Survey attrition after 15 years of tracking children in four developing countries: the Young Lives study[1]

Alan Sánchez[2] Javier Escobal[3]

September 20, 2019

Data availability statement: The data used for this analysis come from the five waves of data collection of the Young Lives Study. The data are publicly available in the UK Data Service at http://doi.org/10.5255/UKDA-SN-5307-3, reference number 5307 (Round 1), http://doi.org/10.5255/UKDA-SN-6852-3, reference number 6852 (Round 2), http://doi.org/10.5255/UKDA-SN-6853-3, reference number 6853 (Round 3), http://doi.org/10.5255/UKDA-SN-7931-2, reference number 7931 (Round 4), and http://doi.org/10.5255/UKDA-SN-8357-1, reference number 8357 (Round 5)

[1] We are very grateful to Juan Ferrer and Alessandra Hidalgo for excellent research assistance. We also thank Marta Favara, Jere Behrman and two anonymous referees for providing detailed comments on an earlier version of this paper. GRADE is one of the partners of the Young Lives Study. We would like to state that GRADE, and the authors in particular, have been involved in the data collection process of the Young Lives Study in Peru.
[2] Grupo de Análisis para el Desarrollo. Corresponding author. Postal address: Av. Almirante Grau 915. Barranco, Lima 15063 - Perú. E-mail: asanchez@grade.org.pe
[3] Grupo de Análisis para el Desarrollo (Lima, Perú).

Abstract

Young Lives (YL) is a multi-country, birth cohort study that over a period of 15 years followed close to 8,000 individuals born in 2001-2 and 4,000 individuals born in 1994-5 in Ethiopia, India (states of Andhra Pradesh and Telangana), Peru and Vietnam. This study aims to document attrition in the YL samples, to identify the correlates of attrition, and to test for attrition bias. The cumulative year-to-year attrition rate is 0.5% and 0.8% for the younger and older cohort, respectively, among the lowest attrition rates in longitudinal studies in low-and-middle-income countries. Attrition rates vary by household wealth, area of residence, and ethnicity/caste across countries. Attrited individuals are more likely to be poorer in the Peruvian sample and wealthier in the other country samples. When analyzing a set of nutritional and cognitive outcomes, there is suggestive evidence of attrition bias, part of which is driven by child mortality. Even though attrition is very low, our findings highlight the importance of controlling for household socio-economic characteristics when performing statistical analysis of the YL samples. The study also highlights key insights that can be helpful to reduce attrition or ameliorate its effects in other longitudinal studies.

1. Introduction

In developing countries, birth cohort studies –those that track a group of individuals born in a given year or period, sometimes since birth or early childhood– remain scarce.[1] The Young Lives (YL) study is the largest existing multi-country, birth cohort, longitudinal study focused on low-and-middle-income countries. Between 2002 and 2016, it has collected data for approximately 12,000 individuals in four countries: Ethiopia, India (states of Andhra Pradesh and Telangana), Peru and Vietnam. Data from YL has been used in more than two-hundred peer-reviewed papers.[2] However, after 15 years of data collection, there is little evidence about the nature of attrition in the YL samples.[3]

Longitudinal studies try to keep attrition rates low in order to retain statistical power and to reduce the possibility of attrition bias. Attrition bias might occur under certain conditions if attrited and non-attrited individuals are different and no further measures are taken. Previous studies that use follow-up data from birth cohorts or household panels in developing countries show that while it is common that attrited and non-attrited observations differ in their characteristics, this does not necessarily lead to attrition bias (Alderman et al., 2000; Falaris, 2003; Dercon and Outes-Leon, 2009). However, an absence of attrition bias cannot be established *a priori*.

Within this context, this document has three aims. First, to document the attrition rates observed in the YL samples and their evolution between 2002 and 2016. Second, to identify which individual and household characteristics predict attrition, and how results vary across country samples. Third, to establish whether the attrition observed lead to bias estimates of key outcomes.[4] We expect our results to contribute to the scarce literature on the nature of attrition bias in birth cohort studies in developing countries (Alderman et al., 2000; Thomas et al., 2001; Falaris, 2003; Norris et al., 2007; Dercon and Outes-Leon, 2008), and to inform

[1] According to the Institute of Fiscal Studies, by 2018 there were 175 longitudinal studies operating in the developing world, from which only 12 are birth cohort studies following members since birth or childhood according to our analysis.
[2] On the YL website, 204 journal articles that use YL data appear listed.
[3] Dercon and Outes-Leon (2008) used data from the first two rounds of the YL study to document attrition patterns and found no evidence of attrition bias. While this is reassuring, after three additional rounds of data collection attrition levels must have increased and attrition patterns might have changed over time.
[4] The data used for this study is publicly archived at the UK Data Service. See Jones (2018), Boyden (2018a and 2018b), Woldehanna et al. (2018), and Sánchez et al. (2018).

researchers that use YL data to investigate child development in low-and-middle-income countries.

The rest of the document is structured as follows. Section 2 provides a summary of the design of the YL study, and an explanation of fieldwork procedures. Section 3 reports YL attrition rates by country sample and cohort, and Section 4 shows graphical evidence of their evolution over time. Sections 5 and 6 explain the methodology and data used, respectively. Section 7 reports the results obtained, and Section 8 reports our conclusions and suggestions.

2. Young Lives Study

General design

YL is a multi-country, birth cohort, longitudinal study designed to follow children born into poverty (and some born non-poor) over a period of 15 years with the objective of tracing the causes and consequence of childhood poverty (Harpham, 2002; Barnett et al, 2013). Countries from different continents were chosen to reflect the diversity of the developing world: Ethiopia, the state of Andhra Pradesh in India (currently split into the states of Andhra Pradesh and Telangana), Peru and Vietnam. Two birth cohorts were selected in each country: a younger cohort (YC) born in 2001-2, and an older cohort (OC) born in 1994-5. The first visit took place in 2002, when children in each cohort aged (approximately) 1 and 8 years old, respectively. More details about the sample design are explained below.

The YL study differs from a standard household panel, for several reasons. First, the unit of interest is the child, and only one per family was selected. From here onwards, we use the acronym “YLI” (which stands for YL individual) to refer to this person. The focus on YLI means that, were her to become part of a new family, the previous family was no longer followed; and, if YLI died, her family was no longer followed. Second, in case of moving/migration, the protocol was to follow YLI as long as she remained within the national territory (or within the state boundaries in the case of Andhra Pradesh in India). Third, YL does not have refreshment samples, as it is common practice in household panels.

Context of the study countries

Some key characteristics of the selected countries are reported in **Table 1**. At the beginning of the Millennium, Ethiopia, India and Vietnam were classified as low-income countries by the World Bank, whereas Peru was classified as a middle-income country. Consistent with this, in 2004 Peru had a relatively low level of monetary poverty compared to Ethiopia. Another key difference was the level of urbanization: in 2002, in Peru, around 3 out of 4 lived in urban areas versus 1 out of 4 in the other countries. Despite these differences, in 2002 Vietnam reported the highest level of educational attainment with Peru very close.

Table 1: country characteristics

Country	Rural population (%)		Poverty rate (%) *(USD 1.90 USD per day, PPP 2011)*				Gini coefficient	Primary completion rate (%)	
	2002	2016	1999	2004	2010-11	2015-16	2010-11	2002	2015
Ethiopia	85	80	56	36	34	27	33	31	54
India	72	67		38	21		35	76	98
Peru	26	21	18	14	6	4	46	97	99
Vietnam	74	66		27	4	3	39	101	104

Source: World Bank Development Indicators.

Between 2002 and 2016 the four countries enjoyed high rates of economic growth (9.2 per cent in Ethiopia, 7.4 per cent in India, 5.6 per cent in Peru and 6.4 per cent in Vietnam, according to the World Bank Development indicators). Aligned with this, by 2016 Vietnam and Peru had been re-classified as upper-middle income countries and India as lower-middle income. Despite an exceptional reduction in monetary poverty during the study period, Ethiopia remains as a low-income country. Among the four countries, Peru appears as the most unequal according to the Gini coefficient (using information from 2010-11).

Sample design

Sample design in each country was guided by three principles (Wilson et al., 2006). First, there would be two samples in each country, one with 2,000 YLIs aged between 6 and 18 months (YC), and another one with 1,000 YLIs aged between 7 and 8 years (OC), in the first visit. Second, YLIs were to be sampled from 20 sentinel sites (from here onwards, clusters), in turn selected to reflect the geographic, ethnic and socio-economic diversity of each selected country. In each sample, 100 were to be selected for YC, and 50 for OC. Third, each country sample had to be composed mainly (but not uniquely) of YLIs from poor households. In other words, poor households had to be over-sampled. Samples were not required to be nationally representative, but they had to be informative of the diversity of living conditions in each country. Following these principles, in Ethiopia, India, and Vietnam, a purposive (i.e., non-probabilistic) approach was used to select clusters, whereas in Peru they were randomly selected. Specifically, in Ethiopia, India and Vietnam, clusters (*kebeles*, *mandalas* and *communes*, respectively) were selected using a purposive approach oriented to reflect the geographic, ethnic and socio-economic diversity, to have a mixed of urban/rural areas and poor/non-poor areas, and to over-sample the poor. In the case of Peru, 20 clusters (or *districts)* were chosen randomly from the universe of *districts*, excluding the wealthiest 5%. After the clusters were selected, in each country a specific area of the cluster was randomly selected to start the process of data collection. More information about the sample design followed in each country is reported in **Table A.1** in the Online Annex. For technical details, see Outes and Sánchez (2008), Kumra (2008), Escobal and Flores (2008), and Nguyen (2008).

Given the fact that the samples are not statistically representative of a specific population on each country, it follows that the attrition patterns evaluated in this paper are associated to the potential differences in the relationships between observable variables and outcomes appearing in the original samples and those relationships appearing in the follow-up samples after years of data collection. As such, this analysis does not pretend to infer to any relationships in the country population for which the YL study does not claim any statistical representativeness.

Fieldwork procedures for enrollment and follow-up

For the enrollment of participants, a door-to-door screening phase took place in all countries to search for YLIs from the expected age range (Alemu et al., 2003; Escobal et al., 2003; Galab et al., 2003; Tuan et al., 2003). Contact took place through the main caregivers of prospective participants, whom were told that participation was voluntary and informed of the following. First, the study would last 15 years, with visits spaced between 3 and 4 years. Second, in the first visit a survey would be administered to collect information about: (i) socio-economic aspects of the family; (ii) health, education and time-use of YLI; anthropometric measures of YLI; and, cognitive achievement of YLI (for OC). Third, there would be no monetary incentive or compensation to be part of the study. Fourth, it would be possible to drop from the study at any point in time. Those caregivers that agreed to participate under these terms were asked to sign a consent agreement (oral agreement if the person was illiterate). In the case of OC, YLIs were also asked for assent. Country teams did not report problems with refusals in the enrolment phase.[5]

For the follow-up, fieldwork teams were asked to track YLIs by all means possible, as long as they were within the country borders (or within the borders of Andhra Pradesh and Telangana in the case of India). For illustrative purposes, the follow-up process is exemplified in **Figure 1**, for the most common case in which YLI is a minor. At the beginning of round *x*, the team of enumerators received the contact information (addresses and phone numbers) of all YLIs that were alive, and whose caregivers agreed to continue as part of the study in the previous round. For simplicity, consider the case of round 2 (x=2), when all original members were tracked. For a given YLI, in round 2 enumerators searched for either the main caregiver or another adult family member. If none of them was found, YLI was classified as not located; if one of them was found, several possible outcomes arise, as depicted in the diagram. Only when consent for interview was obtained YLI is considered as non-attrited; in all other cases (not located, died, moved to another country, refused to be interviewed this time or continue as part of the study) YLI is classified as attrited. Attrited participants in round *x* might become non-attrited in subsequent rounds, because the fieldwork team is asked to continue searching for all YLIs as long as they are known to be alive and did not refuse to continue as part of the study. In other words, it is possible to re-enter the sample.

[5] Vietnam reported 36 (out of 3000) refusals (Tuan et al., 2003); India reported 14 refusals, presumably out of 3000 (Galab et al., 2003); Ethiopia declared no major problem with refusals (Alemu et al., 2003).

INSERT FIGURE 1 HERE

The process sketched in **Figure 1** reflects the sources of attrition when the participant was a minor during data collection. If YLI was an adult (which was the case in the last two rounds for OC), the protocol varies according to whether YLI depends economically from her original family (or from a new family that has ties with the original family), is economically independent, and/or has formed a new family. In the last two cases, only consent from YLI was required to proceed with the interview.

3. Attrition rates in the Young Lives study

For this analysis, attrition in round *x* is defined as the ratio of total number of YLIs that died, were not located, moved to another country, or for whom consent was not obtained, relative to the total number of YLIs in round 1. Following this definition, **Table 2** (Column (6)) reports the cumulative attrition rates observed in the fifth visit in 2016 (round 5). The average attrition rate is 6.8% in YC and 12.6% in OC. To facilitate comparison with other longitudinal studies, this equates to 0.5% and 0.8% per year, respectively. For YC, the smallest attrition rate is observed in Vietnam (3.1% in total, 0.2% per year) and the largest in Peru (9.4% in total, 0.6% per year). For OC, the smallest rate is observed in India and Vietnam (8.5%, 0.6% per year) and the largest in Ethiopia (18.6%, 1.2% per year).

Our definition of attrition considers dead children as attrited.[6] As a reference, if dead children are excluded from the calculation (Column (7)), average attrition rates reduce to 4.9% in YC and 11.8% in OC (0.3% and 0.8% per year, respectively). The reduction in attrition is most notorious for YC, which is consistent with the fact that infant mortality is higher than mortality in other periods of childhood. The largest contribution of mortality to attrition (in percentage points) is observed for YC in Ethiopia, whereas the smallest contribution is observed in Vietnam. This pattern is consistent with the poverty and educational rankings depicted in **Table 1**.

[6] Although considering those YLI that died part of the attrited sample is open to debate, we do so because it clearly affects our ability to infer on the relationships between observables and outcomes for the original sample. Obviously, the reasons behind the death of a YLI are different from the reasons why a YLI may migrate out of the country or been part of a family that refuses to continue in the study.

Table 2: attrition rates in the round 5 of the Young Lives country samples

Block A: ***younger cohort***

	Original sample (1)	Gave consent in round 5 (2)	Died (3)	Refused/ Not located (4)	Total attrited in round 5 (5)	Attrition rates in round 5 (6)		Attrition rates without dead YLIs (7)	
						Total	By year	Total	By year
Ethiopia	1999	1812	85	102	102	9.4%	0.6%	5.4%	0.4%
India	2011	1900	43	68	68	5.5%	0.4%	3.5%	0.2%
Peru	2052	1831	25	167	167	9.4%	0.6%	8.2%	0.6%
Vietnam	2000	1938	12	48	48	3.1%	0.2%	2.4%	0.2%
Total:	8062	7481	165	385	385	6.8%	0.5%	4.9%	0.3%

Block B: ***older cohort***

	Original sample (1)	Gave consent in round 5 (2)	Died (3)	Refused/ Not located (4)	Total attrited in round 5 (5)	Attrition rates in round 5 (6)		Attrition rates without dead YLIs (7)	
						Total	By year	Total	By year
Ethiopia	1000	814	11	175	175	18.6%	1.2%	17.8%	1.2%
India	1008	922	19	67	67	8.5%	0.6%	6.9%	0.5%
Peru	714	597	6	100	100	14.8%	1.0%	14.1%	0.9%
Vietnam	1000	910	5	85	85	8.5%	0.6%	8.5%	0.6%
Total:	3722	3243	41	427	427	12.6%	0.8%	11.8%	0.8%

Source: reports from Young Lives country teams and own calculations. In Column (3), death reports correspond to information collected between rounds 2 and 5. In Column (4), YLIs that refused or that were not located were typically classified as such based on fieldwork efforts made in round 5, except for cases of families that requested to be excluded from the study at some point between rounds 2 and 4. In columns (7), attrition rates reported exclude dead YLIs from both the numerator and the denominator.

Attrition rates appear to be low, more so considering the challenges involved in collecting longitudinal data in low-income country settings. To put results in context, we proceed to compare the attrition rates observed in the YL country samples with those observed in other birth cohort, longitudinal studies that follow children from birth or early childhood in other developing countries. To this end, we used as a reference information collected by the Institute of Fiscal Studies (IFS). IFS have gathered information on 175 longitudinal studies in the developing world, though few of them are cohort studies. We were able to identify 12 similar studies.[7] Since two of them do not report attrition rates, we only report results for 10 studies. **Figure 2** summarizes the average year-to-year attrition rates of the selected studies and of the YL country samples (for more information, see **Table A.2** in the Online Annex). Annualized attrition rates in the literature range between 0.8% and 3.0% (per year), with an average of 1.3%. Comparatively, the attrition rates observed in the YL country samples for YC are the lowest in the group. Although OC is not directly comparable to these studies because children from these cohorts are not tracked since early childhood, even in this case attrition rates are at the lower end of longitudinal studies in the developing world.

[7] These studies fulfill three conditions: (i) have tracked children since infancy; (ii) have lasted at least 5 years; and, (iii) have reported results from follow-up studies within the last 15 years.

INSERT FIGURE 2 HERE

4. Cumulative attrition rates over time

YLIs were re-visited in 2006 (round 2), 2009 (round 3), 2013 (round 4) and 2016 (round 5). **Figure 3** shows the evolution of the cumulative attrition rate across rounds, separately for each cohort. Results do not change in any substantive way if dead children are excluded from the definition of attrition (see **Figure A.1** in the Online Annex).

Cumulative attrition rates have increased over time, with the only exception of Vietnam between rounds 4 and 5.[8] For OC, it is evident that the speed at which attrition grew increased after round 3 in all country samples, except for India. This pattern, which is not observed for YC, is likely associated with the age of the cohort, which reached the majority of age in round 4. This matters for several reasons. First, when becoming adults, YLIs may have felt empowered to decide whether to remain in the study –which impact on refusal rates is ambiguous. Second, some YLIs became economically independent and/or formed new families (had children, got married/cohabitated, or both). One consequence of this is that the length of the questionnaire increased for them –compared to a previous scenario where they only had to answer the individual questionnaire, in this case they would also have to answer the household questionnaire. This is likely to increase refusal rates. Third, the transition to adulthood likely increased the probability of migration, which increases the difficulty to locate study participants.

In **Figure 4**, cumulative attrition rates over time are reported separately for YLIs that were in the bottom-tercile of wealth during the first round, relative to those in the top-tercile. For YC, YLIs from the top-tercile report higher attrition rates (on average) in the Ethiopian, Indian and Vietnamese samples. A similar pattern is observed for OC, with differences being larger in this case. In contrast, for the Peruvian sample higher attrition rates (on average) are observed for YLIs from the bottom-tercile in OC –in YC, differences are not discernable. The differences in attrition rates between YLIs from poorer and wealthier households may be the effect of several factors. On the one hand, one might expect lower mortality rates in the top-tercile. On the other hand, if it is more likely for those that are better off to migrate, mobility

[8] For this country sample, a decrease in attrition is observed between rounds 4 and 5, which implies that some YLIs not interviewed in round 4 re-entered the sample in round 5.

could generate higher attrition rates in the top-tercile. Opportunity costs of YLIs and their families may generate also that the continuing participation rate may be different across the wealth spectrum.

Figure 5 depicts the evolution of cumulative attrition rates by area of residence. For YC, in all country samples attrition rates are higher (on average) for YLIs that were selected from urban areas. This is consistent with findings from the literature that urban households are more likely to disappear from longitudinal studies (Rayan and Petesch, 2007). The fact that the study does not follow-up YLIs that move to other countries (or other states in India), which are more likely to be urban, may also contribute to higher attrition rates for urban areas. The same pattern is observed for OC, with the only exception of the Peruvian sample, where in fact the opposite is observed –attrition rates are higher for YLIs from rural areas.

Finally, **Figure 6** depicts the temporal evolution of the cumulative attrition rates by gender. For YC no relevant differences arise. In the case of OC, attrition rates are higher for males in Vietnam and Peru, and higher for females in the case of Ethiopia. Transition into adulthood may be related to the exploration of educational and labor opportunities away from original place of residence; and these opportunities tend to be different for men and women (Heckert, 2015; Valentine et al., 2017).

INSERT FIGURE 3 HERE

INSERT FIGURE 4 HERE

INSERT FIGURE 5 HERE

INSERT FIGURE 6 HERE

5. Methodology

Attrition rates in the YL samples are low by international experiences, especially in YC, but they are not negligible (cumulative attrition in round 5 can go up to 18% in OC for Ethiopia). Therefore, it is important to establish if there is selective attrition and, in that case, if it leads to biased estimates of key outcome indicators. While it can be expected that selective attrition leads to biased estimates, the two phenomena do not need to concur. A set of studies that analyze attrition patterns in longitudinal data from the U.S. find that despite attrition is high and correlated with observable characteristics it does not lead to biased estimates. A similar conclusion is reached by Alderman et al. (2000) after analyzing longitudinal data from Bolivia, Kenya and South Africa, and by Falaris (2003) using longitudinal data from Peru, Cote d'Ivoire and Vietnam.

Following Fitzgerald et al. (1998), the concept of attrition bias can be linked to the concept of sample selection bias (Heckman, 1979). For clarity of exposition, we follow a terminology proposed by Alderman et al. (2000) –which derives from Fitzgerald et al. (1998) but that is closer to the statistical literature— and consider the following types of attrition: missing completely at random, attrition that exhibits selection on unobservables, and attrition that exhibits selection on observables. Consider A_t as an indicator of attrition that takes the value of 1 if an observation is missing due to attrition, and 0 otherwise. The model of interest is as follows:

$$y_t = \beta_0 + \beta_1 x_t + \varepsilon_t \qquad (1)$$

where y_t is an outcome of interest for a researcher (for instance, child´s height-for-age) observed if $A_t = 0$ (i.e., for the non-attritors), x_t is a variable (or vector) observed for both attritors and non-attritors (for instance, a set of variables observed in the first visit), and ε_t is the error term with mean zero. Now consider the following linearization of the attrition process:

$$A_t^* = \delta_0 + \delta_1 x_t + \delta_2 z_t + \mu_t \qquad (2)$$

where A_t^* is a latent variable such that $A_t = 1$ if $A_t^* > 0$ and $A_t = 0$ if $A_t^* \leq 0$; μ_t is the error term with mean zero; and, z_t is an auxiliary variable (or vector) observed for both attritors and non-attritors that is not part of x_t. There are several alternatives for z_t, including, for instance, characteristics of the interviewer and/or of the participant that might be correlated with attrition but that are not of interest in equation (1).

The case of attrition that exhibits selection on unobservables occurs when $\mu_t|x_t$ is not independent from $\varepsilon_t|x_t$. This means there are unobserved factors that explain attrition and that simultaneously explain the outcome of interest. Solving attrition bias of this nature is challenging due to the exclusionary restriction required (i.e., a suitable variable z_t such that $z_t|x_t$ is independent from $\varepsilon_t|x_t$). For simplicity, from here onwards we assume that $\mu_t|x_t$ is independent from $\varepsilon_t|x_t$.

Attrition due to selection on observables might not generate bias (*ignorable* selection on observables) if one of the following conditions hold: (a) $z_t|x_t$ is independent from $\varepsilon_t|x_t$; or, (b) δ_2=0. Moreover, if $\delta_2 = \delta_1$= 0 there will be completely at random attrition. On the other hand, if neither (a) nor (b) hold, there will be *non-ignorable* selection on unobservables. When this happens, information from $f(y_t|x_t, A_t = 0)$ is not enough to recover information from $f(y_t|x_t)$.

For instance, suppose that the outcome of interest is height-for-age. A child's height-for-age is expected to be negatively correlated with the altitude of the place in which the child was born (Pomeroy et al., 2014), however altitude might not be part of the model of interest for the researcher. At the same time, it is plausible that altitude is correlated with attrition under the assumption that communities located at higher altitudes are more difficult to reach. The combination of these two circumstances would lead to (*non-ignorable*) selection on observables. In this case, suitable alternatives need to be considered. Fitzgerald et al., (1998) and Alderman et al. (2000) show that using weighted least squares regressions, with the expression $\Pr(A_t = 0|x_t)/\Pr(A_t = 0|z_t, x_t)$ as weights, allow to recover the parameters from $f(y_t|x_t)$. Other alternatives suggested in the wider literature include the use of fixed effects estimates to absorb the effects of attrition (Verbeek and Nijman, 1996; Ziliak and Kniesner, 1998), and multiple imputation methods (Schafer and Graham, 2002).

In our analysis, we are interested in two aspects. First, to identify which individual and household characteristics predict attrition. Second, to test if there might be attrition bias due to *non-ignorable* selection on observables in the YL samples. For the first objective, we use t-tests to establish whether attrition varies according to individual and household characteristics that would typically be considered in vector x_t for models of determinants of health, nutrition, and educational outcomes. For the second objective, we rely on two strategies. First, we estimate models that regress A_t on x_t and z_t, and then test the importance of z_t (attrition probit tests). Following Fitzgerald et al. (1998) suggestion that a lagged value of y_t can play the role of z_t, we use y_1 as a proxy for z_t. Second, we consider the so called BLGW test proposed by Becketti et al. (1998). The objective of this test is to verify whether the coefficients of models to predict an outcome of interest vary between attrited and non-attrited participants. This is verified by including an interaction between all the selected predictors for the child outcome (included in vector x_t) and whether the child is attrited. The joint significance of these interactions is then assessed by means of a F-Test.

6. Data

We use data from the first visit of the YL study combined with an attrition indicator that takes the value of 1 if YLI is attrited in the last visit (round 5), 0 otherwise. The following individual and household characteristics measured in the first visit are considered as potential correlates of attrition: sex of YLI; whether YLI is the first born (OC only); caregiver basic demographic characteristics (sex, age, and marital status); household characteristics (household wealth, area of residence and caregiver´s education); caregiver´s ethnicity or caste, and household size. Household wealth is measured considering the location of the household across terciles of the distribution of the wealth index.[9] As for ethnicity/caste, in Ethiopia we consider the caregiver´s region of origin (Amhara, Gurage, Hadiva, Oromo, Sidama, and Tigrian); in India, whether the caregiver belongs to one of the following castes: Schedules Caste, Scheduled Tribe, Backward Class, and Other Class; in Peru, whether Spanish is the caregiver's native tongue –as opposed to Quechua, Aymara and other dialects; in Vietnam, whether the caregiver belongs to Kinh, the main ethnic group –as opposed to other ethnic origins (Chinese, Tay, H'Mong, Nung, Ede, Thai, Dao and Giay).

As for YLI outcomes, we use those available in round 1. For nutritional status, we choose three indicators based on anthropometric data that are available for YC and OC (at ages 1 and 8, respectively): height-for-age Z-scores, weight-for-age Z-scores and stunting. Z-

[9] The wealth index is a composite measure that comprises information related to household access to services (electricity, drinking water, toilet facilities), housing quality (main material used for floor, wall and roof) and durable goods (including TV, radio, fridge, and computer, among others).

scores are based on the reference growth curves of the World Health Organization (WHO). Stunting is defined as being two standard deviations below the height of the median child of a given age and sex according to the latest WHO standard definition, and is associated with chronic undernutrition. For educational achievement, we focus on OC only and look at school enrolment and literacy indicators at age 8. We use a reading task that tests the child's ability to read letters, a simple word and a simple sentence; and a writing task that determines a child's ability to write a simple sentence. Specifically, for this analysis we focus on the proportion of children able to read words or sentences, and able to write without difficulty/errors (treating each as separate outcomes). **Table 3** reports averages of stunting and the educational variables across countries and cohorts, whereas the distributions of height-for-age and weight-for-age for YC and OC are reported in **Figure A.2** and **Figure A.3** (respectively) in the Online Annex.

Table 3: selected outcomes of the country samples

	Stunting (in %)		**Writes without difficulty/errors (in %)**	**Reads words or sentences (in %)**	**School enrolment (in %)**
	Younger Cohort	*Older Cohort*	*Older Cohort*		
Ethiopia	43	33	23	27	66
India	31	33	53	65	98
Peru	29	27	55	85	99
Vietnam	21	29	75	92	99

Source: Young Lives data, authors' elaboration.

7. Results

Comparison between attrited and non-attrited YLIs using t-tests

We start by describing differences between attrited and non-attrited YLIs as informed by t-tests (results are reported in **Table A.3** and **Table A.4** for YC and OC, respectively, in the Online Annex). For the younger cohort, in Ethiopia, India, and Vietnam, attrited YLIs are more likely to be from urban areas and from the top-tercile of wealth, as had been suggested by the patterns reported in **Figure 3** and **Figure 4**. For the Peruvian sample, differences by household wealth and area of residence are not statistically significant. In addition, attrited YLIs are more likely to have caregivers with a higher level of education in India, and more likely to have caregivers that are single in Ethiopia, Peru and Vietnam.

Differences by ethnicity/caste also appear relevant. In the Peruvian sample, the proportion of caregivers whose native tongue is Spanish (which are more likely to come from wealthier households) is lower among the attrited group. In the Vietnamese sample, children from the largest ethnic group (Kinh) are less likely to be attrited. In the Indian sample, attrited children are less likely to be from BC (Backward Classes) and more likely to be from OC (Other Class).

For the older cohort, females are more likely to be attrited in the Ethiopian sample, whereas males are more likely to be attrited in the Peruvian sample. In addition, YLI that are the first born are less likely to be attrited in the Indian sample. Looking at caregiver and household characteristics, attrited YLIs are more likely to come from better-off families in Ethiopia, India and Vietnam. For instance, attrited children are more likely to be from urban areas in India and Vietnam, less likely to be from the bottom-tercile in Ethiopia, more likely to be from the top-tercile in India, and their caregivers are more likely to have a higher level of education in India. Conversely, for Peru the opposite seems to occur: attrited children are more likely to be from rural areas, to come from larger households, and are less likely to have Spanish as maternal language. In addition, for this cohort some differences according to YLI characteristics are found. Confirming patterns observed in the graphical analysis,

We also explore differences in YLI outcomes by attrition status. For the younger cohort, in the Ethiopian sample attrited participants have lower height-for-age, lower weight-for-age, and higher levels of stunting. In contrast, for the older cohort where mortality is much lower, in the Ethiopian sample attrited participants have a higher height-for-age, are more likely to read words and sentences, and more likely to read without difficulties compared to non-attrited participants. In addition, stunting levels are higher in the Peruvian sample among attrited participants (younger cohort), whereas in the Indian and Vietnamese samples attrited participants are less likely to be school enrolled (older cohort).

Testing for attrition bias

Previous results show that attrition is correlated with observed variables, Therefore, it is important to establish whether there could be attrition bias driven by observable characteristics. We test whether YLI outcomes observed in the first visit predict attrition in the last visit, conditional on individual and household characteristics (attrition probit tests).

Results are reported in **Table 4** (Panel A). All estimations consider the clustered nature of the samples.

Table 4: Attrition probit tests (marginal effects)

	Younger Cohort				Older Cohort			
	Ethiopia Attrition	India Attrition	Peru Attrition	Vietnam Attrition	Ethiopia Attrition	India Attrition	Peru Attrition	Vietnam Attrition
PANEL A: using preferred definition of attrition								
Height-for-age z-score	-0.013***	-0.004	-0.004	-0.004	0.020**	-0.008	0.018	0.001
	(0.004)	(0.003)	(0.007)	(0.003)	(0.010)	(0.006)	(0.013)	(0.010)
Weight-for-age z-score	-0.016***	-0.006	-0.003	-0.003	0.017	-0.007	0.016	0.007
	(0.004)	(0.004)	(0.008)	(0.003)	(0.011)	(0.007)	(0.016)	(0.008)
Stunting	0.040***	0.017	0.028**	0.013	-0.033	-0.004	-0.009	0.002
	(0.013)	(0.011)	(0.014)	(0.008)	(0.029)	(0.014)	(0.027)	(0.022)
Reads words or sentences					0.016	-0.026*	-0.002	-0.009
					(0.056)	(0.014)	(0.030)	(0.037)
Writes without difficulty					0.043	-0.024	0.001	0.001
					(0.038)	(0.024)	(0.027)	(0.017)
Child is school enrolled					0.013	-0.147*	-0.050	-0.201**
					(0.034)	(0.084)	(0.102)	(0.078)
PANEL B: attrition excluding dead children								
Height-for-age z-score	-0.003	-0.000	0.002	-0.001	0.020**	-0.001	0.017	0.005
	(0.003)	(0.002)	(0.006)	(0.003)	(0.009)	(0.007)	(0.013)	(0.010)
Weight-for-age z-score	0.001	-0.000	0.003	0.001	0.012	-0.001	0.016	0.008
	(0.003)	(0.002)	(0.008)	(0.003)	(0.009)	(0.006)	(0.017)	(0.007)
Stunting	0.010	0.000	0.011	0.003	-0.034	-0.014	-0.008	-0.009
	(0.010)	(0.006)	(0.012)	(0.005)	(0.026)	(0.015)	(0.027)	(0.021)
Reads words or sentences					0.020	-0.016	0.015	-0.001
					(0.052)	(0.016)	(0.025)	(0.032)
Writes without difficulty					0.044	-0.018	0.003	-0.001
					(0.035)	(0.022)	(0.025)	(0.017)
Child is school enrolled					0.040	-0.071	-0.065	-0.152***
					(0.030)	(0.079)	(0.104)	(0.059)

Notes: Each cell reports marginal effects from different probit model specifications. In all cases, the dependent variable takes the value of 1 if the YLI is attrited in the last visit, 0 otherwise. YLI outcomes are observed at age 1 and 8 for the younger cohort and older cohort, respectively. Estimates control for the following characteristics: sex of YLI, YLI is the first born (OC only); caregiver's sex, age, and marital status; household wealth, area of residence; caregiver´s education, ethnicity or region (caste in India):, household size.. Robust standard errors (clustered at the cluster level) are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Out of 12 outcomes observed for the younger cohort, YLI lagged outcomes predict attrition in 4 cases (3 at the 1% level, 1 at the 5% level). In the Ethiopian sample the three available outcomes (height-for-age, weight-for-age, and stunting), predict attrition. Similarly, in the Peruvian sample, stunting predicts attrition. In both cases, YLIs with worse nutritional status are more likely to be attrited. For the older cohort, YLI lagged outcomes predict attrition in 4 out of 24 outcomes (2 at the 5% level, 2 at the 10% level). There is evidence of attrition bias for height-for-age in the Ethiopian sample, for the reading task in the Indian sample, and for school enrolment in the Indian and Vietnamese samples. While this is not conclusive evidence of systematic *non-ignorable* selection on observables in the YL samples, results show that this might occur in some cases. Our definition of attrition includes dead participants. When considering an alternative definition in which dead children are excluded

(Panel B), results become statistically insignificant for all outcomes in the younger cohort and for the Indian sample in the older cohort. This suggests that to a large extent child mortality explains these biases, when they occur.

Finally, we report results of BGLW tests for the younger cohort and older cohort. In total, we estimate 12 models for the younger cohort (3 outcomes for each country sample) and 24 for the older cohort (6 outcomes for each country sample). For each estimation, the indicator of interest is the F-test (derived from the Wald test) for the joint effect of attrition on all coefficients, including and excluding the attrition coefficient. See **Table 5** and **Table 6**, respectively (full estimates are reported in the Online Annex, **tables A.5.1-A.5.3** and **tables A.6.1-A.6.6**, respectively). One key consideration is that F-tests must also consider the clustered nature of the YL samples (Cameron and Miller, 2015). At the same time, the number of clusters in each country sample is small. The combination of both elements limits the number of constraints one can test simultaneously. Specifically, since there are 20 clusters, we can only test the joint significance of up to 19 coefficients. Although we have less than 19 constraints in each country sample (11 in Peru and Vietnam, 13 in India, 17 in Ethiopia), this limits the degrees of freedom available. To deal with this aspect, we report results with an alternative clustering at the community level, where the community is the smallest geographical unit observed in each sample. This makes a substantial difference in India and Peru, where there are 100 and 81 communities, but not in Ethiopia and Vietnam with 26 and 31 communities each. For this reason, for the BLGW tests we focus our attention in the Indian and Peruvian samples.

For the younger cohort, results show suggestive evidence of attrition bias for stunting in India and Peru. In both cases, the null hypothesis that all slope coefficients (excluding the constant) are equal for attritors and non-attritors is rejected at the 1% level. Similar evidence is observed for height-for-age in India (at the 1% level) and weight-for-age in Peru (at the 5% level). Results from the full model show that coefficients that are different between attrited and non-attrited participants are those associated to household wealth, area of residence, ethnicity in Peru, caste in India. The attrition dummy is not statistically significant in any of these models, which suggests that bias is already captured by the control variables included in the model.

For the older cohort, among the nutritional outcomes, the null hypothesis that all slope coefficients are equal for attritors and non-attritors is rejected for weight-for-age in India and Peru (at the 1% and 5% level, respectively), and for height-for-age in Peru (at the 10% level). Among the educational outcomes, similar evidence is found for the reading task in both country samples, and for the writing task in the Indian sample (at the 1% level). While for the younger cohort household wealth, area of residence, and ethnicity appear as crucial dimensions, for the older cohort coefficients that are different between attrited and non-attrited participants (especially for educational outcomes) are those related to caregiver´s education, caregiver's age and participant's order of birth arise. As in the previous case, the attrition dummy is not statistically significant in these models, suggesting bias is already captured by the control variables included in the model.

Table 5: BGLW tests, Younger Cohort

	Ethiopia		**India**		**Peru**		**Vietnam**	
Height for age								
N	*1,942*	*1,942*	*1,992*	*1,992*	*2,036*	*2,036*	*1,992*	*1,992*
F (excluding constant)	1.18	3.86	1.12	2.13	1.36	0.99	1.03	1.95
Pvalue	0.2993	0.0030	0.3396	0.0284	0.2106	0.4485	0.4120	0.0891
F (including constant)	2.55	7.53	1.15	1.96	1.30	0.98	1.16	1.90
Pvalue	0.0034	0.0000	0.3213	0.0409	0.2334	0.4594	0.3163	0.0906
Weight for age								
N	*1,852*	*1,852*	*1,994*	*1,994*	*2,037*	*2,037*	*1,992*	*1,992*
F (excluding constant)	1.61	4.08	1.05	1.18	2.65	2.19	1.37	4.42
Pvalue	0.0980	0.0021	0.3987	0.3161	0.0069	0.0362	0.2053	0.0013
F (including constant)	3.05	6.07	1.22	1.24	2.39	1.97	1.30	3.93
Pvalue	0.0005	0.0001	0.2686	0.2710	0.0108	0.0534	0.2294	0.0022
Stunting								
N	*1,942*	*1,942*	*1,992*	*1,992*	*2,036*	*2,036*	*1,992*	*1,992*
F (excluding constant)	1.05	4.06	1.62	3.42	1.48	2.69	2.15	4.80
Pvalue	0.3962	0.0021	0.0957	0.0007	0.1607	0.0111	0.0288	0.0007
F (including constant)	1.89	7.18	1.69	3.85	1.72	2.54	2.11	6.30
Pvalue	0.0363	0.0000	0.0697	0.0001	0.0787	0.0127	0.0261	0.0001
Clustered standard errors	No	Yes	No	Yes	No	Yes	No	Yes
Number of clusters	-	26	-	100	-	83	-	31

Notes: Each country-outcome cell reports results from a model specification in which the country-outcomes is regressed on individual and household characteristics observed in the first visit, the attrition indicator, and the interaction between the attrition indicator and all individual and household characteristics. The F-statistic (excluding constant) corresponds to the null hypothesis that all slope coefficients are equal for attritors and non-attritors, whereas the F-statistic (including constant) corresponds to the null hypothesis that all coefficients are equal for attritors and non-attritors.

Table 6: BGLW tests, Older Cohort

	Ethiopia		India		Peru		Vietnam	
Height for age								
N	*963*	*963*	*1,008*	*1,008*	*709*	*709*	*1,000*	*1,000*
F (excluding constant)	1.26	11.90	0.73	1.13	0.82	1.86	1.00	3.72
Pvalue	0.2177	0.0000	0.7321	0.3458	0.6218	0.0639	0.4425	0.0021
F (including constant)	1.42	11.59	0.76	1.29	0.86	2.42	0.92	4.39
Pvalue	0.1259	0.0000	0.7155	0.2298	0.5913	0.0129	0.5215	0.0005
Weight for age								
N	*928*	*928*	*1,008*	*1,008*	*710*	*710*	*1,000*	*1,000*
F (excluding constant)	1.87	28.82	0.69	2.49	1.24	2.29	0.64	3.08
Pvalue	0.0224	0.0000	0.7739	0.0056	0.2539	0.0210	0.7934	0.0070
F (including constant)	1.87	46.46	0.70	2.38	1.24	2.35	0.65	2.94
Pvalue	0.0197	0.0000	0.7734	0.0067	0.2541	0.0154	0.8033	0.0082
Stunting								
N	*963*	*963*	*1,008*	*1,008*	*709*	*709*	*1,000*	*1,000*
F (excluding constant)	1.35	41.24	0.57	1.01	0.69	1.08	1.50	6.38
Pvalue	0.1678	0.0000	0.8783	0.4476	0.7520	0.3915	0.1239	0.0000
F (including constant)	1.36	48.02	0.54	1.16	0.64	1.39	1.38	7.10
Pvalue	0.1563	0.0000	0.9123	0.3194	0.8094	0.1986	0.1687	0.0000
Reads words/sentences								
N	*988*	*988*	*997*	*997*	*685*	*685*	*994*	*994*
F (excluding constant)	2.66	7.86	0.89	3.19	1.63	3.03	0.62	1.41
Pvalue	0.0006	0.0000	0.5634	0.0005	0.0851	0.0029	0.8115	0.2179
F (including constant)	2.51	14.25	0.95	3.39	1.50	2.92	0.57	2.65
Pvalue	0.0009	0.0000	0.5025	0.0002	0.1190	0.0032	0.8667	0.0152
Writes without difficulty								
N	*988*	*988*	*997*	*997*	*685*	*685*	*994*	*994*
F (excluding constant)	2.59	17.98	1.19	4.84	0.65	0.96	0.82	1.66
Pvalue	0.0008	0.0000	0.2798	0.0000	0.7840	0.4880	0.6235	0.1331
F (including constant)	2.57	16.87	1.24	5.31	0.60	0.90	0.75	1.52
Pvalue	0.0007	0.0000	0.2365	0.0000	0.8448	0.5484	0.6987	0.1717
Enrolled at school								
N	*999*	*999*	*1,008*	*1,008*	*713*	*713*	*1,000*	*1,000*
F (excluding constant)	2.06	48.94	2.73	0.83	0.72	0.39	3.85	10.26
Pvalue	0.0100	0.0000	0.0008	0.6328	0.7197	0.9529	0.0000	0.0000
F (including constant)	1.95	47.20	2.90	0.91	0.68	0.38	4.01	9.41
Pvalue	0.0140	0.0000	0.0003	0.5467	0.7676	0.9640	0.0000	0.0000
Clustered standard errors	No	Yes	No	Yes	No	Yes	No	Yes
Number of clusters	-	26	-	100	-	83	-	31

Notes: Each country-outcome cell reports results from a model specification in which the country-outcomes is regressed on individual and household characteristics observed in the first visit, the attrition indicator, and the interaction between the attrition indicator and all individual and household characteristics. The F-statistic (excluding constant) corresponds to the null hypothesis that all slope coefficients are equal for attritors and non-attritors, whereas the F-statistic (including constant) corresponds to the null hypothesis that all coefficients are equal for attritors and non-attritors.

8. Conclusions

This study looks at attrition patterns in the YL study. We document that the attrition rates observed for the YL samples are among the lowest observed attrition rates in longitudinal studies in low-and-middle-income countries. The older cohort has a larger attrition rate than the younger cohort, which is possibly related to increased migration rates and increased refusal rates among individuals that are transiting to adulthood. While attrition rates are low by international experiences, our results show that attrition in the YL samples is correlated with aspects such as wealth, area of residence, ethnicity and caste (among other socio-economic characteristics of the household and caregiver). Given that 15 years of data collection have passed, this result is not surprising. Evidence suggests that better-off families (e.g., households from urban areas and/or from the top-tercile of wealth) are more likely to be attrited in Ethiopia, India and Vietnam. On the contrary, for the Peruvian sample, evidence shows that rural households are more likely to be attrited.

Using attrition probit tests, we find evidence of attrition bias in certain outcomes in some country samples (4 out of 12 in the younger cohort and 4 out of 24 in the older cohort). To a large extent this result is driven by child mortality. This is true especially for the younger cohort in the Ethiopian sample. If children that died are excluded, there is no evidence of bias for the younger cohort, and very little evidence for the older cohort (2 out of 24 outcomes). We also performed BGLW tests, however this type of test can only be implemented in the Indian and Peruvian samples because only in those cases the number of communities is large enough. Using this strategy, we detect evidence of attrition bias, which exact nature depends on the country sample, the cohort, and the outcome.

A research question that has not been considered in this paper and should be addressed in future work is the reasons for YL having lower attrition rates than those appearing in other comparable, longitudinal, studies. Are these lower attrition rates related to characteristics of the chosen countries, characteristics of the children or families or can be explained by the protocols used to track the individuals? Penny et al. (2012) suggests that, at least for the case of Peru, fieldworker recruitment and training processes may partially explain low attrition rates. In addition, the fact that all countries implemented tracking exercises between rounds could also be part of the explanation. Having collected fieldworker's characteristics across rounds should allow us to test whether differences in fieldworker characteristics (age, gender, ethnicity, education, experience) has influenced attrition rates across or within countries.

Recommendations that arise from our analysis can be divided in two strands. One related to YL study and the other related, in general, to other cohort studies. One specific recommendation is that when performing statistical analysis of the YL country samples it is important to control for household and caregiver socio-economics characteristics observed in the first visit. Other strategies often applied to longitudinal data, such as individual and household fixed effects estimates can also be thought as alternative ways to deal with potential attrition bias. The evidence found suggest that, when there is bias, it is driven mostly by wealth and area of residence. Given the fact that studies using the YL sample control for these characteristics in their statistical analysis, this strategy should suffice to limit the effects of attrition bias

Our results show that there are some patterns that are common to all samples and there are some patterns that are country specific. For other birth cohort studies, it may be useful to distinguish between those correlates that consistently affect attrition in all countries (and have also been shown as relevant in the literature) from those correlates which relevance vary across cohorts or across countries. In the first case, the recommendation involves the use of sampling strategies to reduce attrition in subpopulations that are likely to be exposed to attrition. Similarly, the evidence suggests that the effort for tracking a cohort of children needs to be increased as the children enter into adulthood, as they have a larger say in whether to remain participating in the study and their chances for increased mobility rise.

The evidence also suggests that different contexts may generate different attrition patterns and if we cannot reduce attrition ex-ante by increasing sample size, we still need to confront the potential biases its presence may generate. For example, attrition rates between YLIs from poorer and wealthier households may be different across countries. In this case there is no single recommendation to tackle high attrition rates. Rather there is a need to review its presence and control for wealth status accordingly. Similarly, females may be more likely to be attrited in one country (YL Ethiopia), whereas males may be more likely to be attrited in other countries (YL Peru). Controlling for sex should suffice to limit the effects of this attrition bias.

Even if attrition is low, we could expect attrition bias when a sample is followed for a long time. As children grow and as their families are exposed to changing conditions, reasons for splitting, moving arise and these reasons are usually more likely to affect to specific subgroups of the tracked sample. It is important to control for these observables when performing statistical analysis as a way to deal with potential attrition bias.

References

Alderman, H., Behrman, J.R., Kohler, H-P., Maluccio, J.A., Cotts Watkins, S. (2001) "Attrition in longitudinal household survey data", Demographic Research, 5(4), p. 79-124.

Barnett, I., Ariana, P., Petrou, S., Penny, M., Duc le, T., Galab, S., Woldehanna, T., Escobal, J., Plugge, E., Boyden, J. (2013) "Cohort profile: the Young Lives study", Int J Epidemiol, 42(3), p. 701-8.

Becketti, S., W. Gould, L. Lillard and F. Welch (1988) "The panel study of income dynamics after fourteen years: an evaluation", Journal of Labour Economics 6, p. 472-92.

Boyden, J. (2018a). Young Lives: an International Study of Childhood Poverty: Round 2, 2006. [data collection]. 3rd Edition. UK Data Service. SN: 6852, http://doi.org/10.5255/UKDA-SN-6852-3

Boyden, J. (2018b). Young Lives: an International Study of Childhood Poverty: Round 3, 2009. [data collection]. 3rd Edition. UK Data Service. SN: 6853, http://doi.org/10.5255/UKDA-SN-6853-3

Cameron, C., Miller, D (2015). "A practitioner's guide to cluster-robust inference". Mimeo.

Dercon, S., Outes-Leon, I. (2008) "Survey attrition and attrition bias in Young Lives". Technical Note Number 5. Oxford: Young Lives.

Escobal, J., Flores, E. (2008). "An assessment of the Young Lives sampling approach in Peru". Technical Note 3. Oxford: Young Lives.

Falaris, E. (2003) "The effect of survey attrition in longitudinal surveys: evidence from Peru, Côte d'Ivoire and Vietnam", Journal of Development Economics 70, p. 133–157.

Fitzgerald, J., Gottshalk P., Moffitt, R. (1998) "An analysis of sample attrition in panel data: the Michigan panel study on income dynamics", Journal of Human Resources 33: 251-99.

Heckert, J. (2015). "New perspective on youth migration: motives and family investment patterns", Demographic Research, 33, 765-800.

Heckman, J, (1979) "Sample selection bias as a specification error", Econometrica (47), p. 153-162.

Jones, N., Huttly, S. (2018). Young Lives: an International Study of Childhood Poverty: Round 1, 2002. [data collection]. 6th Edition. UK Data Service. SN: 5307, http://doi.org/10.5255/UKDA-SN-5307-3

Kumra, N. (2008). “An assessment of the Young Lives sampling approach in Andhra Pradesh, India”. Technical Note 2. Oxford: Young Lives.

Nguyen, N. (2008). “An assessment of the Young Lives sampling approach in Vietnam”. Technical Note 4. Oxford: Young Lives.

Outes-Leon, I., Sánchez, A. (2008). “An assessment of the Young Lives sampling approach in Ethiopia”. Technical Note 1. Oxford: Young Lives.

Penny, M., S. Madrid, B. Oré (2012). “Selection and Induction of Supervisors for Fieldwork: Experiences from Young Lives in Peru” Technical Note 24. Oxford: Young Lives.

Pomeroy, E., Stock, J., Stanojevic, S., Miranda J., Cole, T., Wells, J. (2014) “Stunting, adiposity, and the individual-level “dual burden” among urban lowland and rural highland peruvian children”. Am J Hum Biol, 26(4), p. 481–90.

Rayan, D., P. Petesch. (2007). “Moving Out of Poverty (Volume 1) Cross-disciplinary Perspectives on Mobility”, Palgrave Macmillan and The World Bank, Washington D.C.

Sánchez, A., Penny, M., Woldehanna, T., Galab, S., Boyden, J., Duc, L. Thuc. (2018). Young Lives: an International Study of Childhood Poverty: Round 5, 2016. [data collection]. UK Data Service. SN: 8357,

Schafer, J., Graham, J. (2002) “Missing data: our view of the state of the art” Psychological Methods, 7(2), p. 147-177.

Valentine, J. L., Barham, B., Gitter, S., Nobles, J. (2017). “Migration and the pursuit of education in southern Mexico”, Comparative Education Review, 61(1), p. 141-175.

Verbeek M., Nijman T., (1996) “Incomplete panels and selection bias”. In The Econometrics of Panel Data, eds. L. Maty as and P. Sevestre, Dordrecht: Kluwer.

Wilson, I., Huttly, S., Fenn, B. (2006), A Case Study of Sample Design for Longitudinal Research: Young Lives. Int. J. Social Research Methodology, Vol. 9, No. 5, December 2006, pp. 351–365

Woldehanna, T., Galab, S., Sanchez, A., Penny, M., Duc, L. Thuc, Boyden, J. (2018). Young Lives: an International Study of Childhood Poverty: Round 4, 2013-2014. [data collection]. 2nd Edition. UK Data Service. SN: 7931, http://doi.org/10.5255/UKDA-SN-7931-2

Ziliak, J., Kniesner, T. (1998). “The importance of sample attrition in life cycle labour supply estimarion”. Journal of Human Resources, 33, p. 507-30.