lichand^{*†} and Sharon Wolf[§]

April 14, 2025

Abstract: Measuring child labor accurately is a major challenge: parents' and children's reports tend to differ dramatically, and there is typically no way to verify whose reports are truthful (if any). To overcome this challenge, this paper uses novel data from a cocoa certifier in Côte d'Ivoire that draws on satellite imagery to minimize under-reporting. Concretely, aerial photos allow them to select remote and hard-to-reach communities – where parents typically have not been sensitized by government or NGOs, averting social desirability biases – and to visit these communities while cocoa is being harvested – precisely when children in employment are very visible, making it easier for enumerators to impute it if parents still fail to report it. We compare their figures with those obtained from business-as-usual surveys with parents and children in these regions, and find that (1) reporting inconsistencies between parents and their children in fact decrease with household remoteness; (2) adults dramatically under-report child labor relative to the certifier data, by a factor of at least 60%; and (3) in turn, children self-reports are statistically identical to the latter. Taking advantage of an experiment that randomly assigned a text-message campaign to discourage child labor, we further show that parents' reports not only underestimate its prevalence, but can even lead to the wrong conclusions about the effects of policy interventions.

^{*} Graduate School of Education, Stanford University; glichand@stanford.edu.

[§] Graduate School of Education, University of Pennsylvania.

[†] Both authors contributed equally to this work.

1 Introduction

Child labor is a pervasive practice in agriculture, especially in West Africa, where the global cocoa industry sources roughly half of its produce. Defined according to international standards and national legislation, child labor typically characterizes any work by children under the age of 12, and excessive or risky work by 12–17 year-olds (see Section 2.1.1). 2020 data from the International Labor Organization (ILO) documented 160 million child workers worldwide – 9.6% of all 5–17 year-old children [1]. Strikingly, that figure was over 2-fold in Sub-Saharan Africa (23.9%). In Côte d’Ivoire, the setting of our study, 15% of the cocoa industry employees that year were children [2].

While such figures have mobilized massive global attention and resources to combat child labor, a key challenge in interpreting them is reliability. This is because while official statistics on child labor are essentially based on adult reports (see Section 2.1.1), there are many reasons why parents or caregivers may fail to truthfully disclose children in employment. They might fear that admitting to child labor (e.g., that their under-12 year-old children help during the harvest season) could lead them to face legal action from child protective services, or adversely impact their livelihoods (e.g., if the companies sourcing from them are punished for child labor in their supply chain, which could trickle down to lower prices or even to discontinuation of their income; see [3]). Importantly, if parents omit children in employment, child labor will be under-estimated in the data – since child labor is *contained* in children in employment.

This paper uses novel data from a cocoa certifier in Côte d’Ivoire to overcome this challenge. ENVERITAS, a global NGO that certifies smallholder coffee and cocoa farms when it comes to agricultural best practices – including the absence of child labor –, draws on satellite imagery to minimize under-reporting. Concretely, aerial photos allow them (1) to select remote and hard-to-reach communities, where parents typically have not been sensitized by government or NGO campaigns to discourage child labor, averting social desirability biases; and (2) to visit these communities during harvest season, precisely when child labor is very visible, making it easier for enumerators to mark it as present even if parents still fail to report it. We compare their figures with those obtained from business-as-usual surveys with parents and children in these regions, which partially overlap with the remote villages where ENVERITAS certifies farming practices. Our sample focuses on primary school children, nearly all under 12 years old – for whom any hours in employment are classified as child labor by national legislation.

Before turning to the certifier data, we start by documenting patterns in reporting discrepancies in our own surveys that lend credibility to our empirical strategy. Restricting attention to households in our sample with valid GPS location data, 29.3% of adults admit that their children worked at least an hour in cocoa fields in the month before the survey. In contrast, 51.8% of children report that they did so – a striking 77% average difference in reporting rates. First, we document that such discrepancies indeed decrease with the distance from each household to the school their children attend: the difference between children’s and parents’ reports decrease significantly with every additional km from the school, consistent with ENVERITAS’s claim that village remoteness is inversely related to reporting biases. Second, we estimate that children’s self-reports do *not* increase with distance from the school, minimizing concerns with child labor in more remote villages misrepresenting its typical prevalence. We further document that not only are discrepancies between children’s and parents’ reports much lower when it comes to domestic work – a more socially acceptable activity and typically not the subject of advocacy campaigns, in contrast to labor in cocoa fields –, but also, that such discrepancies do *not* decrease with

household remoteness. All in all, results back up the certifiers’ sampling strategy as a way of measuring the ground truth for the prevalence of child labor, which we can then use to validate parents’ and children’s reports in the sub-sample for which the two datasets overlap.

Our main contribution is to document first-hand that adult surveys indeed under-report the prevalence of child labor, and the extent of under-reporting. In regions with subsequent third-party verification, 45.5% of children reported having worked in cocoa plantations in the previous month, matching almost exactly the 47.8% prevalence indicated by the certifier; in contrast, only 16.2% of parents in those regions reported children in employment – a nearly 2/3 reporting gap. Across regions, under-reporting ranged from 60% to 85%.

Previous studies were not able to document the extent of under-reporting by parents. These papers rely on three types of comparisons. First, comparisons between parents’ reports under different conditions for social desirability bias [e.g., 4], to document that such biases exist – but unable to pin down their magnitude. Second, comparisons between adults’ and children’s reports [5, 6], to document that discrepancies exist – but unable to pin down what the ‘ground truth’ is. Third, comparisons between an objective measure of children in employment and children’s self-reports [7], to document that the latter are accurate – but unable to pin down under-reporting by parents. Similarly to ours, the latter study documents that children accurately report the number of hours they work, using logs from GPS trackers worn by different household members to verify data from surveys and activity diaries. The study, however, actively refrained from having enumerators ask adults about child labor. In the absence of parents’ reports, it could not document whether parents under-report children in employment and, if so, by how much.

Our second contribution is to document that basing child labor accounts on surveys with parents not only underestimates its prevalence, but can also bias evaluations of how it responds to policy interventions. Taking advantage of a randomized control trial that assigned some Ivorian parents to messages discouraging child labor in cocoa fields to study whether the estimated impact of the intervention on child labor depends on how the latter is measured, we find that while messages had no effect on children in employment according to children themselves, they significantly increased children in employment according to parents (by 55.1%). Presumably, the reason for such discrepancy is that the intervention tried to foster investments in children that would reduce children’s participation in labor activities *without explicitly condemning child labor* – partially deterring social desirability biases. Once again consistent with the claim that such biases are stronger in communities that have been previously sensitized, treatment effects on child labor according to parents sharply decay with our measure of remoteness.

These findings have key implications for how child labor should be measured. One possibility is to survey children directly (the ‘*who*’). Our results indicate that this would most likely yield accurate estimates of children in employment, but it might also involve complex technical and ethical dilemmas. It might be hard to ensure a unified understanding of what characterizes employment, especially among younger respondents. Most importantly, participation might put children at risk if it triggers backlash by parents. An alternative is to survey adults, as usual, but to focus on hard-to-reach communities (the ‘*where*’) during harvest season (the ‘*when*’) – leveraging technological advances such as satellite imagery. We further discuss the implications and limitations of our findings in Section 4.

2 Materials and methods

2.1 How child labor is measured

There are striking differences between children’s and adults’ reports in the few settings where both have been asked independently about children in employment. According to NORC (a research institute based at the University of Chicago), which surveys children directly about the number of hours worked in cocoa fields, 38% of 5-17 year olds in Côte d’Ivoire reported to have worked in 2018-19 [2]; in contrast, the ILO figure for 2016 – based on adult reports for children employment – was only 23% [1]. Even worse, different sources tell very different stories about child labor’s recent trends. While ILO data indicate a 38% decrease in child labor worldwide since 2000 [8], NORC data record a nearly 65% *increase* in child labor since 2008-09 [2].

While these differences are suggestive that official statistics based on adult reports are biased, it is hard to be sure; besides differences in whom is surveyed, discrepancies across surveys might also accrue to differences in their geographical coverage or in the timing of data collection, or to other methodological differences. Moreover, in the absence of verification, it is unclear whether children’s self-reports do not suffer from reporting biases too.

2.1.1 ILO methodology

The International Labour Organization (ILO) follows the Convention on the Rights of the Child, the ILO Minimum Age for Admission to Employment Convention (No. 138), and the ILO Worst Forms of Child Labour Convention (No. 182) to define child labor [1]. According to these conventions, whether children in employment characterizes child labor depends on the child’s age, the number of hours dedicated to work, and the work conditions. For children 11 years old or younger, any employment characterizes child labor. For those between 12 and 14 years old, 15 or more weekly work hours and/or hazardous work conditions (for any number of hours) characterize child labor. Last, for those between 15 and 17 years old, child labor applies in case of 43 or more weekly work hours and/or hazardous work conditions. In Supporting Information, we illustrate all conditions used by the ILO to define child labor for children of different age groups. Importantly, the ILO encourages countries to modify and complement these guidelines with national legislation that makes figures contextually appropriate. In Brazil, for example, local laws establish that any employment characterizes child labor for children 14 years old or younger – a stricter standard relative to the ILO guidelines.

Statistics on children in employment, number of hours worked, and work conditions come from different surveys around the globe. The ILO does not collect the data itself but, rather, harmonizes data from these different sources to compute the prevalence of child labor according to its methodology. All leading international organizations currently follow the ILO methodology. For example, UNICEF adapted its Multiple Indicator Cluster surveys after 2013 to match ILO guidelines. The World Bank also tracks child labor following the same methodology, only for 7-14 year-olds rather than 5-17 year-olds.

In the ILO methodology, child labor is computed by *combining* surveys with adults and those with children. Concretely, the ILO methodology *only uses the adult questionnaire* to compute children in employment. In turn, the children questionnaire is used to assess hazardous work conditions *only among those who are employed* (“[a]s in the previous rounds, the current round of the Global Estimates of Child Labor uses data obtained from the adult questionnaire, except for conditions of work, where the information from the child

questionnaire is deemed to be more reliable”; [9], p. 59). In other words, child labor is a subset of children in employment – which is exclusively based on parents’ reports. As such, official statistics on child labor depend crucially on the accuracy of the latter.

2.1.2 NORC methodology

NORC, a research institution at the University of Chicago, has tracked child labor in the cocoa industry for Ghana and Côte d’Ivoire since 2015, building on Tulane University’s work in the region dating back to 2008. It reports statistics associated with child labor, for cocoa production in particular and for agricultural activities more broadly [2]. NORC defines child labor based on the number of work hours and work conditions for children of different age groups, consistent with the ILO methodology (see Supporting Information). Different from the ILO, however, NORC surveys children about the number of hours worked *directly*, and defines child labor based on *children’s self-reports* (“[u]sing the responses of children relating to engagement in cocoa production, we generated estimates of children’s engagement in child labor and in hazardous child labor in cocoa production-related activities”; [2], p.61).

Comparing the different data sources – which differ according to the reporting sources used for computing children in employment – is telling. 2016 ILO data for Côte d’Ivoire (<https://ilostat ilo org/topics/child-labour/>) indicated that 17.5% of 5-17 year-old children engaged in economic activity and household chores. In contrast, NORC data indicated that, as recently as 2018-19, 64% of all 5-17 year-olds in cocoa-producing regions of Côte d’Ivoire worked in the past 7 days, and 78% in the past 12 months.

While differences are striking, the NORC surveys cover different geographies and years than those used to compute official statistics, making it hard to attribute the gaps to under-reporting by adults in the ILO, UNICEF and World Bank data. Moreover, even if one would accept that parents report children in employment to a lesser extent than children themselves (see, for instance, [5, 6]) or that parents’ reports are prone to social desirability biases [4], it could be that children’s self-reports are similarly unreliable. Without additional data to *verify the actual prevalence* of children in employment, one simply cannot tell.

2.1.3 ENVERITAS methodology

ENVERITAS is a not-for-profit NGO that certifies coffee and (more recently) cocoa companies by verifying farming practices in their supply chain – from chemical usage to child labor (see, for instance, [10]). Since 2020, their methodology has been aligned with the ILO standards.

The certifier has specialized in reaching small, hard-to-reach and remote farmers, who often cannot even be located by traditional certifiers. To do that, ENVERITAS relies on partnerships to access fresh acquisitions of 50cm-resolution satellite data filtered for quality (with alleged maximum cloud covers of 15%, “crucial for finding (...) farms in cloudy equatorial regions”; [11]), combined with machine learning models to identify specific crops. Additional details on how such models have been successfully used to map coffee-growing households in different geographies can be found in [12]. In Côte d’Ivoire, the certifier has applied this methodology to identify constraints to quality education and early childhood development – including child labor – in cocoa-growing communities [13].

The combination of satellite imagery with machine learning models generates GPS coordinates for each ENVERITAS field team. Pins in Google Maps assign plots to be surveyed to each enumerator. Michael Kra, country lead for ENVERITAS, points out that

“When one surveys farmers, sometimes they do not tell you the truth. The pins will tell us to go where the truth is”. The idea is that farmers in hard-to-reach communities are less sensitive to social desirability biases. Michael Kra explains: *“We work in very rural communities. The road is often not good, requiring enumerators to use boats to cross rivers, motorcycles, etc. These communities being much less accessible, farmers have been rarely sensitized about child labor. This is why it is easier for them to openly disclose child labor. We can only find these communities because of satellite imagery”*. Supporting Information includes a sample satellite picture to illustrate how they are able to locate cocoa plantations based on aerial imagery. Supporting Information also compiles pictures from enumerators in the field as they survey some of these remote communities, documenting how they often have to cross dirt roads and flooded paths in order to reach them.

ENVERITAS further monitors cocoa harvests in the country with the help of satellite imagery, surveying farmers specifically during harvest season. With harvesting activities ongoing at the time of the survey, the potential presence of child labor becomes apparent: not only it is much harder for parents to falsely deny children in employment when it is more visible, but also, ENVERITAS complements survey data with direct field observation [12]. Michael Kra explains how surveying farmers during harvest time make it harder to under-report child labor: *“Enumerators can mark child labor as present even if a parent fails to report it (although this happens in less than 25% of surveys). They can add comments with additional information they have observed in the farm: e.g., ‘child is helping, climbing cocoa trees’”*.

2.2 Background for this study

2.2.1 Study sample

Our study takes place in the cocoa-producing regions of Aboisso and Bouaflé in Côte d’Ivoire. Along with Ghana, the country hosts almost 2/3 of the world’s cocoa production. This has been linked to one of the highest incidences of child labor worldwide, with nearly 1.6 million children employed in cocoa fields (<https://foodtank.com/news/2021/02/norc-report>).

We collected child labor data in the context of a broader research project, focused on evaluating different communication interventions to prevent student dropouts (see [14]). As such, we focus on a sample of primary school children, all of whom were enrolled at baseline. We discuss the implications of that sampling restriction for the generalizability of our findings in Section 4.

Supporting Information provides descriptive statistics for our study sample. Almost all (92%) of participating children are 5-11 years old – for whom any form of employment is considered child labor according to ILO guidelines. Half of children in our sample are girls and live in rural areas (defined according to their parents’ income source), and slightly over half of them are enrolled in the first primary cycle (CP2). Nearly a quarter (22%) of households in our sample are extremely poor – at baseline, they made at most a little over 1 USD a day. For only 18% of households income from all sources was more than 6 USD a day at baseline.

2.2.2 Campaign to discourage child labor

[14] evaluated a communication campaign (Eduq+, implemented by Brazilian EdTech Movva) that delivered nudges (reminders and motivational messages meant to make children’s school life ‘top of mind’) directly to parents’ mobile phones. The intervention was implemented over the 2018-19 school year. Nudges were organized in thematic sequences

– comprised of four messages –, with two messages delivered each week. Content was catered to each students’ age group. Messages tried to encourage parents to participate more actively in their children’s school life. There was some emphasis on showing up to school, especially in the context of parent-teacher meetings, and on discouraging harmful practices like corporal punishment as a disciplining strategy.

Several sequences explicitly discouraged child labor in cocoa fields, describing how it might detract from child development and learning. The language was careful, in an attempt to openly discuss the issues without creating stigma or setting social expectations that ultimately make it harder to track whether children work on the fields. To illustrate the approach, during the intervention, parents received a text message stating “*It is important that you child complete her/his education! School can provide a better future not only for her/him, but for your whole family*”. A few days later, another text encouraged them to “*Talk to your child about the importance of focusing on her/his education. Equally important as learning family traditions is learning the values and skills that only the school can teach*”. See [15] and [14] for additional details about the campaign.

In the experiment, nudges to parents were cross-randomized with nudges to teachers, aimed at increasing their attendance and time-on-task while teaching. For simplicity, we focus on discrepancies between parents’ and children’s reports between the treatment condition that had only parents nudged and the control group (which did not receive any messages) in the main text. Supporting Information compiles results for the other treatment arms.

2.3 Data and outcomes

2.3.1 Survey data

This research was approved by the University of Zürich Institutional Review Board (Protocols OEC IRB # 2018-035 and OEC IRB # 2019-052). Côte d’Ivoire did not require local IRB approval for non-medical research at the time. Our study comprises 198 CP2 and CE2 classrooms (second and fourth grades, respectively) across 99 Ivorian public schools in the cocoa-producing regions of Aboisso and Bouaflé. Within each school, we randomly drew 13 CP2 students and 12 CE2 students to be surveyed at baseline (at the beginning of the school year, from 01/10/2018) and end line (at the end of the school year, from 01/06/2019). Importantly, children and parents were surveyed independently. Enumerators ensured that this was the case, especially since we also tested children’s numeracy and literacy skills as part of the broader project, which required them to sit by themselves – only accompanied by our survey team.

We also conducted an additional follow-up survey at the beginning of the following school year (from 01/10/2019). In this follow-up, we surveyed all teachers and only one parent per classroom, but no children. This follow-up data focused on collecting additional information about work conditions for children in employment, but it also provides a measure of children in employment according to parents already after the growing season [16] and much closer to the timeline of the ENVERITAS data collection. Data were accessed for research purposes in 01/12/2019.

The sample comprises 1,285 CP2 students and 1,190 CE2 students surveyed at baseline along with their parents, in addition to 198 teachers. We were able to track all teachers, 1,157 CP2 students (90.0%) and 1,086 CE2 students (91.3%) at end line. We assigned replacement households in case the ones drawn could not be tracked by enumerators. All participants or their legal guardian(s) verbally consented to participate; minors still had to assent. All participants were free to discontinue their participation at any point

without sanctions. At the end line, no children – and less than 3.5% of their parents – refused to answer about their employment conditions. We discard child-parent pairs involving refusals, focusing on the 2,500 observations for whom we can compare children’s and parents’ reports. Out of those, we have information on all baseline characteristics that we use as controls in some of our specifications for 2,246 observations. Furthermore, we have information on home GPS location (which we use to compute a measure of remoteness) for 1,790 of those. Missing data for the latter often involves parents and children whom were surveyed at the school or in some other location.

Supporting Information compiles the survey questions related to children in employment in each wave. As indicated, we asked parents and children the same question about children in employment in cocoa fields at both baseline and end line (“*In the last month, have you [any of your children] engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation*”). This aligned with the data collected by ENVERITAS and NORC, and follows the ILO methodology for measuring children in employment. Since parents are not asked *specifically* about the child from whom we elicited employment information independently, that should decrease our ability to detect under-reporting by parents (e.g., if some answer affirmatively about older children, or about children who are already out of school).

Even though we also surveyed parents and children about other forms of employment (e.g., domestic work and construction), we do not analyze these data in this paper because we have no way to verify such reports. These additional measures of children in employment are described and analyzed in [14].

2.3.2 Certifier data

We use data on children in employment collected by ENVERITAS over January 2020, during the harvest season in Aboisso and in two sub-regions of Bouafle (Bouafle 2 and Tiapoum Adiake). ENVERITAS’s sampling frame relied on geographical units of 10,000 farmers, identified via satellite imagery. They randomly drew 125 farmers to be surveyed in each unit. 8,150 households were approached by ENVERITAS, 7,402 of which were successfully surveyed. The certifier also surveyed schools in regions outside of our study sample; see Supporting Information. We do not use data in these other regions, or collected prior to 2020 (before the survey instrument was consistent with the ILO methodology regarding the definition of child labor).

Adults surveyed by ENVERITAS were asked “*Do any of your children between 6 and 16 years old help you work on the cocoa farm?*”. Children were not surveyed directly by ENVERITAS; thus far, they have only piloted surveys with children in Tonkpi, a region outside our study sample.

2.4 Empirical strategy

2.4.1 Assessing the certifier claim that remoteness is inversely related to reporting biases

To assess whether remoteness is indeed inversely related to reporting biases, we estimate how children’s and parents’ reports vary with their household’s distance to the school (in km). Supporting Information includes the histogram for this measure, computed as the linear distance between the school and home GPS locations. 70% of households with valid home GPS data are within 1km of the school. While some are as far as 20km from the latter, the vast majority are within a 2.5 km radius of it.

Taking advantage of natural variation in distance to the school, we estimate a linear model for the association between distance and children’s reports, and between the former and the discrepancy between children’s and parents’ reports. We restrict attention to endline data, since that survey was the closest to harvest time – precisely when child labor is supposed to take place and, hence, around the same time the certifier enumerators were also in the field. As such, we further drop from the analyses households targeted by nudges to parents, as we later document treatment effects on reporting discrepancies.

When analyzing discrepancies, we also control for the school-level share of children in employment, and allow its coefficient to vary with whether the child has reported to work – all of which might influence the extent of reporting biases. We assess the sensitivity of our estimates to outlier observations and to specification choices by plotting the relation between reporting discrepancies and distance from the school, dropping observations further than 1.5 km from the school.

Importantly, we estimate those relations not only for child labor in cocoa fields, but also, for domestic work (helping with household chores, etc.) – much less sensitive to social stigma, and which has not been systematically targeted by NGOs and international organizations promoting children and adolescents’ rights in the region. Domestic work provides a counterfactual to assess the claim that discrepancies between parents’ and children’s reports are indeed driven by social desirability biases, and a placebo test for the claim that village remoteness should decrease the extent of reporting biases by the adults. Concretely, we hypothesize that (1) differences in children in employment across parents’ and children’s reports should be significantly lower for domestic work relative to labor in cocoa fields, and that (2) the magnitude of discrepancies should decrease with the distance to the school only for labor in cocoa fields.

2.4.2 Assessing the accuracy of different reporting sources

We assess the accuracy of children in employment reports according to parents and according to children themselves by comparing our survey data to the certifier data within the regions for which the two datasets overlap.

We report p-values from tests of differences in proportions of children in employment according to each source, considering equal population variances when comparing parents’ and children’s reports (through an Ordinary Least Squares regression, given the paired design), and unequal population variances when comparing any of them to certifier data – since the latter was collected from a different sample. We cluster standard errors at the regional level.

2.4.3 Assessing the sensitivity of the campaign’s effect sizes to different reporting sources

To document whether different reporting sources might lead to bias in evaluating the impacts of interventions to discourage child labor, we contrast effect sizes of nudges to parents on the prevalence of children in employment in cocoa fields based on children’s self-reports or on adult reports. We also allow treatment effects based on parents’ reports to vary with the distance to the school, to assess whether results are consistent with the spatial patterns we document for discrepancies between parents’ and children’s reports.

3 Results and discussion

3.1 Descriptive statistics based on children’s self-reports

Supporting Information documents the aggregate prevalence of children in employment, and that by student characteristics, according to children’s self-reports at baseline. 38.1% of children reported to have worked at least one hour in cocoa fields over the previous month. As a point of comparison, by the end of the school year (closer to harvest season), this figure was up to 50% (see Supporting Information). The prevalence of child labor was only slightly higher among fourth graders than that among second graders (39.6% vs. 36.8% at baseline, and 51% vs. 48% at end line). Boys were nearly 50% more likely to work in cocoa fields at baseline than girls (44.4% vs. 31.2%). The baseline prevalence of children in employment was higher for the bottom income bracket (38.9%), but not low even for the top income bracket (23.9%). Naturally, child employment in cocoa fields was much higher in rural areas (52.4% vs. 23.9%). For an account of adult reports about children’s work conditions, elicited in the follow-up survey, see Supporting Information.

Supporting Information also illustrates correlations between classroom-level prevalence of child labor and educational outcomes. Consistent with common sense, classrooms with a higher share of children in employment at baseline feature lower test scores by the beginning of the school year, and higher dropout rates over the course of the school year. Estimating a linear relation between the variables in each case suggests that moving from 0% to 20% children in employment is associated with about 0.08 s.d. lower test scores – what children tend to learn in one school quarter, and the magnitude of effect sizes of many educational interventions, such as nudges to parents evaluated in this setting [15]. Similarly, moving from 0% to 40% children in employment is associated with roughly doubling dropout rates.

While these associations are not causal, they help understand the centrality of the issue for governments and international organizations monitoring children’s rights. This is why accurate measurement is key.

3.2 Household remoteness and reporting discrepancies

We start by investigating the association between a measure of household remoteness – the distance from each household to the school where the child participating in the study was enrolled – with discrepancies between parents’ and children’s reports. For these analyses, we restrict attention to the 1,790 households in our sample with valid GPS location data, and drop observations targeted by nudges to parents (in combination with or independently of nudges to teachers) – as we later show that the intervention affects reporting (see Section 3.4) –, resulting in 1,363 observations.

If remote communities are indeed less often or less intensely sensitized, a challenge is that remoteness might affect both under-reporting *and* the true prevalence of child labor. To separate these potential effects, Table 1 starts by assessing whether children’s self-reports vary systematically with our measure of remoteness (in column 1); and columns (2) to (4) assess how the discrepancy between children’s and parents’ reports varies with that measure. Panel A focuses on labor in cocoa fields; and Panel B turns to domestic labor. All columns control for children’s self-reports; column (3) further controls for the school-level prevalence of child labor (according to children) and its interaction with each child’s self-report, as those might also affect the extent of reporting biases; column (4) additionally controls for baseline characteristics. Because we include a school-level measure of prevalence, we cluster standard errors at the school level.

In Panel A, column (1) documents that distance from school is not associated with

children’s self-reports. The average end-line prevalence of children in employment within the sub-sample with valid home GPS data was 50.6%; such prevalence did not systematically increase with the household’s distance to the school. This is consistent with the idea that more remote communities do not necessarily feature more children working in cocoa fields. In turn, columns (2) to (4) document that such distance was systematically associated with lower discrepancies between children’s and parents’ reports. While, in this sub-sample, the average discrepancy between children’s and parent’s reports was 22.9 p.p., we estimate that it decreased by 1.7 p.p. (significant at the 5% level) with every km away from school, 7.4% of the average discrepancy. Incidentally, other estimated coefficients are also informative: consistent with our previous discussion, the more common children in employment is at the school, the lower the reporting discrepancies – consistent with the role of social expectations. Results support the certifier claim that remoteness is inversely associated with reporting biases, without necessarily being associated with children in employment itself.

Figure 1 documents that the relation between household remoteness and reporting discrepancies is not an artifact of outlier observations or specification choices. Dropping observations further than 1.5 km from the school from the analyses, the binscatter plot shows that discrepancies indeed decline systematically with distance. If anything, the relation is even stronger within that range: a linear regression estimates that discrepancies decrease 4.5 p.p. per km, from a baseline of 25% (an effect size significant at the 10% level).

One concern is that our measure of remoteness captures some other underlying differences in these communities other than those linked to prior exposure to sensitization campaigns – and corresponding differences in the prevalence of social desirability biases as a result. To tackle that concern, Panel B of Table 1 replicates the previous analyses to domestic labor in an attempt to rule that out. Importantly, Figure S10 in Supporting Information documents that, in the full sample, discrepancies are significantly lower when it comes to domestic labor. Here we are interested in whether such discrepancies are *not* negatively associated with our measure of household remoteness. Panel B documents that, different from labor in cocoa fields, differences between parents’ and children’s reports do not systematically change with the distance from the school.

All in all, results lend confidence to the claim that the certifier’s sampling strategy might indeed reveal the ‘ground truth’ when it comes to the local prevalence of child labor. Moreover, they suggest children’s self-reports in the remote villages surveyed by ENVERITAS are likely representative of those in other villages. As such, we can confidently use that data to validate the prevalence of child labor according to different reporting sources in the sample for which the two datasets overlap, in the next section.

Table 1: Association between household distance to school, children’s self-reported employment, and within-household reporting discrepancies between parents and children

	Children in employment (self-report)	Discrepancy (child - parent)		
	(1)	(2)	(3)	(4)
Panel A - Work in cocoa fields				
Distance from school (km)	0.014 (0.014)	-0.014 (0.012)	-0.017** (0.007)	-0.017*** (0.006)
Child labor (self-report)		0.633*** (0.037)	1.063*** (0.066)	1.073*** (0.072)
Average child labor at school			-0.406*** (0.078)	-0.388*** (0.086)
Child labor × Average child labor			-0.512*** (0.122)	-0.511*** (0.131)
Sample mean	0.506	0.229	0.229	0.228
R-squared	0.001	0.376	0.471	0.477
Observations	1363	1362	1362	1235
Panel B - Domestic work				
Distance from school (km)	-0.003 (0.007)	0.010 (0.012)	0.009 (0.011)	0.000 (0.009)
Domestic work (self-report)		0.617*** (0.046)	0.523 (0.432)	0.383 (0.461)
Average domestic work at school			-0.412 (0.474)	-0.485 (0.524)
Domestic work × Average domestic work			0.136 (0.507)	0.339 (0.536)
Sample mean	0.886	0.082	0.082	0.088
R-squared	0.000	0.212	0.215	0.243
Observations	1363	1361	1361	1234
Baseline controls	No	No	No	Yes

Notes: Endline data excluding households targeted by nudges to parents (in combination with or independently of nudges to teachers). In column (1), Children in employment (self-report) = 1 if the child answered affirmatively to the question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation” in Panel A, and “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. In columns (2) to (4), Discrepancy (child - parent) = self-reported employment - parent’s reported children in employment, whereby the latter = 1 if the parent answered affirmatively to the question “In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, in Panel A, and “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. Distance from school (km) is the linear distance between the household and their child’s school GPS location, in kilometers. School-level % children working in the field averages (self-reported) and School-level % children in domestic work are the percentage of children surveyed in each school who answered affirmatively to each question. Baseline controls include child gender, grade indicators, standardized test scores (averaged across numeracy and literacy); and summary measures of parental engagement, student effort, socio-emotional skills, working memory, visual attention, impulsivity, self-esteem, and mindset (see [14]. Summary measures computed following [17], standardizing each component by normalizing values by the mean and standard deviation of the control group at the baseline survey within each grade, and then averaging over all standardized components. Sample restricted to observations with valid home GPS coordinates. All regressions estimated through Ordinary Least Squares, with standard errors clustered at the school level.

*** p<0.01, ** p<0.05, * p<0.10

Figure 1: Within-household discrepancies of children in employment in cocoa fields (self- vs. adult reports), by distance from school (in km)

Notes: The y-axis displays the difference between children in employment in cocoa fields reported by children themselves and that reported by their parents, in response to the question “In the last month, have you / has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”. The x-axis displays distance from the child’s home to the school (in kilometers). Observations with distance greater than or equal to 1.5 km were dropped, resulting in the exclusion of 12% of observations (490 households). Data divided into 30 quantiles. Fitted line estimated through Ordinary Least Squares (OLS), with no control variables; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

3.3 Validating survey data with certifier data

We now turn to the comparisons between independent reports by parents and children in our end-line survey data, and ENVERITAS data. Figure 2 presents the prevalence of children in employment according to each reporting source, along with p-values for pair-wise statistical tests of differences in proportions. We restrict attention to the aggregate prevalence based on each reporting source across the regions where both datasets overlap in the main text. Moreover, we also restrict attention to households in the ENVERITAS data whose children were *all* enrolled in school at the time of the survey, to make its sampling frame consistent to that of our surveys. Supporting Information further documents comparisons within each region, and relaxing that sample restriction. Our findings are very robust across the different sets of comparisons.

In the absence of under-reporting, the rate of children in employment according to either parents’ or children’ reports should be identical (in fact, since the question directed to parents was about *any* of their children, the former should be weakly greater than the latter). This is, however, strictly at odds with what we find. Figure 2 documents that, in regions with subsequent verification, 45.5% of children reported to have worked in cocoa plantations in the previous month, matching almost exactly the 47.8% prevalence indicated by the certifier (p-value of the difference = 0.746). In contrast, only 16.2% of parents in those regions reported employing children – a nearly 2/3 reporting gap ($p = 0.000$). Figure S13 in Supporting Information shows that, across regions, under-reporting by adults was striking, ranging from 50% to 90% ($p = 0.000$ in each case).

Figure 2: Validation of child labor measures using third-party data, households with all children in school

Share of students who worked in cocoa plantations, by reporting source

Notes: Columns show the share of children who worked at least an hour in cocoa fields over the previous month, according to children (in red), parents (in blue) and ENVERITAS (in green), along with 90% confidence bars. In the ENVERITAS data, we restrict attention to households whose *all* children between 6 to 11 years old were enrolled in school at the time of the survey (in Supporting Information, we relax that sample restriction). The figure reports the average prevalence across all regions for which survey data overlaps with ENVERITAS data. Children answered the following question at end line: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question at end line: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. In the survey conducted by ENVERITAS during the harvest season (identified through satellite imagery), farmers answered the following question: “Do any of your children between 6 and 16 years old help you work on the cocoa farm?”. For the measures reported by children and parents, observations are restricted to the control group of the communication intervention. P-values from tests of proportions with unequal population variances (when children’s or parents’ reports are compared to ENVERITAS data; accounting for each source’s intra-cluster correlation computed at the regional level), and from tests of proportions with equal population variances (when comparing children’s and parents’ reports; through Ordinary Least Squares regressions, clustering standard errors at regional level). Sample sizes are as follows: (i) Children: 475; (ii) Parents: 475; and (iii) ENVERITAS: 232.

3.3.1 Robustness to differences in the timing of the surveys

One potential caveat of the results above is that our end-line survey (when both children and parents were asked about children in employment) dates from June 2019, almost 6 months prior to the ENVERITAS data collection. Here we assess whether such difference matters for our findings.

To do that, we take advantage of our follow-up wave, conducted in October 2019 (much closer to January 2020, when ENVERITAS conducted its survey). While children were not surveyed in that wave, we can compare parents' reports across the end-line and follow-up surveys to gauge whether adult reports converged to children's self-reports as harvest season was approaching. Supporting Information documents that this was *not* the case. At the follow-up wave, only 28% of parents admitted that children worked at all in cocoa fields during the school year. This figure was still only about half that reported by children at the end-line survey – even though, at the follow-up wave, we asked parents to report on work performed by children in cocoa fields *at any point during the previous school year*. Supporting Information further documents that teacher reports of child labor over the course of the previous school year were statistically identical to parents' reports at the end line.

3.4 How adult reports can also bias policy evaluations

Last, we present evidence that relying on parents' reports for children in employment might not only lead to inaccurate estimates about the prevalence of child labor, but also, to potentially incorrect conclusions about the effects of interventions to discourage it. Table 2 estimates treatment effects of nudges to parents on children in employment at end line, relative to the control group. Column (1) uses children's self-reports as the dependent variable, while columns (2) to (4) use parents' reports. Columns (1) and (2) use the full sample for whom we have children's and parents' reports; columns (3) and (4) restrict attention to those with valid home GPS information. Column (4) allows treatment effects to vary with our measure of remoteness – the linear distance from the household to the school where the participating child was enrolled (in km). The idea is that part of the effects of the intervention on parents' reports might play out through its effects on reporting biases (especially since its goal was to openly address the issue without creating stigma). All columns control for baseline characteristics. Because the over-arching intervention evaluated in [14] was centered around teachers, we cluster standard errors at the classroom level.

In Panel A of Table 2, column (1) documents that the intervention did not systematically affect children in employment according to self-reports. In contrast, based on parents' reports, one would have concluded that the intervention *increased* the prevalence of children in employment (column 2). The effect is sizeable: an 8.9 p.p. increase, significant at the 10% level – 40.3% of its end-line prevalence in the control group. Column (3) documents that the same holds within the sub-sample with valid home GPS information. Most importantly, column (4) documents that this is likely an artifact of the interplay of the intervention with social desirability biases in communities that have been previously sensitized: very close to the school, the estimated effect size is even larger (a 14.5 p.p. increase, significant at the 5% level), but it sharply decreases with distance to school, fading out at an approximately 3 km radius, based on our linear estimate.

Table S3 in Supporting Information further documents that results are robust to outlier observations: the coefficient of nudges to parents in column (4) is virtually unchanged if we winsorize distance from school at the 99th percentile.

Table 2: Treatment effects of nudges, by reporting source and household remoteness

	Children in employment (self-report)	Children in employment (parent's report)		
	(1)	(2)	(3)	(4)
Panel A - Work in cocoa fields				
Nudge to parents	0.054 (0.040)	0.089* (0.045)	0.120** (0.056)	0.145** (0.061)
Distance to school (km)				0.016 (0.012)
Nudge to parents \times Distance				-0.061* (0.031)
Control mean	0.498	0.221	0.230	0.230
R-squared	0.115	0.091	0.094	0.096
Observations	2,246	2,176	1,608	1,608
Panel B - Domestic work				
Nudge to parents	0.046** (0.020)	0.051** (0.022)	0.066** (0.026)	0.072** (0.028)
Distance from school (km)				-0.003 (0.012)
Nudge to parents \times Distance				-0.015 (0.031)
Control mean	0.878	0.789	0.794	0.794
R-squared	0.034	0.044	0.050	0.050
Observations	2,246	2,175	1,607	1,607
Controls (baseline survey)	Yes	Yes	Yes	Yes

Notes: In column (1), Children in employment (self-report) = 1 if the child answered affirmatively to the working question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation” in Panel A, “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. In columns (2) to (4), Children in employment (parent’s report) = 1 if the parent answered affirmatively to the question “In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, in Panel A, and “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. Distance from school (km) is the linear distance between the home and their child’s school GPS locations, in kilometers; Nudge to parents = 1 if the parent (but not teachers) was assigned to nudges, and 0 otherwise. Baseline controls include treatments arms, child gender, grade indicators, standardized test scores (averaged across numeracy and literacy); and summary measures of parental engagement, student effort, socio-emotional skills, working memory, visual attention, impulsivity, self-esteem, and mindset (see [14]. Summary measures computed following [17], standardizing each component by normalizing values by the mean and standard deviation of the control group at the baseline survey within each grade, and then averaging over all standardized components. In columns (3) and (4), sample restricted to observations with valid home GPS coordinates. All regressions estimated through Ordinary Least Squares with standard errors clustered at the classroom level. *** p<0.01, ** p<0.05, * p<0.10

Reassuringly, when it comes to domestic work (Panel B), not only were effect sizes very similar regardless of reporting sources (an increase of 4.6 p.p. according to children, in column 1, and of 5.1 p.p. according to parents, in column 2, both significant at the 5% level), but also, treatment effects did *not* vary systematically with distance to the school.

All in all, results are consistent with the interpretation that the social stigma around child labor in cocoa fields (in contrast to domestic work) might introduce bias in official statistics when it comes to evaluating the impacts of policies to discourage child labor whose impacts might at least partially affect its social acceptability – and, hence, magnify or attenuate reporting biases in parents’ reports.

4 Conclusions

Data on the prevalence of child labor across space and over time is critical for governments and international organizations committed to ensuring children’s rights. In the cocoa industry, particularly intensive in child labor given its low rate of mechanization, data issues have posed important challenges to monitoring and enforcement efforts by international organizations and policymakers over the years.

Our finding that child labor statistics following the ILO methodology not only misrepresent the prevalence of child labor, but also mischaracterize its trends (where interventions were in place), raises critical concerns. While the general sentiment of the literature on child labor is that substantial progress has been achieved in recent decades (“[a]n important lesson from all the literature reviewed herein is that child labor can change dramatically and quickly in countries as a result of changes in the economic and policy environment.”; [18], pp. 27-28), our results call that sentiment into question. Concretely, in Supporting Information, we use our estimates to calibrate predictions of the bias-adjusted prevalence of child labor in Côte d’Ivoire and Ghana. Under the stated assumptions, child labor could affect nearly 6.9 million children in these countries, 50% more than in the latest World Development Indicators’ statistics. Naturally, this does not invalidate the claim that child labor might be caused by lack of economic development or that it can respond to policy changes [18]; however, it does raise caution about carefully interpreting changes in adult reports in contexts of fast transitions – in light of how such changes might impact not only the economic fundamentals that generate trade-offs between employing children vis-a-vis investing in their human capital, but also, reporting biases themselves (e.g., by changing parents’ beliefs about what is socially appropriate).

Based on our results, asking children independently about whether they work (and, if so, how many hours) could yield child labor indicators consistent with costly-verification data. There are, however, technical challenges involved in interviewing children, including whether children are asked inside or outside the household, and variations in the understanding of what exactly characterizes ‘work’ by children in different countries (see [19] and [6] for a broader discussion of different framing issues). While the experiences of NORC and Tulane University in Côte d’Ivoire and Ghana since 2008-09 could inform the replication of self-reports elicitation of child labor by other agencies moving forward, surveying children directly involves some ethical challenges as well. In particular, children might be exposed to violence if interviews trigger backlash from parents, even when they are conducted in a different setting (e.g., schools). These concerns also limit the potential of technologies worn by children, such as GPS trackers, despite their validated accuracy [7].

For those reasons, finding ways to limit reporting biases in adult surveys may be a superior (and necessary) alternative. Progress in this space has, however, been slow. Most studies focus on indirect elicitation methods (i.e., list experiments, in various forms; e.g., [3]),

but these methods generate estimates that are typically imprecise and not necessarily closer to the ‘ground truth’ (see, for instance, <https://statmodeling.stat.columbia.edu/2014/04/23/thinking-list-experiment-heres-list-reasons-think/>). In contrast, leveraging technologies to survey adults in communities less subject to previous sensitization and during periods when it is easier to observe child labor might be a much better way forward.

Naturally, that also involves its own challenges. Although we have documented that remoteness is not systematically associated with the true prevalence of child labor (based on children’s self-reports) in the close vicinity of schools, it might be that the communities surveyed based on satellite imagery are *not* representative of the territories or the populations of interest, leading to biased aggregate estimates of children in employment and child labor. Moreover, if remote sensing through satellite imagery is imperfect (e.g., if the geographical coverage of the available data is selective either when it comes to which regions it covers or which farms it surveys within each region), then the estimates will again be biased. There are also cost considerations. Obtaining access to high-frequency data that can inform prediction models (and hiring staff or external vendors to train and update such models) might be expensive. Such requirements might limit the ability to learn timely about how child labor evolves across space, particularly in response to policy interventions.

While we expect our contributions to generalize beyond the Ivorian cocoa industry, a limitation is that our sample consists entirely of school children – all of whom were enrolled at the time of our surveys. This excludes children who had dropped out of school (or never enrolled in the first place), presumably more likely to work in cocoa fields. While their parents might be less sensitive to social desirability biases (since their children are not in school), there are other potential sources of social pressure to under-report children in employment – particularly economic pressures linked to restrictions to child labor in global supply chains. Investigating if parents’ under-reporting changes with children’s school enrollment remains a promising avenue for future work.

Acknowledgments

We thank Maxime Deroubaix and Michael Kra from ENVERITAS for explaining the survey methodology in detail and kindly granting access to their geo-referenced child labor data. We also acknowledge helpful comments from Andrew Dillon and Megan Passey (International Cocoa Initiative), excellent research assistance by Laura Ogando, Gabriel de Campos and Thiago da Costa, and field supervision by Nicolo Tomaselli and Innovations for Poverty Action. All remaining errors are ours.

Supporting Information

Supporting Information file contains S1-S18 Figures and S1-S5 Tables.

References

- [1] ILO, UNICEF. Child Labour: Global estimates 2020, trends and the road forward. 2021.
- [2] Sadhu S, Kysia K, Onyango L, Zinnes C, Lord S, Monnard A, et al. NORC Final Report: Assessing Progress in Reducing Child Labor in Cocoa Production in Cocoa Growing Areas of Côte d’Ivoire and Ghana; 2020.
- [3] Lépine A, Toh W, Treibich C. Estimating the prevalence of child labour in the cocoa industry via indirect elicitation methods: A mixed-methods study. 2022.
- [4] Jouvin M. Addressing social desirability bias in child labor measurement : an application to cocoa farms in Côte d’Ivoire. 2021;(2021-08).
- [5] Galdo J, Dammer A, Abebaw D. Gender Bias in Agricultural Child Labor: Evidence from Survey Design Experiments. 2020;(13826).
- [6] Dillon A. Measuring Child Labor: Comparisons between Hours Data and Subjective Measures. Research in Labor Economics. 2010;(31):135-69.
- [7] Dillon A, Hanicotte H, Valenza M. Pilot data analysis & methodological insights on labor measurement. Innovations for Poverty Action; 2017.
- [8] Christian Science Monitor. Child labor in decline: Life gets better for 94 million children; 2019. Last accessed February 3, 2022. Available from: <https://www.csmonitor.com/World/Points-of-Progress/2019/0828/Child-labor-in-decline-Life-gets-better-for-94-million-children>.
- [9] International Labour Organization. Methodology of the global estimates of child labour, 2012-2016. 2017.
- [10] Tran H, Nathan S, Ilmma A, Burkiewicz M, Wisana I. Identifying limiting factors for feasible productivity improvement for smallholder farmers in coffee sector in Indonesia. Asian Journal of Agriculture. 2021;5(2):53-60.
- [11] Enveritas. Supply Chain Verification: Finding Small-Holder Coffee Farms; 2020. Last accessed 8 February 2020. Available from: <https://storymaps.arcgis.com/stories/10ef0006e4d94b86b228748c7082cf85/print>.
- [12] Tran H, Nathan S, Ilmma A, Burkiewicz M, Wisana DGK. Identifying limiting factors for feasible productivity improvement for smallholder farmers in coffee sector in Indonesia. Asian Journal of Agriculture. 2021;5(2).
- [13] TRECC. Transforming Education in Côte d’Ivoire – Reflections, Lessons Learned, Next Steps; 2021. Last accessed March 4, 2022. Available from: <https://www.cocoainitiative.org/sites/default/files/resources/TRECC-5-YEAR-REPORT-4.pdf>.
- [14] Wolf S, Lichand G. Nudging Parents and Teachers to Improve Learning and Reduce Child Labor in Côte d’Ivoire. mimeo. 2023.
- [15] Lichand G, Wolf S. Arm-wrestling in the classroom: the non-monotonic effects of monitoring teachers. University of Zurich, Department of Economics, Working Paper. 2022;(357).

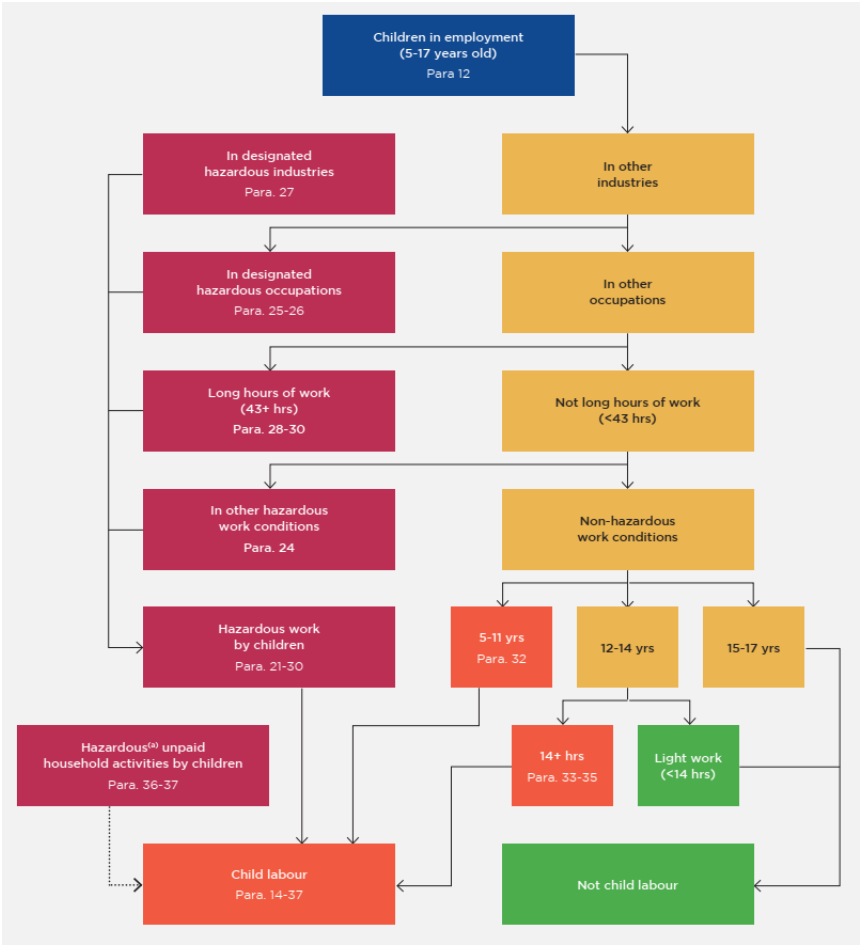
- [16] Yoroba F, Kouassi BK, Diawara A, Yapo LA, Kouadio K, Tiemoko DT, et al. Evaluation of Rainfall and Temperature Conditions for a Perennial Crop in Tropical Wetland: a Case Study of Cocoa in Côte d'Ivoire. 2019.
- [17] Kling J, Liebman J, Katz L. Experimental analysis of neighborhood effects. *Econometrica*. 2007;75(1):83-119.
- [18] Edmonds EV, Theoharides C. Child Labor and Economic Development. In: *Handbook of Labor, Human Resources and Population Economics*. Springer; 2021. p. 1-29.
- [19] Guarcello K, Kovrova I, Lyon S, Manacorda M, Rosatti F. Towards consistency in child labour measurement: assessing the comparability of estimates generated by different survey instruments. 2010;(54).

Additional details on classification methodologies

This Appendix compiles additional details on how the International Labor Organization (ILO), NORC and ENVERITAS compile data on children’s work hours and work conditions, and on how these data is used to compute children in employment and child labor statistics in each case.

ILO

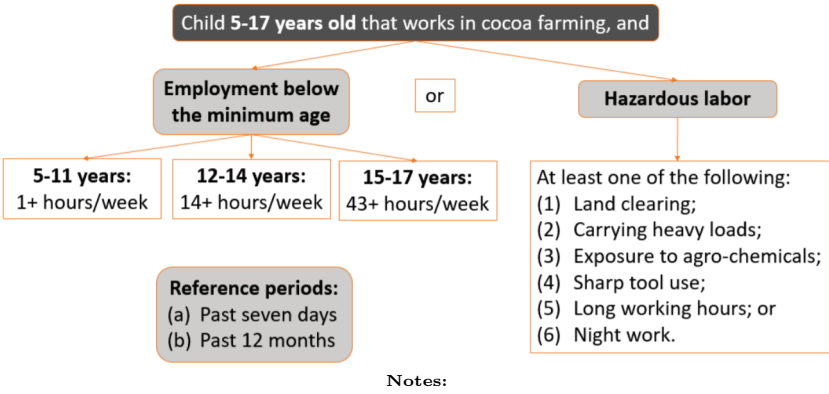
Figure S1: Conceptual framework of the ILO global estimation of child labor



Notes: extracted from [1].


NORC

Figure S2: Conceptual framework of the NORC estimation of child labor in cocoa farming



ENVERITAS

Figure S3: Geographical coverage of ENVERITAS data



Notes: Elaborated by the authors.

Figure S4: Satellite picture of a community visited by ENVERITAS



Notes: Satellite picture of Baho-Brousse, a community visited by ENVERITAS on Jan/2020. The surrounding cocoa fields are clearly visible in the picture (those not arranged in rows, which, in turn, are rubber plantations).

Figure S5: Enumerators surveying hard-to-reach communities (Jan/2020)



Figure S6: Enumerators surveying hard-to-reach communities (Jan/2020)



Survey instruments

This Appendix compiles the questions used to assess children in employment through our business-as-usual surveys, organized by wave and by whom was asked in each case. For all details on the survey instruments used in the context of the over-arching project, see [15].

Table S1: Survey questions about children in employment

Questions	Timeline	Respondent
“In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”	Baseline	Children
“I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”	Baseline	Parents
“In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”	End line	Children
“I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”	End line	Parents
“Did (child name) work in cocoa fields more than 10 hours a week in a typical week last year?”	Follow-up	Parents
“Did (child name) work in cocoa fields more than 10 hours a week during holidays last year?”	Follow-up	Parents
“To the best of your knowledge in this list of students, can you point out the students who worked in cocoa fields over 10 hours a week in a typical week last year?”	Follow-up	Teachers
“To the best of your knowledge in this list of students, can you point out the students who worked in cocoa fields over 10 hours a week in a typical week during vacation last year?”	Follow-up	Teachers
“To the best of your knowledge in this list of students, can you point out the students who worked at all in cocoa fields during the school year last year?”	Follow-up	Teachers

Notes:

Descriptive statistics

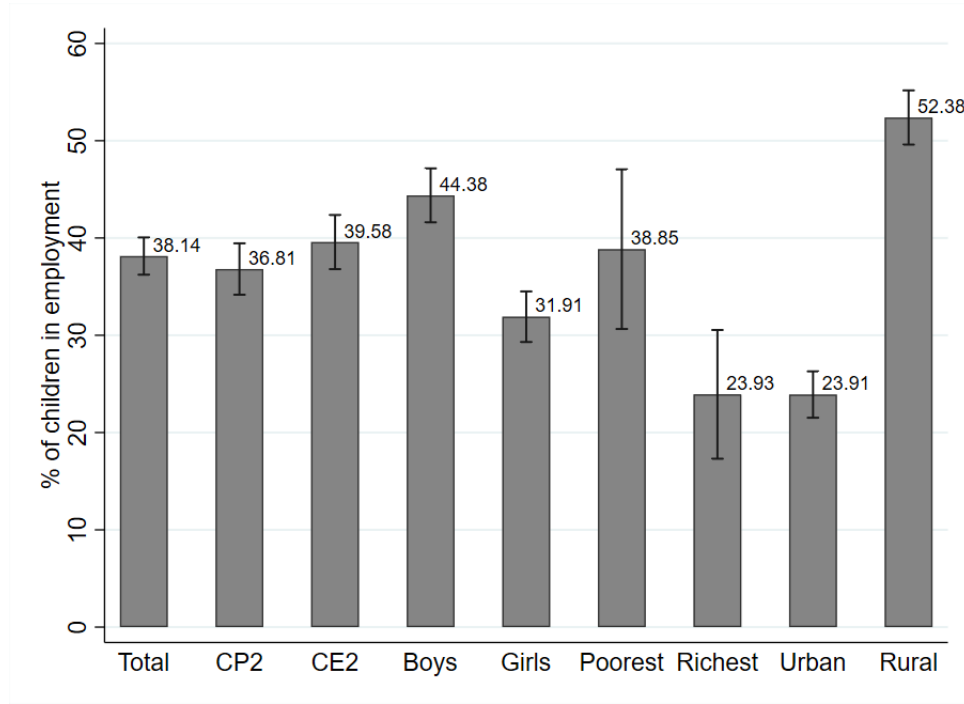
This Appendix compiles descriptive statistics of the study sample, based on our baseline survey. Table S2 provides summary statistics for student and household characteristics. Figure S7 showcases the share of children in employment by student and household characteristics. Next, Figure S8 displays classroom-level correlations between the baseline share of children in employment (based on children’s self-reports) and baseline standardized test scores (Panel A) and student dropout rates by the end of the school year (Panel B; restricting attention the control group of the intervention). Figure S9 displays the histogram of households’ distance to the school where their children were enrolled at the time of the survey (the measure of remoteness we explore in the main text). Last, Figure S10 showcases the prevalence of children in employment by activity and reporting source.

Table S2: Descriptive statistics (baseline survey)

	Mean	S.D.	Obs
Child is a girl	0.50	0.50	2,475
Child age			2,150
Under 5 years old	0.00	0.04	
5-11 years old	0.92	0.27	
12-14 years old	0.07	0.26	
15 years old and above	0.01	0.08	
Enrolled in 1st primary cycle (CP2)	0.52	0.50	2,475
Rural household	0.50	0.50	2,471
Household monthly income (in 2015 USD)			2,177
Less than USD 19	0.06	0.24	
USD 19-37	0.16	0.37	
USD 37-55	0.16	0.36	
USD 55-92	0.21	0.41	
USD 92-185	0.23	0.42	
USD 185-370	0.11	0.31	
More than USD 370	0.07	0.26	

Notes: CP2 is the second grade for the 1st primary cycle in Côte D’Ivoire education. Rural areas are defined according to parents’ main occupation: agricultural or plantation activities are defined as rural. Household monthly income was reported in CFCA and converted to 2015 USD.

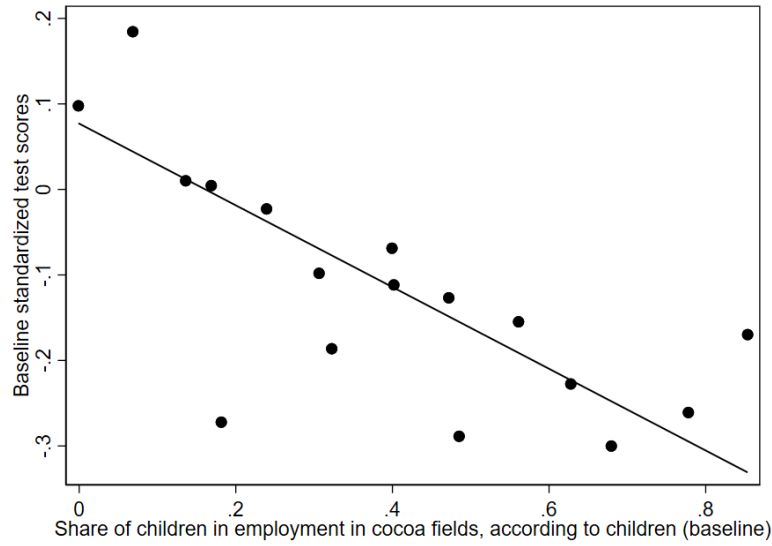
Figure S7: Share of students who worked for at least one hour in cocoa plantations over the last month, according to children



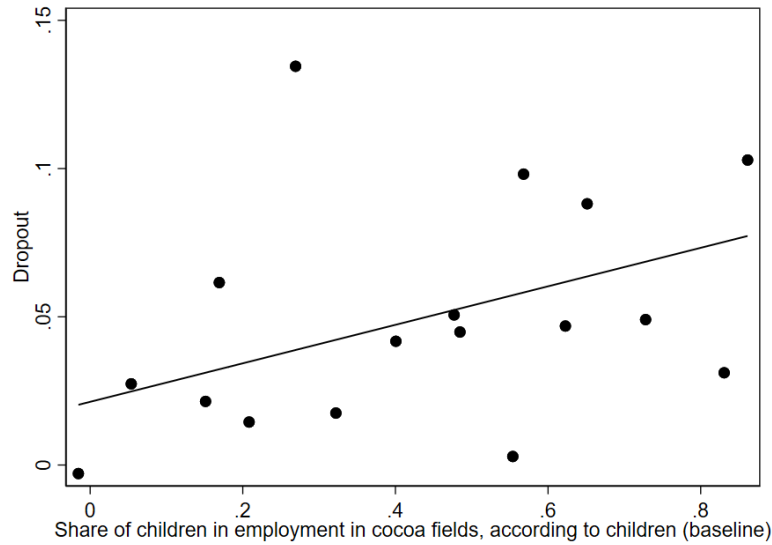
Notes: Figure S7 shows the share of students who report to have worked in cocoa plantations in the last month for one hour or more at baseline, in response to the following question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. The first bar comprises the whole sample in [15]; additional bars consider the indicated sub-samples. “Poorest” comprises households with monthly income reported by parents below 10,000 CFA (~ 19 USD), while “Richest” comprises those with monthly income reported by parents above 200,000 CFA (~ 372 USD). Rural and urban areas are defined according to parents’ main occupation (agricultural or plantation activities are assigned to the former). Samples sizes are the following: (i) Total: 2,475; (ii) CP2: 1,285; (iii) CE2: 1,190; (iv) Boys: 1,237; (v) Girls: 1,238; (vi) Poorest: 139; (vii) Richest: 163; (viii) Urban: 1,234; (ix) Rural: 1,237.

Figure S8: Baseline correlation between child labor and educational outcomes

Panel A: Correlation between test scores and child labor




Panel B: Correlation between student dropouts and child labor



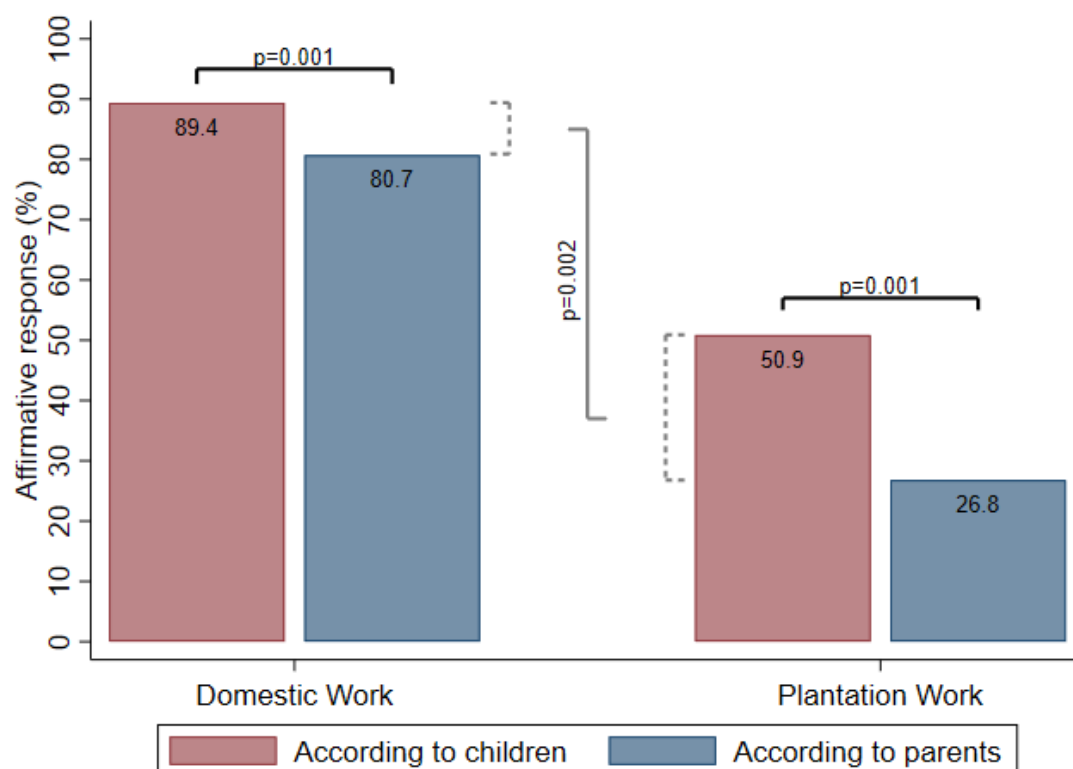
Notes: Panel A reports a bin-scatter plot of baseline standardized test scores as a function of baseline children in employment in cocoa fields (self-reported by children). Standardized test scores are a summary measure of numeracy and literacy test scores (averaging across each component, normalized by their mean and standard deviation in the control group), following [17]. Children in employment stands for the baseline share of students who report to have worked in cocoa plantations in the last month for one hour or more, in response to the following question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Test scores and the prevalence of children in employment are averaged at the classroom level. Panel B reports a bin-scatter plot of student dropout rates, based on administrative data (see [15]), as a function of baseline children in employment in cocoa fields (self-reported by children). Student dropouts and the prevalence of child labor are averaged at classroom level. Because student dropouts are defined at the end line, Panel B restricts observations to the control group of the intervention.

Figure S9: Histogram of the distance from households to the school (in km)



Notes: Distribution for the sub-sample of children with valid GPS coordinates associate with their household location.

Figure S10: Share of children in employment, by activity and reporting source




Notes: Bar shows the share of respondents that had an affirmative response to the works question in the endline survey. Specifically, parents were asked "In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives?" and "Work in a cocoa plantation". Children were asked "In the last month, have you engaged in one or more of the following activities, for one hour or more? Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives?" and "Work in a cocoa plantation". The p-value in the figure is relative to the paired t-test clustering standard errors at the regional level. The vertical p-value is relative to the test on the differences in responses given between parents and children between working conditions. Sample sizes are the following: (i) Children 2500; (ii) Parents: 2418 for domestic work question and 2019 for plantation work.

Prevalence of children in employment and hazardous work conditions measured in the follow-up survey

This Appendix compiles results on children’s work hours and work conditions based on the follow-up survey (conducted in October 2019). Figure S11 plots the share of affirmative answers to different questions about children in employment. Figure S12 adds to the comparisons in the main text what teachers report about each student in the follow-up survey, when they were asked to report whether each of them had worked in cocoa fields at any point over the course of the previous school year. As the figure shows, teachers’ and parents’ reports (the latter, collected at the end-line survey) are statistically identical ($p=0.23$), and significantly under-estimate the end-line prevalence of children in employment according to children ($p<0.001$ in each case).


Figure S11: Parents’ reports on children in employment and hazardous work conditions at the follow-up wave (October 2019)



Notes: Bars show the share of students who parents have reported to be in the following conditions (yes/no questions): 1) Help family: “Did (child name) work at all to help you around home, assist in a family business or earn pocket money outside school hours under adult supervision during the school year last year?”; 2) Help family 1 hour: “Did (child name) work at all to help you around home, assist in a family business or earn pocket money outside school hours under adult supervision over 1 hour a week during the school year last year?”; 3) Cocoa: “Did (child name) work at all in cocoa fields during the school year last year?”; 4) Cocoa > 10 hrs: “Did (child name) work in cocoa fields more than 10 hours a week in a typical week last year?”; 5) Cocoa > 10 hrs vacation: “Did (child name) work in cocoa fields more than 10 hours a week during vacation last year?”; 6) Dangerous work: “Was (child name) involved in activities in cocoa fields such as clearing of forests and felling of trees, bush burning, manipulating agrochemicals or using sharp tools during the school year last year?”; 7) Night work: “Did (child name) work between 7 p.m. and 7 a.m. during the school year last year?”; 8) Heavy work: “Was (child name) engaged in heavy physical labor in a typical week last year?”; 9) Machete: “Did (child name) used a machete while working in the fields last year?” 10) Hurt: “Did (child name) get hurt at least once while working in the fields last year?”. All measures were collected at the follow-up surveys [15]. Across all bars, the sample is restricted

to the control group.

Figure S12: Share of students who worked in cocoa plantations during the school year, according to different sources



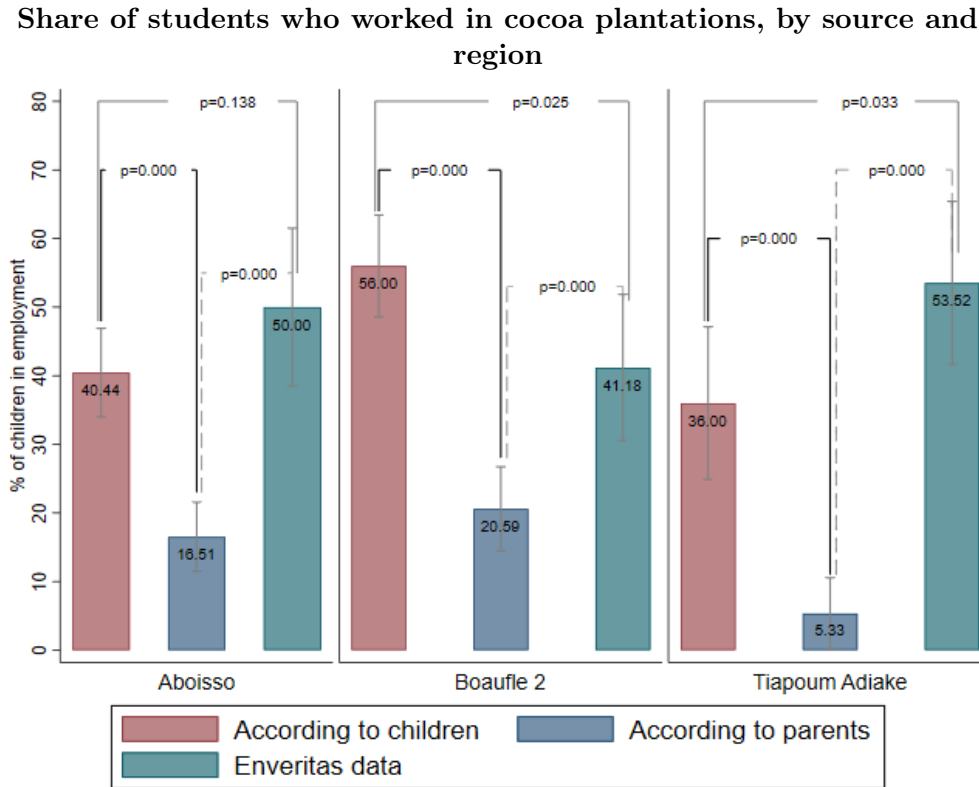
Notes: Bars show the share of students who have worked in cocoa plantations according to children (in red), parents (in blue), or teachers (in green). Children and parents answer at end line if the child engaged in one hour or more in cocoa plantation activities over the last month, as described in Table S1. Teachers answer at follow-up if each of their students worked in cocoa plantation at all over the last school year, as described in Table S1. Across all bars, the sample is restricted to the control group. The first set of bars comprises the whole sample in [15], while the additional ones split the sample by primary grades (CP2 and CE2). Sample sizes are the following: (i) Total: 625 ; (ii) CP2: 323; (iii) CE2: 302.

Additional results

This Appendix compiles additional results. Figures S13, S14 and S15 complement our main results not only by breaking down our analyses of reporting discrepancies by region for which our surveys overlap with the certifier data, but also by expanding the sample to include additional households surveyed by ENVERITAS whose children were not necessarily all enrolled in school. The patterns we document are very robust across sub-samples and sample restrictions.

Next, Figure S16 illustrates how under-reporting patterns vary by urban status, at the classroom level. For almost all classrooms, adult reports are higher than children's self-reports (only 6 out of 198 observations are below the 45-degree line). The red line estimates a quadratic relation for mostly-rural classrooms (those with less than 20% of children living in urban areas), and the blue line, for mostly-urban classrooms (with over 80% of children living in urban areas). No mostly-urban classroom has a prevalence of children in employment greater than 20%. Within that lower range, however, under-reporting can be very high; for instance, the estimated relation predicts that if parents in these classrooms report a 10% prevalence of children in employment, the bias-adjusted prevalence is actually closer to 30%. For mostly-rural classrooms, under-reporting is predicted to be even the more striking the lower the prevalence in adult reports is. In these classrooms, a 10% reported prevalence would correspond to a nearly 45% bias-adjusted prevalence of children in employment. Naturally, in both cases, there is less room for under-reporting as prevalence according to adult reports increases. Next, Figure S17 documents the share of students reported to work over 10 hours a week at any point during the previous school year, according to both parents and teachers. Teachers were more conservative than parents when the question focused on employment during the school year, but less so if it focused on school holidays. Concretely, teachers identified child labor for only 2.8% of students during the school year (8%, according to parents) but for 12% of students during school holidays (10%, according to parents). Last, Figure S18 documents that, depending on reporting sources, the correlation between children in employment and educational outcomes at end line also systematically differs across treatment cells.


Figure S13: Validation of child labor measures using third-party data, households with all children in school




Notes: Columns show the share of children who worked at least an hour in cocoa fields over the previous month, according to children (in red), parents (in blue) and ENVERITAS (in green), along with 90% confidence bars. In the ENVERITAS data, we restrict attention to households whose *all* children between 6 to 11 years old were enrolled in school at the time of the survey. The figure reports the average prevalence for each region for which survey data overlaps with ENVERITAS data. Children answered the following question at end line: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question at end line: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. In the survey conducted by ENVERITAS during the harvest season (identified through satellite imagery), farmers answered the following question: “Do any of your children between 6 and 16 years old help you work on the cocoa farm?”. For the measures reported by children and parents, observations are restricted to the control group of the communication intervention. P-values from tests of proportions with unequal population variances (when children’s or parents’ reports are compared to ENVERITAS data; accounting for each source’s intra-cluster correlation computed at the regional level), and from tests of proportions with equal population variances (when comparing children’s and parents’ reports) through Ordinary Least Squares regressions. Sample sizes are as follows: (i) Children: 475; (ii) Parents: 475; and (iii) ENVERITAS: 340.

Figure S14: Validation of child labor measures using third-party data, households with at least some children in school

Panel A: Share of students who worked in cocoa plantations, by reporting source



Panel B: Share of students who worked in cocoa plantations, by source and region




Notes: Columns show the share of children who worked at least an hour in cocoa fields over the previous month, according to children (in red), parents (in blue) and ENVERITAS (in green), along with 90% confidence bars. In the ENVERITAS data, we restrict attention to households with *at least one child*


between 6 to 11 years old enrolled in school at the time of the survey. Panel A reports the average prevalence across all regions for which survey data overlaps with ENVERITAS data. Panel B breaks down prevalence by region. Children answered the following question at end line: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question at end line: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. In the survey conducted by ENVERITAS during the harvest season (identified through satellite imagery), farmers answered the following question: “Do any of your children between 6 and 16 years old help you work on the cocoa farm?”. For the measures reported by children and parents, observations are restricted to the control group of the communication intervention. P-values from tests of proportions with unequal population variances (when children’s or parents’ reports are compared to ENVERITAS data; accounting for each source’s intra-cluster correlation computed at the regional level, in Panel A), and from tests of proportions with equal population variances (when comparing children’s and parents’ reports) through Ordinary Least Squares regressions (clustering standard errors at regional level, in Panel A). Sample sizes are as follows: (i) Children: 475; (ii) Parents: 475; and (iii) ENVERITAS: 340.

Figure S15: Validation of child labor measures using third-party data, no sample restrictions

Panel A: Share of students who worked in cocoa plantations, by reporting source



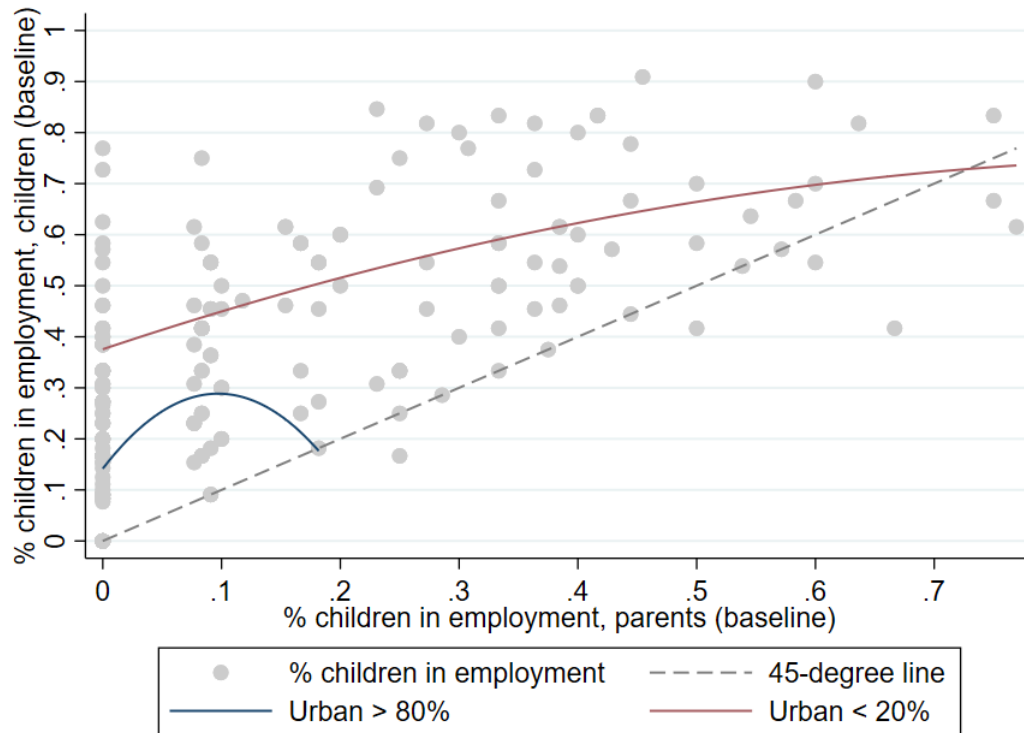
Panel B: Share of students who worked in cocoa plantations, by source and region



Notes: Columns show the share of children who worked at least an hour in cocoa fields over the previous month, according to children (in red), parents (in blue) and ENVERITAS (in green), along with 90% confidence bars. Panel A reports the average prevalence across all regions for which survey data overlaps

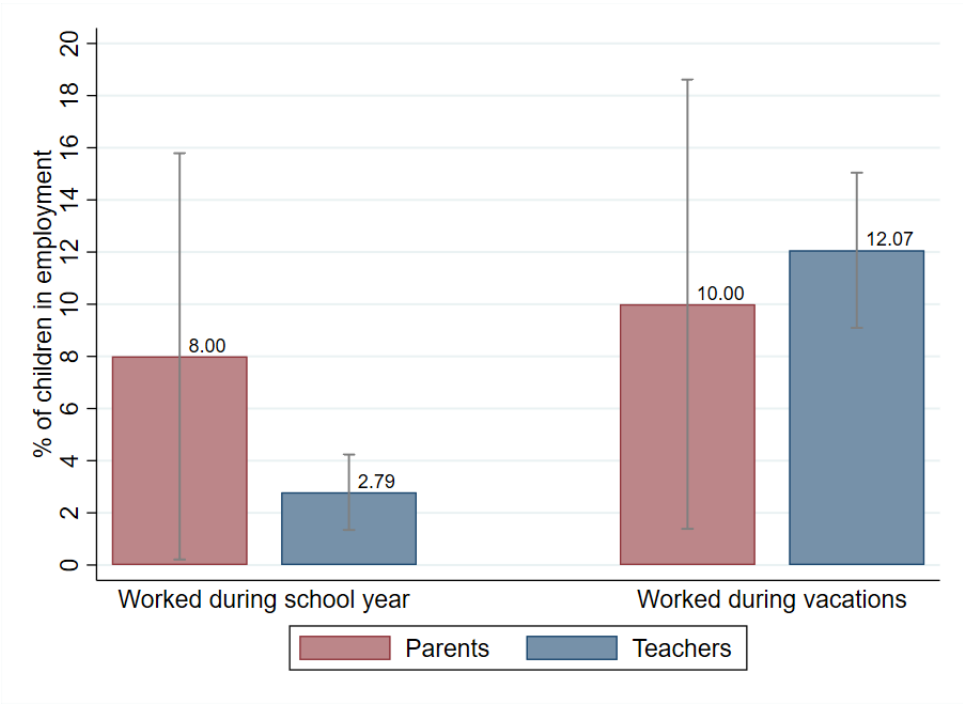
with ENVERITAS data. Panel B breaks down prevalence by region. Children answered the following question at end line: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question at end line: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. In the survey conducted by ENVERITAS during the harvest season (identified through satellite imagery), farmers answered the following question: “Do any of your children between 6 and 16 years old help you work on the cocoa farm?”. For the measures reported by children and parents, observations are restricted to the control group of the communication intervention. P-values from tests of proportions with unequal population variances (when children’s or parents’ reports are compared to ENVERITAS data; accounting for each source’s intra-cluster correlation computed at the regional level, in Panel A), and from tests of proportions with equal population variances (when comparing children’s and parents’ reports) through Ordinary Least Squares regressions (clustering standard errors at regional level, in Panel A). Sample sizes are as follows: (i) Children: 475; (ii) Parents: 475; and (iii) ENVERITAS: 395.

Figure S16: Correlation between parents' and children's answer according to urbanization level



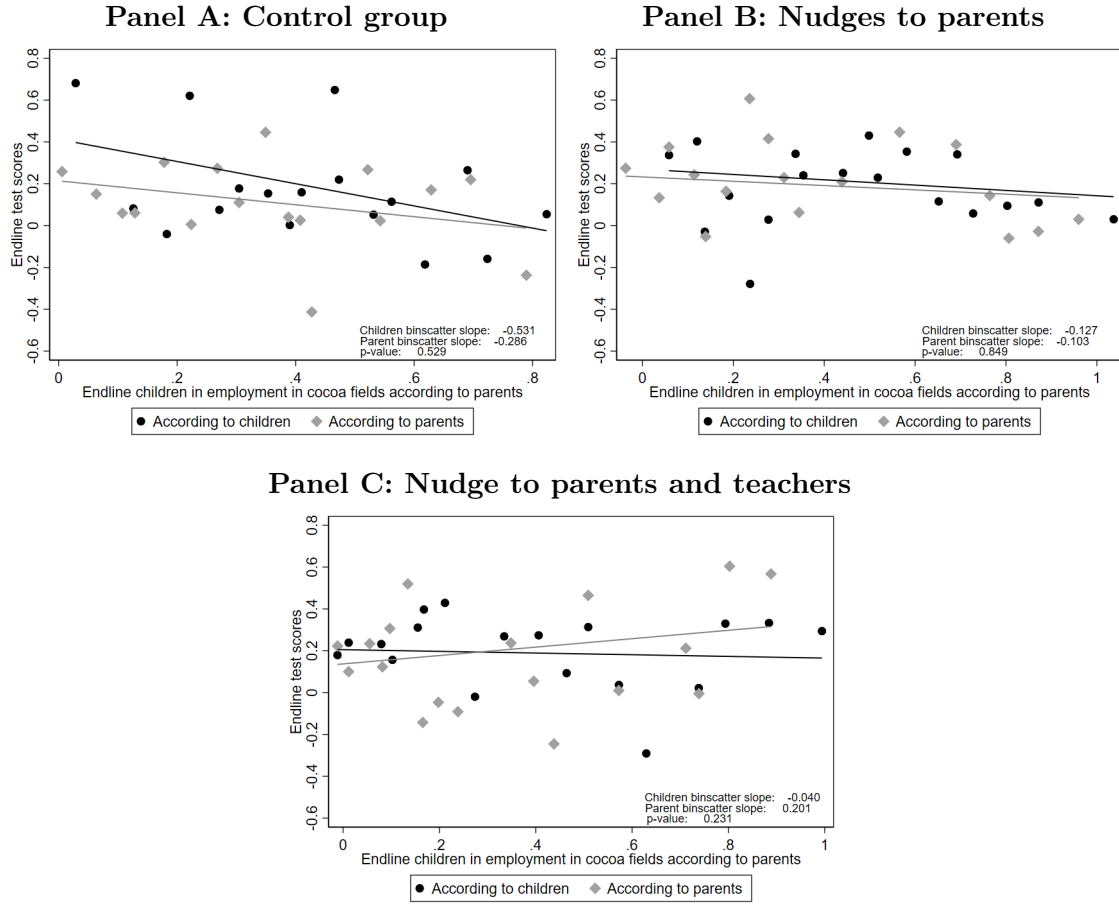
Notes: Data is aggregated at classroom level: average child labor according to children and parents were correlated according to composition on the classroom urbanization. “Urban > 80%” represents the quadratic fit between children’s and parent’s answer on child labor to classrooms that have more than 80% of their students in the urban category, while “Urban < 20%” represents the quadratic fit between children’s and parent’s answer on child labor to classrooms that have less than 20% of their students in the urban category, i.e., majority of students are from rural areas. Children answered the following question at baseline: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question at baseline: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Rural and urban areas are defined according to parents’ main occupation (agricultural or plantation activities are assigned to the former).

Figure S17: Share of students who worked in cocoa plantations at least 10 hours/week, during the school year and during the school holidays



Notes: Bars show the share of students who worked in cocoa plantations during different periods and according to two different sources. Parents (in red) and teachers (in blue) answered whether the child worked 10 hours/week or more in cocoa fields during a typical week (LHS bars) and whether the child worked 10 hours/week or more in cocoa fields during school holidays in the previous school year, as described in Table S1. All measures were collected at the follow-up surveys [15]. Across all bars, the sample is restricted to the control group. Sample sizes are the following: (i) Parents: 200; and (ii) Teachers: 2,500.

Figure S18: Correlation between end-line test scores and child labor, by source and treatment status



Notes: All panels show bin-scatter plots of end-line standardized test scores as a function of children in employment in cocoa fields at end line. Each panel estimates linear relation between the variables according to different measures of child labor; that based on children’s self-reports is shown in black, and that based on parents’ reports is shown in gray. Children answered the following question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Parents answered the following question: “I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, as described in Table S1. Standardized test scores are a summary measure of numeracy and literacy test scores (averaging across each component, normalized by their mean and standard deviation in the control group), following [17]. Panel A restricts the sample to the control group of the intervention; Panel B, to parents who randomly assigned to nudges (via text or or audio messages) in schools where teachers were not assigned to the intervention; and Panel C, to parents assigned to nudges (via text or audio messages) in schools where teachers were also assigned to nudges (via text messages); see [15].

Table S3: Robustness of treatment effects of nudges by reporting source and household remoteness, winsorizing outlier distance values

	Children in employment (self-report)	Children in employment (parent's report)		
	(1)	(2)	(3)	(4)
Panel A - Work in cocoa fields				
Nudge to parents	0.054 (0.040)	0.089* (0.045)	0.120** (0.056)	0.135** (0.065)
Distance to school (km)				-0.009 (0.035)
Nudge to parents \times Distance				-0.037 (0.058)
Sample mean	0.498	0.221	0.230	0.230
R-squared	0.115	0.091	0.094	0.095
Observations	2,246	2,176	1,608	1,608
Panel B - Domestic work				
Nudge to parents	0.046** (0.020)	0.051** (0.022)	0.066** (0.026)	0.065** (0.032)
Distance to School (km)				-0.042 (0.029)
Nudge to Parents \times Distance				0.004 (0.051)
Sample mean	0.878	0.789	0.794	0.794
R-squared	0.034	0.044	0.050	0.052
Observations	2,246	2,175	1,607	1,607
Controls (baseline survey)	Yes	Yes	Yes	Yes

Notes: In this specification, we winsorize outliers with respect of distance to school at the 99th percentile (where $distance_{99th\ percent} \approx 3.88$), hence, outliers values are substituted by the distance value at the 99th percentile. In column (1), Children in employment (self-report) = 1 if the child answered affirmatively to the working question: “In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation” in Panel A, “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. In columns (2) to (4), Children in employment (parent's report) = 1 if the parent answered affirmatively to the question “In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation”, in Panel A, and “Domestic work, such as buying food or cooking, cleaning the house, do the laundry, take care of children or other sick or old relatives”, in Panel B. Distance from school (km) is the linear distance between the home and their child's school GPS locations, in kilometers; Nudge to parents = 1 if the parent (but not teachers) was assigned to nudges, and 0 otherwise. Baseline controls include treatments arms, child gender, grade indicators, standardized test scores (averaged across numeracy and literacy); and summary measures of parental engagement, student effort, socio-emotional skills, working memory, visual attention, impulsivity, self-esteem, and mindset (see [14]. Summary measures computed following [17], standardizing each component by normalizing values by the mean and standard deviation of the control group at the baseline survey within each grade, and then averaging over all standardized components. In columns (3) and (4), sample restricted to observations with valid home GPS coordinates. All regressions estimated through Ordinary Least Squares with standard errors clustered at the classroom level.

*** p<0.01, ** p<0.05, * p<0.10

Predicting the bias-adjusted prevalence of child labor

Informed by the statistical equivalence between ENVERITAS data (treated as the ground truth) and children's reports of the prevalence of children in employment, we can estimate the relation between adult reports (the ILO standard measure) and children's self-reports not only to compute the extent of under-reporting in our data, but also to predict what accurate reports would have been in geographical units without ENVERITAS data. In this Appendix, we show how estimating this *bias-adjustment factor* for children in employment allows us to compute the bias-adjusted prevalence of child labor for Côte d'Ivoire and Ghana, under the assumption that this relation is the same conditional on the local characteristics we can measure. We focus on these two countries because 2/3 of the global cocoa production originates within their borders and because our surveys asked about children in employment specifically in cocoa farms.

According to the ILO definition:

$$CL_i = CE_i \times HC_i, \quad (1)$$

where CL_i is the share of child labor in country i , CE_i is the share of children in employment in country i , and HC_i is the share of children in employment in country i who work in hazardous conditions ($= 1$ for children under 12, and $= \%$ of those working long hours and/or under heavy or dangerous work conditions, otherwise).

Once we determine that children's self-reports are accurate and, as such, can be used as the ground truth, in our data we observe both CE_j , the share of children in employment reported by parents in classroom j , and CE_j^* , the ground truth for children in employment in classroom j . We are interested in predicting CE_i^* for all countries $i \in I$. To do that, we first estimate:

$$CE_j^* = f(CE_j, X_j) + \varepsilon_j, \quad (2)$$

where $f(CE_j, X_j)$ is a function of children in employment reported by parents in classroom j and other classroom characteristics X_j , and ε_j is an error term.

We then use our estimates to compute:

$$\hat{CE}_i^* = \hat{\alpha}_i \times CE_i, \quad (3)$$

where $\hat{\alpha}_i = \frac{\hat{f}(CE_i, X_i)}{CE_i} \geq 1$.

Those same estimates can be used to recover CL_i^* , the ground truth for child labor in country i . This is thanks to the fact that HC_i is free of bias (since $= 1$ for 5-11 year olds, and elicited directly from 12-14 year-old children), and to the multiplicative nature of its definition:

$$CL_i^* = CE_i^* \times HC_i = \hat{\alpha}_i \times CE_i \times HC_i = \hat{\alpha}_i \times CL_i \quad (4)$$

Concretely, we estimate a quadratic polynomial, allowing the extent of under-reporting to vary with both the local level of children in employment reported by parents and with the local urbanization rate:

$$CE_j^c = \sum_{k=0}^2 \beta_k (CE_j^p)^k + \sum_{k=0}^2 \gamma_k \left[(CE_j^p)^k \times urban_j \right] + \varepsilon_j, \quad (5)$$

where CE_j^c is % child labor according to children in classroom j , CE_j^p , that according to parents in classroom j , and $urban_j$ is the % of students from urban areas in classroom j .

We only estimate heterogeneity with respect to children in employment reported by adults and urbanization rates because these are available both in our data and in the World Development Indicators. Other variables, like annual per capita income, are not available in our data.

Last, we compute $\hat{\alpha}_i = \max \left\{ 1, \frac{\hat{CE}_i^c}{CE_i^c} \right\}$, where \hat{CE}_i^c is the predicted value of CE^c for country i using equation (5). As the formula indicates, we constrain the estimated bias-adjustment factor (as well as its confidence interval) to be greater or equal to 1. In practice, this constraint does not affect Côte d’Ivoire or Ghana.

World Development Indicators track child labor for 97 countries (those where the issue is considered to be relevant by the data collection organizations, such as UNICEF country offices), focusing on 7-14 year-old children. With the bias-adjustment factor that we predict for Côte d’Ivoire and Ghana using the procedure described above, we compute country-level bias-adjusted prevalence of child labor and number of child workers (and accompanying 95% CIs) using data from the most recent year available for each country in the World Development Indicators. To arrive at the bias-adjusted prevalence figures, we have to rely, in addition, on World Development Indicators data on the number of 0-14 year old children by country, and subtract these from the UNICEF figures for under-5 children by country. For both population counts, we use 2020 data – assuming that child labor indicators remained constant since their most recent measurement.

Table S4: Relation between children's and parents' reports at the classroom-level

	(1)
Child labor according to parents	0.463 (0.294)
Urban	-0.361*** (0.056)
(Child labor according to parents) ²	-0.069 (0.439)
(Child labor according to parents)*Urban	0.850 (0.633)
(Child labor according to parents) ² *Urban	-1.491 (1.270)
Constant	0.477*** (0.041)
Observations	198
Adjusted R^2	0.570

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Estimation based on the following equation: $CL_c = \alpha + \beta_1 CL_p + \beta_2 Urban + \beta_3 CL_p^2 + \beta_4 CL_p * Urban + \beta_5 CL_p^2 * Urban + \varepsilon$, in which: CL_c is child labor according to children, CL_p is child labor according to parents, and $Urban$ the percentage of students living in urban areas. Results are for classroom-level averages. Rural and urban areas are defined according to parents' main occupation (agricultural or plantation activities are assigned to the former). Child labor according to parents corresponds to the answer on the following question "I will now ask you some questions about activities that your children might have recently performed. In the last month, has any of your children engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation", as described in Table S1. Child labor according to children stands for the baseline share of students who report to have worked in cocoa plantations in the last month for one hour or more, in response to the following question: "In the last month, have you engaged in one or more of the following activities, for one hour or more? Work in a cocoa plantation", as described in Table S1. All variables correspond to baseline answers.

Table S5: Bias-adjusted prevalence of child labor by country

Country	Prevalence			Children in employment		
	WDI*	Bias-adjusted	95% CI	WDI**	Bias-adjusted	95% CI
Côte D'Ivoire	36.50%	50.75%	[44.80%; 56.71%]	2,488,650	3,460,540	[3,054,542; 3,866,538]
Ghana	28.80%	46.67%	[40.74%; 52.61%]	2,122,198	3,439,275	[3,001,906; 3,876,644]

Notes: * using the latest available WDI figures, following [18]; ** using WDI and UNICEF population figures for 0-14 and under-5 children for 2020. Bias-adjusted prevalence of child labor based on Table S4's estimates for country-specific adjustment factors.