Poverty at Higher Frequency

Joshua D. Merfeld* Jonathan Morduch†

September 20, 2026

Abstract

Low incomes are often unsteady within the year. By design, however, poverty headcounts in national statistics ignore the instability. We introduce the *timecount*, an annual measure that explicitly captures the shifting durations and intensities of poverty within the year. We show that common survey methods unintentionally re-shape the math of poverty measures with the consequence that timecounts have become the *de facto* poverty rates in India and other countries, having effectively replaced conventional headcounts. Recognizing the *de facto* role of the timecount in national statistics—and the absence of headcounts—broadens policy strategies. In monthly longitudinal data from rural India, the timecount is 28% larger than the conventional headcount. It registers the months in poverty that are experienced by roughly half of "non-poor" households, and it better predicts child health.

Keywords: timecount, income volatility, consumption smoothing, poverty measurement, seasonal poverty, weighted timecount, household expenditure
JEL Codes: I32, G51, D14, D15

*School of Economics, University of Queensland, St Lucia, QLD, 4072, Australia and IZA; j.merfeld@uq.edu.au
†Robert F. Wagner Graduate School of Public Service, New York University, New York 10003, USA; jonathan.morduch@nyu.edu

1 Introduction

National poverty rates are meant to capture how many people subsist on annual incomes that are insufficient to meet their basic needs (Ravallion, 2016; Atkinson, 2019). We show, however, that national poverty rates in India and other low- and middle-income countries are not headcounts as commonly understood. Without intention or recognition, the measures have been systematically transformed by the nature of the underlying survey data. As a result, today the national statistics instead capture a fundamentally different notion of poverty that encompasses a broader range of people and that responds to a wider set of policies.

The *de facto* national poverty rates approximate a poverty measure that we define as the *timecount*. The statistic still captures a simple idea of deprivation, but, unlike the headcount, the timecount captures the shifting durations and intensities of poverty within the year. It registers people's most difficult moments rather than their average conditions. The timecount has a simple interpretation as the average fraction of the year in which people are poor, measured with household consumption. Alternatively, the timecount can be seen as the share of all periods in a year that are experienced as being in poverty, counted across all periods and people. If a person's resources are below the poverty line for all four quarters of the year, for example, they contribute 1 to the timecount. If their resources are below the threshold for only a single quarter of the year, they contribute 0.25. A related measure, the *weighted timecount*, puts extra weight on lean seasons and other months with below-average resources. By being sensitive to the duration and timing of deprivation within the year, the timecount highlights income instability and the limits to consumption smoothing within the year.

Our focus is on India and low- and middle-income countries that base poverty measurement on expenditure surveys. Together, the countries are home to 99 percent of people counted as being in extreme poverty (World Bank, 2024). Almost all low-income countries measure poverty with expenditure data rather than income (Mancini and Vecchi, 2022). Ninety percent of lower-middle income countries and 62 percent of upper-middle income countries do as well.[1]

[1]Mancini and Vecchi (2022), Figure 3.2. Data are from a survey of data collection methods in 137 countries. Notable exceptions among middle-income countries outside of Latin America include the Philippines and Malaysia. Just 9 percent of high-income countries use consumption rather than

With within-year longitudinal data from rural India, we show that the timecount is 28% larger than the headcount. The increase is due to capturing all months of poverty experienced by the population, including the 35% of all months in poverty experienced by households that would not conventionally be considered poor. With a lasso regression, we show that the variation in the timecount is economically meaningful: it is a better predictor of child height-for-age than the variation in the headcount.

The analysis does not suggest that countries should necessarily consider replacing headcounts with timecounts. Instead, the analysis shows how, without intention, timecounts have already replaced headcounts in national statistics. We also show the impossibility of directly calculating conventional headcounts with survey data from India and other countries that use related survey methods. Since the World Bank's global measure of extreme poverty (Lønborg et al., 2026) draws on the same national surveys, it too cannot be interpreted as a headcount.

The transformation of headcounts into timecounts has been hidden because poverty measures are typically understood through the mathematical forms of indices (e.g., Sen 1976, Watts 1968, and Foster et al. 1984). We show, however, that the properties and meanings of measures are shaped by the way that statistical agencies collect the underlying data. Statisticians face a double challenge when collecting data on the expenditure and consumption of households. The first challenge is that household consumption typically varies through the year, often marked by seasonal patterns, especially for the poorest residents (Chambers, 1983; Devereux et al., 2012). The second challenge is that memory is fallible (Augenblick et al., 2024; Beegle et al., 2012). As a consequence, expenditure surveys usually focus on households' consumption during the week or two before the interview—and, at most, the past month (Smith et al., 2014; Mancini and Vecchi, 2022). To capture seasonality and to spread logistical burdens, experts advise statistical agencies to interview households across the year, sampling new cross-sections in each period of the year, often stratifying by quarter (FAO and World Bank, 2018).

The surveys thus contain snapshots of consumption by different households at

income. Lora et al. (2024) provide a parallel analysis with data on household income in Colombia. Unlike the analysis with expenditure data, an income-based timecount would require assuming that households do not smooth consumption at all. As a result, the focus of Lora et al. (2024) is on annualizing income data.

different times of year. As a result, what is labeled as "household consumption" in surveys reflects short-term conditions specific to the dates on which particular households are interviewed. We show that when statisticians use standard mathematical formulas for headcounts with the data labeled as "household consumption," the resulting indices are transformed by the timing of data collection and the nature of statistical sampling.

We first illustrate the dependence of measures of household consumption on interview dates with the 68th round (2011-12) of the Indian National Sample Survey, a large, nationally-representative data set used to calculate official poverty statistics. We show that quarterly averages of household expenditure follow seasonal patterns, especially for households with lower socio-economic status. For less educated households, for example, average consumption varies by 8% across quarters, while it varies by just 3% for more educated households—consistent with the former group's greater exposure to income instability and weaker tools to smooth consumption within the year.

Within-year longitudinal data is needed to show the implications for poverty analyses. Since that is unavailable in the Indian National Sample Survey, we turn to the Village Dynamics in South Asia (VDSA) Survey, a balanced panel of household-level data from rural India from 2010-2014 that includes the incomes and expenditures of 945 low-income households collected monthly for at least four years. With the VDSA survey, we simulate the sampling process used to collect the Indian National Sample Survey—consistent as well with the expert guidance for data collection in FAO and World Bank (2018). Specifically, we use the standard formula for the headcount but replace the yearly average of household consumption with data for each household from a single randomly-selected month.

The true headcount in this sample is 29%, but we show that the statisticians' approach to estimating the headcount—calculated 10,000 times with random draws of data—converges to 36.98%, which instead closely approximates the true timecount calculated with data from all months (37%). Section 2 provides the mathematical reasoning.

The VDSA data include anthropometric measures, and we show that, by capturing within-year variation, timecounts add predictive power beyond conventional headcounts. Focusing on child weight and height, we use the lasso, the least absolute shrinkage and selection operator, to show that a household-level indicator that

reflects within-year variation in poverty (as in the timecount) is selected among a set of core predictors of child height, and a binary poverty indicator based on average conditions (as in the headcount) drops from the set.

The timecounts in the VDSA are larger than headcounts mainly due to months of poverty experienced by the 71% of households labeled as non-poor in annual data. Of these "non-poor" households, 47% experience at least one month of poverty. When measuring deprivation with the distribution-sensitive weighted timecount, poverty statistics increase by 40% and 48% relative to conventional analogues. The patterns align with related findings in national-level data analyzed by Gibson et al. (2003) in urban China, Jolliffe and Serajuddin (2018) in Jordan, and Gibson and Alimi (2020) in Nigeria.

The framework builds on independent strands of research that establish building blocks of the argument. Prior studies describe the shifting conditions of poverty during the year, partly through the choices of households and partly due to income instability and illiquidity (Chambers, 1983; Collins et al., 2009; Mullainathan and Shafir, 2013). As noted above, in rural areas, for example, deprivations typically intensify in lean seasons and ease in harvest seasons (e.g., Longhurst et al. 1986, Khandker 2012, Devereux et al. 2012, Dercon 2002, Dostie et al. 2002, Carter and Lybbert 2012, Gilbert et al. 2016, Breza et al. 2021, Bryan et al. 2014, Pomeranz and Kast 2024, Fink et al. 2020, Casaburi and Willis 2018, and Augenblick et al. 2024). Income and consumption in urban areas are often unstable as well, shaped by fluctuations in the availability of work and seasonal shifts in demand (e.g., Collins et al. 2009, Gibson et al. 2003, Jolliffe and Serajuddin 2018, Azevedo and Seitz 2016b). Christian and Dillon (2018) and Barker et al. (2002) show that deprivations specific to lean seasons can have long-term negative consequences for children as they age.[2]

Similarly, we draw on analyses of poverty dynamics which document instability across years (e.g., Bane and Ellwood 1986, Morduch 1994, Jalan and Ravallion 1998, Baulch and Hoddinott 2000, Addison et al. 2009, Christiaensen and Shorrocks 2012,

[2]In this line, Evidence Action (2019) writes that "seasonal hunger and deprivation are perhaps the biggest obstacles to the reduction of global poverty." Similarly, Vaitla et al. (2009) note that "Most of the world's acute hunger and undernutrition occurs not in conflicts and natural disasters but in the annual 'hunger season,' the time of year when the previous year's harvest stocks have dwindled, food prices are high, and jobs are scarce." Evidence Action (2019) estimates that seasonal hunger affects around 600 million of the world's rural poor.

Hoddinott 2006, Foster 2009, and Balboni et al. 2022). This work distinguishes between transient and chronic poverty in ways that are analogous to our within-year analyses. Jalan and Ravallion (1998) and Rodgers and Rodgers (1993), for example, analyze chronic poverty across years with an aggregate measure formally similar to the timecount.

The across-year studies, however, maintain conventional concepts of yearly poverty. Our focus on the dynamics of poverty *within* the year, while parallel to the analysis of year-to-year dynamics, changes how "poverty" itself is conceived and measured, and it sharpens the distinction between household income and household expenditure. It also blurs the division between chronic and transient poverty by recognizing the possibility of "chronic transience"—i.e., that short-term shifts in exposure to poverty can be a stable part of yearly economic life.

The analysis of data collection builds on Scott (1992), Deaton and Grosh (2000), Deaton and Zaidi (2002), Mancini and Vecchi (2022), and Beegle et al. (2012) who analyze the design and sampling methods of household consumption surveys. They document tradeoffs associated with the fallibility of respondents' memories and the need for short recall periods for questions on food and other frequently-purchased consumption items. These insights motivate the expert guidance in FAO and World Bank (2018).

Gibson (2001), Gibson et al. (2003), Gibson and Alimi (2020), and Jolliffe and Serajuddin (2018) were first to describe challenges in calculating headcounts when consumption varies during the year and longitudinal data are not available. They show inaccuracies that arise when headcounts are calculated using extrapolated data from randomly-sampled periods, and they show that poverty measures are then sensitive to within-year fluctuations. We build on their work by introducing timecounts, providing a precise interpretation for the *de facto* measures, and tying the emergence of timecounts to expert guidance for data collection.

We introduce the conceptual framework in section 2 and describe practical considerations in section 3. We then describe the longitudinal data from rural India in section 4 and the results in section 5. Section 6 draws connections to related, earlier evidence from China, Jordan, and Nigeria. Sections 5 and 7 and the online appendices raise empirical concerns including the roles of measurement error, consumer durables, secular patterns of growth, and time-varying preferences and needs. Section 7 concludes with conceptual and empirical reflections on poverty in

global settings.

2 Headcounts and timecounts

2.1 The headcount

To illustrate the transformation of headcounts into timecounts, we start with a year divided into T periods of equal length. In each period t, household i consumes c_{it}, so through the year they consume:

$$c_{i1}, c_{i2}, c_{i3}, ..., c_{iT}, \tag{1}$$

and total consumption for the year is $c_{i1} + c_{i2} + c_{i3} + ... + c_{iT}$. Dividing by T gives $\bar{c}_i$, household i's average consumption for the year.

Poverty rates are sometimes calculated based on conditions at one particular time of year, but we follow the guidance in Deaton and Zaidi (2002), Atkinson (2019), and Mancini and Vecchi (2022) who focus on consumption across the year, reflected here by $\bar{c}_i$.

The poverty headcount, P_H, for the population (i.e., the share of the population that is poor for the year based on their average consumption) is then:

$$P_H = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\bar{c}_i < z}. \tag{2}$$

where N is the total number of households and the indicator $\mathbb{1}_{\bar{c}_i < z}$ yields each household's binary poverty status, equal to 1 if $\bar{c}_i$ is below the poverty line z and 0 if not. (The poverty line reflects a household's minimum needs and is constant through the year.)[3]

In principle, the properties of poverty and inequality indices can be read directly from mathematical formulas like equation (2). Here, for example, the headcount is a nonlinear function of annual resources, with a sharp threshold at the poverty

[3]Our focus is on consumption-based poverty measures, but poverty rates are often measured with income instead. The analogue to the problems here will likely be greater with income given that it is not smoothed within the year and thus tends to show greater volatility.

line. The measure is additively separable, decomposable, and not distributionally sensitive (Ravallion, 2016). In practice, however, the properties of official poverty rates also depend on the data that is fed into the formulas. In the next sections we show that when starting with equation (2) and using survey data collected once a year from each household following expert guidelines (FAO and World Bank, 2018), the resulting measure is more accurately represented by a completely different formula, with a different meaning and scope. What seems to be a headcount as represented by equation (2) is in fact a timecount as described below.

2.2 The timecount

It is rare that households completely smooth consumption. Instead, the consumption series in equation (1) typically rises and falls with seasons and shifting economic conditions. To capture implications for deprivation in a simple way, the poverty status for each household can be determined by applying the poverty function $P(\cdot)$ to each period's consumption before aggregating:

$$P(c_{i1}), P(c_{i2}), P(c_{i3}), ..., P(c_{iT}), \tag{3}$$

where $P(c_{it}) = 0$ if $c_{it} \geq z$. The poverty function $P(\cdot)$ is additively separable, which includes forms that when aggregated across the population yield the headcount, income gap, and the distribution-sensitive measures of Watts (1968) and Foster et al. (1984).

Household i's average poverty status during the year, $\bar{P}^i$, is then:

$$\bar{P}^i = \frac{1}{T} \sum_{t=1}^{T} P(c_{it}). \tag{4}$$

In the most simple case, the poverty function reflects a binary poverty status like the headcount but applied to individual periods rather than yearly averages, so that in each period t household i is poor (1) or not (0) based on the indicator $P(c_{it}) = \mathbb{1}_{c_{it}<z}$. The average for household i is then $\bar{P}^i_T$, the share of the year that household i spends in poverty:

$$\bar{P}_T^i = \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}_{c_{it}<z}, \tag{5}$$

where the subscript T indicates that this is a measure of time in poverty. With a poverty line equivalent to the international extreme poverty line of \$3 per person per day, for example, if a household consumes on average \$1 per person per day for each of six months and an average of \$4 per person per day for each of the other six months, $\bar{P}_T^i$ would register 0.5 of a year in poverty.

The *timecount* poverty measure P_T is the average of $\bar{P}_T^i$ for all N households in the population:

$$P_T = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}_{c_{it}<z}. \tag{6}$$

The measure in equation (6) has a simple interpretation: it is the average fraction of the year that the population experiences poverty, bounded by 0 (no one in the population experiences any poverty at any time) and 1 (everyone is always poor). Between 0 and 1, the timecount reflects a mix of households that are always poor, never poor, or sometimes poor. Equivalently, the timecount is the share of periods in which people experience poverty, looking across all households and all periods. If the headcount can be seen as a "poverty of averages," the timecount can be seen as an "average of poverties."[4]

Ravallion (1988), Rodgers and Rodgers (1993), Jalan and Ravallion (1998), and Foster (2009) use a similar measure applied to longitudinal yearly data to analyze the variability and persistence of poverty across years, decomposing experiences into "chronic" and "transient" components, taking as given the conventional understanding of yearly poverty. Our approach also hinges on variability and persistence but, in contrast, harnesses them in aid of an alternative view of poverty itself.

[4]To simplify notation, we ignore population weights and weights for periods of different lengths. Adding weights would be straightforward; for example, poverty in January would contribute 31/365 to the weighted annual average, poverty in February would contribute 28/365 (except in a leap year), etc.

2.3 How headcounts become timecounts without longitudinal data

Statisticians use the formula in equation (2) to calculate headcounts, but statisticians use proxies for $\bar{c}_i$ when household i's average consumption for the year is not available in the data. When, as described in section 3 below, data are collected in accord with expert guidelines, the proxy is based on short-term household consumption from a randomly-selected period (section 3.1; FAO and World Bank 2018, Mancini and Vecchi 2022), which we denote c_i^*:

$$c_i{}^* = \sum_{t=1}^{T} c_{it} \cdot I_{it}, \tag{7}$$

where I_{it} is an indicator which captures the randomized sampling process; I_{it} is equal to 1 in the period that household i was randomly selected to be surveyed and 0 in the other $(T-1)$ periods of the year.[5]

Statisticians only see c_i^*, so they take the standard headcount formula in equation (2) and use c_i^* in place of $\bar{c}_i$:

$$P_H^* = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{c_i^* < z}. \tag{8}$$

When sampling is stratified by period, households are randomly selected to be interviewed in any given period with probability 1/T. The expected value of $\mathbb{1}_{c_i^* < z}$ for household i is thus the probability-weighted average value of $\mathbb{1}_{c_{it} < z}$ across the T periods. This turns out to be the share of the year that the household's consumption was below the poverty line; i.e., it is the household's contribution to the timecount (from equation 5):

$$E[\mathbb{1}_{c_i^* < z}] = \sum_{t=1}^{T} \frac{1}{T} \mathbb{1}_{c_{it} < z} = \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}_{c_{it} < z} = \bar{P}_T^i. \tag{9}$$

In expectation, the "headcount" in equation (8) is then the population average of

[5]The quantity c_i^* is, in practice, based on expenditures adjacent to the survey date, but c_i^* may be adjusted to include a share of some items that are larger and less-frequently purchased during the year. This is illustrated in the application using the Indian National Sample Survey in section 3.3. In some countries, the value for consumption (and the poverty line) may be scaled up to the quarterly or yearly level, but the underlying variation remains rooted in the period of the interview.

the quantity in equation (9):

$$E[P_H^*] = \frac{1}{N}\sum_{i=1}^{N} E[\mathbb{1}_{c_i^*<z}] = \frac{1}{N}\sum_{i=1}^{N} \bar{P}_T^i = P_T, \tag{10}$$

which is the timecount defined in equation (6).[6]

Importantly, while the exact timecount in equation (6) requires within-year longitudinal data for all households and for all periods, the approximation in equation (10) instead relies on randomized stratified sampling by period and the fact that the timecount is a linear aggregation of period-specific values. This finding is in the spirit of Scott (1992) who shows that data collection of this kind (yielding c_i^* for household i) does not deliver an accurate assessment of any given household's overall consumption, but, when averaged across the N households, the c_i^* approximate the true average consumption for the population. In the same way, the process here does not deliver an accurate assessment of any given household's poverty status based on their annual resources, but, when averaged across the N households, the $P(c_i^*)$ aggregate to approximate the true timecount for the population. Even without longitudinal data, the process yields a coherent, interpretable measure that captures variation in poverty within the year.

A corollary is that the headcount cannot be calculated with these data unless all households completely smooth consumption such that $c_i^* = \bar{c}_i$. When that is not the case, calculating the headcount requires within-year longitudinal data for all periods and all households because it depends on knowing $\bar{c}_i$ for *every* household i, and that requires repeated interviews with the same households when memories are fallible and conditions shift through the year.[7] Approximating the timecount, in contrast, depends only on knowing c_i^*, a single observation of consumption from each household from a randomly-chosen period.

[6] A similar logic applies to distribution-sensitive measures, where the indices of Watts (1968) and Foster et al. (1984) become, in practice, the weighted timecount measures introduced in equations (11) and (12) below.

[7] Another path is to impute values through a modeling exercise as described in the conclusion.

2.4 Time-sensitivity: How headcounts and timecounts differ

The timecount and the headcount provide complementary but distinct information, with the headcount as a special case of the timecount. The headcount is equivalent to the timecount when households steadily consume the same amount over time. Then, $c_{it} = \bar{c}_i$ for all t, and equation (5) collapses to $P_T^i = \mathbb{1}_{\bar{c}_i < z}$ and equation (6) thus simplies to the headcount in equation (2). This completely smoothed world, although hypothetical, is implicitly the subject of conventional poverty analyses (e.g., Ravallion 2016; Atkinson 2019). The gap between the headcount and timecount can then be seen as a measure of the extent by which experiences of poverty could be reduced if households in fact completely smoothed consumption, holding all else constant.

Because it is based on average consumption for the year, the headcount in equation (2) is "time-insensitive" in that measured poverty is unaffected by whether household i consumes the same amount every day of the year or if, instead, consumption has strong seasonal spikes and dips. Time-insensitivity means that the only way to reduce the poverty measure is to increase households' average consumption—e.g., policies to increase economic growth, invest in human capital, and improve jobs (e.g., Ravallion 2016).

In contrast, the timecount is "time-sensitive" in that the timing of consumption matters. Holding all else constant, a redistribution that lifts a household out of poverty in any period t will decrease the poverty measure, even if their average consumption is unchanged. This property does not hold for the headcount and other conventional poverty measures that depend only on yearly sums or conditions at a single time of year.

2.5 The weighted timecount

The distribution-sensitive measures of Sen (1976), Watts (1968), and Foster et al. (1984) build from the headcount to register varying degrees of material deprivation below the poverty line. Like the headcount, they are also based only on each household's average yearly resources, $\bar{c}_i$, and thus they are also time-insensitive as defined in section 2.4.

A generalized form of the timecount is both distribution-sensitive and time-

sensitive in that falling below one's annual average consumption level during the year registers as a greater incidence of poverty. The *weighted timecount* (wT), is defined similarly to equation (6). For example, a form that embeds the structure of Watts (1968) is:

$$P_{wT:Watts} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T}\ln\left(\frac{z}{c_{it}}\right)\mathbb{1}_{c_{it}<z}, \tag{11}$$

and a form that embeds the structure of the squared poverty gap of Foster et al. (1984) is:

$$P_{wT:FGT} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T}\left(\frac{z-c_{it}}{z}\right)^2\mathbb{1}_{c_{it}<z}. \tag{12}$$

These measures are sensitive to households' most challenging times of year. They count differently those times when households fare relatively better and worse, even if the households are always below the poverty line. In this way, the weighted timecount is time-sensitive such that, holding all else constant, a pure transfer of income from a household in a period t where they are below the poverty line to any period where they are richer must increase the poverty measure. This is a within-household version of the Pigou-Dalton transfer axiom adapted by Sen (1976), and it does not hold for standard applications of distribution-sensitive poverty measures to yearly consumption aggregates.

3 Collecting data on household consumption

3.1 Why "annual" data depends on when households are surveyed

Headcounts become timecounts through the interaction of survey methodologies and contexts where low-income households fail to smooth consumption during the year. We discuss surveys, then smoothing.

Smith et al. (2014) was presented at the forty-fifth session of the United Nations Statistical Commission (UNSC) in 2014, providing an assessment of the state of

global household-level consumption data. The preface to FAO and World Bank (2018) concludes that the assessment "painted a bleak picture in terms of heterogeneity in survey design and overall relevance and reliability of the data being collected." The result was a drive to harmonize practices led by the Inter-Agency and Expert Group on Food Security, Agricultural and Rural Statistics, convened by the World Bank and UN Food and Agriculture Organization, and endorsed by the forty-ninth session of the United Nations Statistical Commission in 2018 (FAO and World Bank, 2018). The guidelines are discussed further by Mancini and Vecchi (2022), which in turn updates Deaton and Zaidi (2002). The guidelines were also disseminated as a guidebook for the World Bank's Livings Standards Measurement Survey (LSMS) program.

We focus on the interaction of six methodological choices at the core of the guidelines: (1) basing welfare assessments on household consumption; (2) asking questions on food and other frequently-purchased items with reference to recent expenditures only; (3) surveying households just once during the year when budgets are limited; (4) collecting data on different households across the year to account for seasonality; (5) stratifying samples by period to avoid conflating seasonality with regional differences; and (6) treating the resulting data as being representative of each household's annual consumption.

The first methodology is well-established. The preference for consumption-based measures stems from the challenges of accurately measuring net income from common occupations like agriculture and self-employment in low- and middle-income countries. Because they tend to be more accurate, measures of consumption are generally higher than income for the same households, especially where earners do not receive wages or salaries (Carletto et al., 2021).

The second methodology (focusing on recent expenditures) results from the fallibility of memory (Scott, 1992; De Weerdt et al., 2020; Augenblick et al., 2024; Beegle et al., 2012). The limited ability to recall past food consumption leads Smith et al. (2014) to judge that a recall period longer than two weeks or less is unreliable. Smith et al. (2014) analyze a global survey conducted by the International Household Survey Network and covering data sets mainly collected between 2005-9 from low- and middle-income countries (Appendix A1.3, p. 50). Smith et al. (2014) show that of 96 household-level surveys, 41 percent employed a recall period for food of less than one week, most commonly a single day (using a diary method). Nearly

a quarter used a week exactly, 5 percent used two weeks or half a month, and 7 percent used a month. Taken together, Smith et al. (2014) find that 70 percent thus collect food data with "reliable" recall periods of two weeks or less and 77 percent have recall periods of a month or less. Faith in long-term memories has mostly been abandoned in favor of short-term snapshots.

The third methodology (surveying households just once) is a logistical compromise. The household surveys ask a battery of questions about different forms of expenditure and consumption. The surveys tend to be long and difficult. For a given budget, more households can be surveyed if the questions are asked of each household just once during a survey year. The result is more observations of households but fewer observations of the same households over time.[8] As Mancini and Vecchi (2022) summarize, echoing Deaton and Grosh (2000), "There is abundant evidence that food consumption and expenditure display systematic seasonal variation on a yearly, monthly, and weekly basis. The only way to accurately capture habitual consumption for each household is to survey them multiple times over the year, but this is also the most expensive option, and in practice, it is difficult to implement." (Appendix E, page 159.)

These first three methodological choices will deliver an accurate measure of household consumption if households fully smooth consumption within the year (Jolliffe and Serajuddin, 2018). With completely smoothed consumption, researchers can reliably extrapolate to the year from a single short-term snapshot of a given household's consumption taken from any period (Deaton and Grosh, 2000; Jolliffe and Serajuddin, 2018; Azevedo and Seitz, 2016b). Knowledge of what households consumed during a week in April, say, would deliver a good approximation of their annual consumption when multiplied by 52. But as World Bank economists Azevedo and Seitz (2016b) conclude, "...this assumption hinges on the existence of well-functioning formal and informal household smoothing mechanisms... Indeed, given market imperfections, the fact that markets often do not even exist in some developing countries, and the inability of poor households to smooth labor and non-labor income fluctuations, there are many reasons that we would expect poverty to fluctuate in the short-term." (See also, e.g., Longhurst et al. 1986, Breza et al. 2021, Bryan et al. 2014, Fink et al. 2020, Augenblick et al. 2024). In fact,

[8]Some statistical agencies address respondent fatigue by surveying households on multiple days within a short period, but they do not return to the household later in the year.

the analyses here show that the necessary condition in terms of data collection is stronger. What is needed is not just that households have the ability to completely smooth within the year if they choose (as in Paxson 1993) but that they actually do.

This is a conundrum for statistical agencies in resource-constrained environments, and, recognizing the within-year consumption instability, the expert guidelines add two more recommendations. The fourth methodology (collecting data through the full year) provides a way to account for seasonality and other within-year fluctuations at the regional and population level. The fifth methodology (stratification by period such that statistical agencies draw a new sample each period with replacement) avoids the conflation of time and place as survey teams interview different households at different times and in different places.

What is then aggregated and labeled as "household consumption" in data files is not yearly consumption (nor the yearly average for household i, $\bar{c}_i$). It is instead c_i^*, an observation of consumption rooted in a particular period, a random draw from the series of consumption levels shown in equation (1). As noted above, the six methodological choices together allow researchers to accurately track the population average of household consumption (Scott, 1992), but not nonlinear transformations of annual household consumption like poverty and inequality (Deaton and Grosh, 2000; Gibson, 2001).[9]

Deaton and Grosh (2000) and Deaton and Zaidi (2002) recognize the problem for poverty and inequality statistics and recommend instead asking households about their consumption in "a usual month" as a way to retrieve something closer to annual consumption for each household. About 10 percent of methods surveyed by Smith et al. (2014) did so. But the evidence in the past two decades has tilted strongly against the "usual month" approach, despite its conceptual appeal.[10] Beegle et al. (2012), for example, find that the "usual month" approach leads to significant

[9] Mancini and Vecchi (2022) conclude in their World Bank assessment of data collection practices, "Data collection spread over the year, but with only one interview per household, results in an accurate estimate of average consumption for the population, but with excess variability around the mean; however, it is a viable option in resource-constrained contexts. Regardless of the approach chosen, care should be exercised to ensure that enumeration is equally spread throughout the days of the week and the month..." (pp. 159-160).

[10] In principle, asking households about their consumption in a "usual month" is a way to maintain the year as a reference period and put aside concerns with seasonality, while also keeping the recall period to one month so households are not pressed to remember spending across long horizons (Mancini and Vecchi 2022, footnote 22, p. 28).

underestimation of household spending on food. Asking questions about a "usual month" is complicated (what is a usual month in an unstable context?) and in practice the questions take relatively long to ask and answer. FAO and World Bank (2018) conclude that "the usual month may be a lose-lose proposition if it is less accurate and more cumbersome to implement when compared to a seven-day recall. This is possibly the most important single development in the evidence base since the publication of Deaton and Grosh (2000)" (p. 19, cited by Mancini and Vecchi 2022).

The Beegle et al. (2012) evidence thus cements the practice of focusing on recent consumption of food and other frequently-purchased goods while addressing seasonality through sampling methods. As a result, statistical offices end up with c_i^* for each household i (from equation 7) rather than $\bar{c}_i$.

3.2 Sampling strategies in practice

Section 3.1 noted that when households completely smooth consumption, the timing of surveys does not matter: consumption sampled from a household in any period will roughly predict consumption across periods. If the assumption held, it would justify the data collection practices described below, where (in contrast to expert guidelines) statistical agencies most often survey households for just part of the year and/or do not stratify by period.

There is no readily-accesible information on the details of sampling methods within the meta-data on the household surveys used to calculate poverty and inequality, but Smith et al. (2014) provide a comprehensive accounting of survey practices in a wide range of low- and middle-income economies. They find that 88.5% of surveys involve interviewing households just once in the year. Only 11.5% of the surveys involve repeat interviews with the same households.[11] In the presence of instability, only this minority of surveys yield approximate annual household consumption.

[11] The list includes countries which, with the exception of Nigeria, contribute little to counts of global extreme poverty: Azerbaijan, Belarus, Fiji, Kazakhstan, Mauritania, Nigeria, Papua New Guinea, Philippines, Russia, Ukraine, and Vietnam. The assessment in Smith et al. (2014) includes 100 nationally representative surveys with sufficiently-documented methods. In addition, interviewing households multiple times per year does not necessarily mean that these households are asked about expenditure/consumption patterns multiple times a year; sometimes only certain modules are asked in a given visit.

A large group of surveys (42%) involves interviewing households once and conducting surveys throughout the year (usually 1/12 of the total sample surveyed in each month), in keeping with the expert guidelines of FAO and World Bank (2018). This group includes countries with a large share of households living in extreme poverty, including India, Bangladesh, Ethiopia, Mozambique, and Uganda.[12] This 42% of countries is currently calculating the timecount as their *de facto* poverty rate, depending on how they stratify samples.[13]

The largest group in the Smith et al. (2014) study (47% of the household consumption surveys in low and middle-income countries) interviewed households just once during the year, collecting data for just part of the year. This approach fails to deliver either headcounts or timecounts, and the results are not readily interpretable. However, by collecting data in particular periods, the poverty measures are sensitive to within-year conditions—directly, they reflect conditions in the survey period and, indirectly, they pick up the impact of other parts of the year via the process of consumption smoothing.

3.3 Sampling and instability in India's National Sample Survey

The Indian National Statistical Office implements one of the world's largest household surveys, the National Sample Survey (NSS). The NSS is the basis for official poverty estimates in India, and we examine data from the 68th round of the NSS, which was collected from July 2011 to June 2012—a period that overlaps the timing of the rural data we use in section 5.[14]

[12]The full list of countries, which comprise 41.7 percent of the total sample, are: Afghanistan, Armenia, Bangladesh, Brazil, Burkina Faso, Cambodia, Cape Verde, Chad, Ethiopia, Gambia, Georgia, Ghana, India, Iraq, Kenya, Lao PDR, Latvia, Lesotho, Lithuania, Malawi, Mongolia, Morocco, Mozambique, Nepal, Niger, Pakistan, Romania, Rwanda, Seychelles, Sierra Leone, South Africa, Sri Lanka, Swaziland, Tanzania, Thailand, Timor-Leste, Tunisia, Uganda, Vanuatu, and Yemen. The list is from the base data accompanying Smith et al. (2014), downloaded in October 2023 from http://www.ihsn.org/food.

[13]Smith et al. (2014), however, do not report on which surveys follow the guideline to stratify samples by period, which involves conducting interviews on different samples drawn from across the country each period. Some countries do not stratify by period but work their way through the country, finishing in a given region before moving to the next region, until data collection is finished. Without stratification each period (or without repeat visits to households), the data conflate time and place, and statistical agencies can calculate neither headcounts nor timecounts.

[14]India's data are important in a global perspective. We calculate that India's 21st century drop in poverty rates alone accounts for 40 percent of the global decline in measured poverty. Data from World Bank Indicators, using the current $3 a day global poverty line (at 2021 prices) shows that

Although methods have varied over time, data collection on household expenditures follows in the spirit of the expert guidelines in FAO and World Bank (2018), with data collected through the year (from July through the following June), stratified by quarter.[15] Using the notation of section 3, the sampling method delivers a single, randomly-chosen month of consumption, c_i^* for each household i, rather than their actual average consumption for the year, $\bar{c}_i$. Only with complete smoothing does $c_i^* = \bar{c}_i$.

Each household's randomly-selected monthly consumption value is then used when constructing national headcounts, using a poverty line that is also in monthly terms (Subramanian, 2012). On average in the 2011-12 NSS, 82 percent of the value of household expenditures in the NSS were collected with 30-day recall. More than half of this was food. The remaining 18 percent of expenditures by value were collected with 365-day recall, including durables like clothing as well as large one-off expenses like school tuition and hospital costs.[16]

To illustrate how short-term recall and one-time interviews in randomly-selected months affect measures of household consumption, we focus on two subgroups

between 2004 and 2022, India's poverty rate fell from 46.4% to 5.3% while the global rate fell from 30.4% to 10.5%. Multiplying by population data (and interpreting the data in terms of headcounts) shows that the number of people in poverty fell by 451.55 million in India and globally by 1,134.81 million—a ratio of 40 percent.

[15]The Indian National Sample Survey Organisation (2001) describes the process of stratifying annual data collection by sub-rounds within the year: "In the case of a number of characteristics, seasonality is a factor to be reckoned within data collection. For obtaining adequate reflection of the seasonal pattern of the country's economy in the collected data, the procedure of subdividing the total survey period into several sub-periods is adopted such that the uniform spread of field work over the round and also proper representation of different seasons in the data collected are ensured. The survey period of one year is divided into four or six equal sub-periods called sub-rounds. Normally an equal number of representative sample villages and urban blocks are allotted to each sub-round in such a manner as to obtain valid estimates for each sub-round." (Section 1.7, p. 2). Expert guidelines suggest a somewhat shorter recall period for food (one week rather than one month), but the NSS is largely aligned. For NSS reports, see the Ministry of Statistics and Programme Implementation's website: https://www.mospi.gov.in/download-reports. For methodological debate and the evolution of methods used in later NSS rounds, see, e.g., Deaton and Dreze 2002; Bardhan et al. 2017; Ghatak 2022; Roy and Van der Weide 2025.

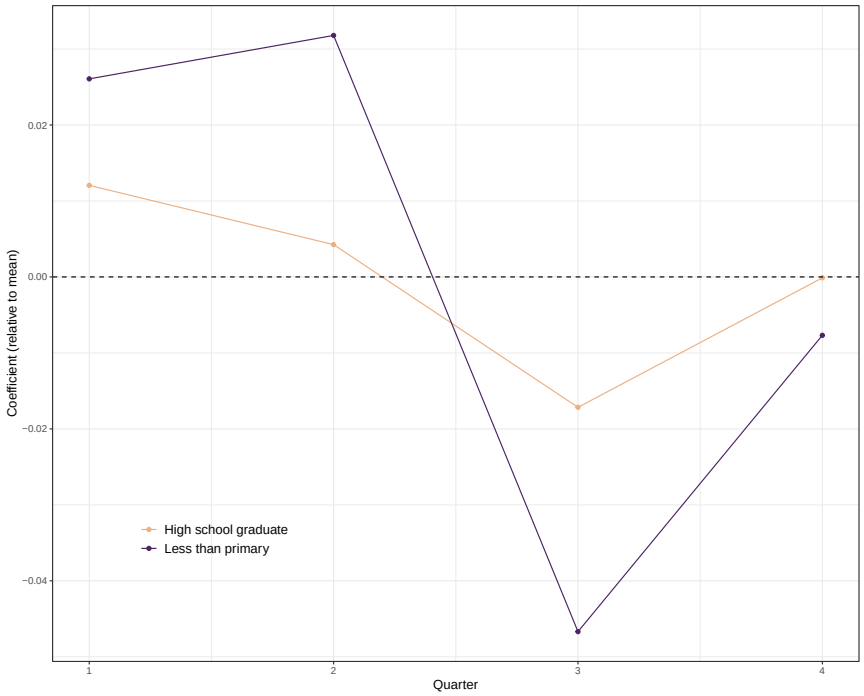
[16]In practice, the poverty rate being measured by household surveys depends on the ratio of households' expenditures on 30-day items relative to 365-day items. If household expenditure is completely comprised of 30-day recall items, poverty corresponds exactly to the timecount. The value of items reported with short-term recall accounts for a larger share of total spending for poorer households: The share of expenditures collected with 30-day recall was 86 percent in the lowest quartile versus 75 percent in the highest quartile. (The food share in total expenditures was 58 percent in the lowest quartile and 39 percent in the highest quartile.) For this reason, we continue to refer to this as the timecount for simplicity and elaboration of the concept.

with different socio-economic conditions: households whose heads completed high school and households whose heads who did not complete primary school. Households with more-educated heads have roughly three times the monthly expenditure of households with less-educated heads, based on the average expenditure across the year.[17] To illustrate variability, we use subgroup means of consumption averaged across the year as benchmarks and normalize each to zero, plotting them as the horizontal line at zero in Figure 1 (due to normalization, the two subgroups have the same benchmark at zero for the year). The horizontal axis marks the quarter of the year in which households were interviewed, and we plot the subgroup averages of monthly expenditure conditional on being surveyed in the given quarter.[18] The coefficients on the vertical axis show the log-difference in quarterly values relative to the overall yearly mean for the respective subgroup.

[17] Thanks to randomized, period-stratified sampling, the survey gives an accurate mean of average monthly consumption for subgroups even with one-time interviews with short-term recall (Scott, 1992). If the survey was collected for just part of the year or without stratified sampling, mean expenditures for the country as a whole would be biased.

[18] The 68th round of the NSS was collected from July 2011 through June 2012, and we arranged the qarters to follow the annual, January-December calendar. Figure 1 thus depicts Q3, Q4, Q1, and Q2 in the order of data collection.

Figure 1: Household expenditure by date of survey in the 68th NSS



Notes: The horizontal axis reflects the four quarters in which data are stratified. The values on the vertical axis pertain to households sampled in the given quarter. The vertical axis gives the average log difference of consumption for households interviewed in the given quarter relative to average consumption across the year (which is normalized at 0). Means are calculated using person weights, created by multiplying the household weight by the household size. All expenditures are deflated to July 2011 (the first month of the NSS 68) prices using the CPI.

As noted earlier, in a world of stability, where all households have completely smoothed consumption, data on c_i^* for each household i, sampled from a randomly-chosen period, would roughly reflect their average consumption for the year, $\bar{c}_i$. In that case, all the data in Figure 1 should cluster along the horizontal line at zero, showing no consumption variation across the year. This is because a random cross-section is drawn each period, and the average for each cross-section should be the same at any point in the year if households have completely smoothed away within-year variation (absent broader trends).

Figure 1, however, shows evidence of incomplete smoothing. Average monthly household expenditure varies systematically with interview dates, consistent with seasonality and illiquidity. To quantify the predictive power of interview timing, we regress (log) of "household expenditures" per capita on quarter-of-interview indicator variables for all households. The regression yields an F-statistic of 6.79

($p = 0.0001$).[19]

Figure 1 shows that the subgroup averages are above the annual means in the first two quarters and that average consumption sharply drops in the third quarter (July to September), the time of the summer monsoon, a pre-harvest period when agricultural jobs are relatively scarce. It is also the most challenging season for health as measured by anthropometric outcomes (Dwivedi et al. 2023, Johnston et al. 2021).

The drop for the subgroup with lower levels of education is much larger than for better-educated households, consistent with greater exposure to instability and weaker tools to smooth consumption. On average, the less-educated subgroup has both lower mean expenditure and greater variation across quarters. For the households with lower levels of education, the difference between the highest quarter (quarter 2) and the lowest quarter (quarter 3) is almost 0.08 in logs, indicating a (statistically significant) roughly eight percent difference in reported monthly household expenditures between surveys in the lowest and highest quarters. For the subgroup with higher levels of education, the difference from highest to lowest quarter is only three percent and not statistically significant. The result is consistent with greater consumption smoothing by better-off households relative to more disadvantaged households—who are doubly disadvantaged: once due to lower mean consumption and again due to greater temporal variation in consumption.

The figure thus illustrates how the expenditure data are conditioned by the randomly-assigned timing of the NSS interviews such that c_i^* does not approximate $\bar{c}_i$ across the year, especially for households less likely to be able to smooth consumption. We next show how this kind of variability transforms poverty statistics, turning what have historically been interpreted as headcounts into timecounts.

[19] The NSS was designed with seasonality in mind, and the figure reinforces the choice to collect data through the year on period-stratified samples.

4 Evidence from rural India

4.1 Village Dynamics in South Asia (VDSA) Survey

Because any given household is only interviewed in a single quarter, the NSS data in section 3.3 do not allow direct comparisons of headcounts and timecounts. For that, we turn to data for the same households over time, using the Longitudinal Village Level Studies of the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT) collected in India between 2010 and 2014.

The region depends on rainfed agriculture, and the year is marked by substantial seasonality, sharpening contrasts between headcounts and timecounts. The data are not representative of India, nor rural India as a whole, but we employ them to demonstrate mechanisms behind poverty analyses.[20]

The households in the sample were drawn as a random sample of the households in each area, stratified by landholdings. We combine modules on production activities, financial transactions, and household expenditure to construct monthly aggregates of expenditures, net income, and wealth for 1,526 households over 60 months, from July 2010 to June 2015. The households come from 30 villages across 15 districts in nine states. Approximately 94% households in the full sample self-identify as Hindu and the other 6% are divided between those who self-identify as Christians, Muslims, Sarnas, and others.

The data collection project, also known as Village Dynamics in South Asia (VDSA), provides statistics at the monthly level. Since the data are monthly, we shift from analyzing quarters (as in section 3.3) to months. Not all households are observed in all 60 months. In some regions, breaks occurred during the first quarter of 2012 and the first quarter of 2014. Additionally, households from the state of Telangana contain only yearly records from the 2014 wave. To create a panel data set with the greatest number of households possible, but with a balanced panel within each given year, we only include households for which we have four or five full years of monthly data.[21]

[20] Despite increasing urbanization globally (World Bank, 2021), Castañeda et al. (2018) find that rural residents comprised 80 percent of the world's population living in extreme poverty, and agriculture accounted for 65 percent of the employment of adults (aged 15 and above) in extreme poverty.

[21] We drop villages in the state of Telangana due to this restriction. However, there are relatively

The demographic variables are defined yearly – they are asked in only the July survey for each year – while the income and expenditure measures are monthly. We use a simple measure of household size, aggregating across all demographic groups in the table, to calculate per capita values for income and expenditures. Monetary measures are deflated to 2011-2012 rupees. The average household in the final sample includes slightly more than six individuals, with the most common demographic group being prime-aged males and females (between 15 and 59 years of age). The average household head is about 50 years of age, with an average of five years of education. The probability that the household head did not complete primary education is 50 percent for the sample with five years of data and 69 percent for the sample with four years.

Net income is a combination of production activities and wages. We do not observe interest paid/received for most financial transactions, so net income is mainly earnings from production and employment. In production activities, we include all the costs and revenues originating from cultivation, employment, and livestock. We record own agricultural income based on when the crop is sold or consumed, not when it is harvested. Importantly, because net income is a combination of revenues and costs, it can be negative in some months, especially during the agricultural planting season when costs are incurred but sales are still several months away. This prevents us from taking logs and from calculating certain poverty measures for income as discussed below.

Expenditures are our main interest, and they are more straightforward. Surveys are implemented each month and consumption is divided by whether it is home-produced, purchased, or received as a gift. Thus any agricultural production that is directly consumed ("own consumption") is given a value based on opportunity costs and aggregated in the total. We then sum across these categories.

Since we use per capita variables, we weight households by household size in order to interpret results as "per person" in the population from which the sample is drawn. In line with the stratification procedure used when collecting the data, we overweight landless households, multiplying the household size by 1.5 for the final sample weights.[22]

few households in Telangana, making up just 1.71 percent of the original sample.

[22] We thank Andrew Foster for providing us with information around the sampling design for these waves of the survey.

Column 1 of Table A1 of the appendix shows summary statistics for the 581 households (24,713 household-month observations) that we drop from the analyses – those with fewer than four full years of data. The second and third columns show statistics from the sample we use. They give data from balanced panels (four full years in column two and five full years in column three). Column 2 contains 116 households and column 3 contains 829 households, and most of the analysis is with these 945 households observed for 55,308 household-months. The second and third columns do not include records from Telangana, leaving 23 villages from eight states. Comparing the first and third columns shows that the excluded households (shown in column one) are poorer (annual expenditure is 28% lower) and less wealthy (42% lower wealth) than those with four or more full years of data.[23]

The food share in annual budgets is 50% for the full sample, 63% for the subsample below the poverty line, and 48% for the non-poor subsample (using yearly data and averaging across years). As we describe in Appendix section A3, durables and semi-durables are a small part of budgets and will matter more in other applications. Appendix section A4 explores expenditure volatility versus expenditure growth over time.

While in most household surveys from low-income regions, expenditure data is more accurate than income data since income tends to be under-reported, especially in rural settings (Carletto et al., 2021), here, agricultural income was a main focus of survey collection and was collected monthly. The measure of median income in colums 3 is 16.7 percent higher than median expenditures, but the means are closer. When we take the weighted average across columns 2 and 3, reflecting the full sample used in the analysis, the average difference between income and expenditure is 5 percent.

We use rural poverty lines by state reported by the Reserve Bank of India (RBI).[24] The poverty lines are in the range of the World Bank $1.90 per person per day (extreme) poverty lines using 2011 PPP prices, the global poverty line that prevailed during the study period.[25] For example, in 2011-12 the rural poverty lines in Andhra

[23]Table A2 in the appendix shows the number of households in our final sample by year and month.

[24]Reserve Bank of India (2021). Handbook of Statistics on the Indian Economy, 2020-21. Table 151 : Number and Percentage of Population Below Poverty Line. https://www.rbi.org.in/scripts/PublicationsView.aspx?id=20556

[25]As of June 2025, the World Bank uses a $3 per day poverty line and 2021 PPP ex-

Pradesh, Madhya Pradesh, and Gujarat were 860, 771, and 932 rupees respectively. The World Bank $1.90 per person per day line is $57 per 30-day month, which, using the 2011 PPP conversion rate to rupees (15.55) is 886 rupees, just above the Madhya Pradesh line.

The mode of per capita expenditure is very close to the poverty line, and many households are clustered on either side of the line. Specifically, 57.6% of the poor sample (and 17.0% of the entire sample) had annual expenditures between 75% and 100% of the poverty line. On the other side, 25.5% of the non-poor sample (and 17.9% of the entire sample) had annual expenditures between 100% and 125% of the poverty line. Inspection of the density (appendix figure A1) shows how for these groups, which together comprise 35% of the sample, poverty status is susceptible to shifts during the year.

When analysing poverty, we use the headcount, timecount, the distribution-sensitive measures of Watts (1968) and Foster et al. (1984), and the related distribution-sensitive and time-sensitive measures in equations (11) and (12). When using income, we only use the headcount poverty measure, but we calculate all measures when using expenditures.[26]

5 Results

5.1 Implications of time-sensitivity

5.1.1 Headcounts and timecounts

Table 1 presents population-weighted poverty summary statistics. The first column presents averages for the entire sample. The second and third columns present averages for households that are defined as poor or not based on their annual

change rates (https://www.worldbank.org/en/news/factsheet/2025/06/05/june-2025-update-to-global-poverty-lines). For comparison to the contemporaneous literature, we continue to use the $1.90 per person per day global threshold for extreme poverty (using 2011 PPP exchange rates).

[26] Since income can take negative values in some agricultural seasons, it is not possible to construct the Watts (1968) index with income on a monthly basis since it depends on logarithms. While it is technically possible to construct squared poverty gaps, the negative income values sometimes lead to very large estimates when squared. Because poverty in India is generally measured with household expenditure, and to avoid the problem of negative incomes, we focus only on expenditure-based measures (and calculate income-based headcounts for comparison).

consumption.

When measuring poverty by yearly household expenditure, poor households comprise 29% of the sample (row 2). As in section 2.4, the headcount can be understood as a timecount when consumption is completely smooth for each household. If households completely smoothed consumption, the timecount would thus also be 29%, but the timecount in row 3 shows that 37% of all household-months are spent in poverty, following equation (6).[27]

The increase over the headcount is largely due to the poverty of the "non-poor". "Non-poor" households, on average, spent 16% of their time below the poverty line.[28] Since non-poor households make up 71% of the sample (Figure A2 in the appendix), their months of poverty add substantially to the total. Just under two-thirds of households experience poverty in at least one month per year, including 47% of "non-poor" households. Altogether, 35% of all months in poverty are attributable to deprivations experienced by households that would not conventionally be considered poor.[29] At the same time, the experiences of "poor" households reduce the timecount: as a group, they spent just 86% of the year below the poverty line on average (equivalent to spending 1.7 months above the poverty line), reducing the timecount accordingly. The timecount thus reflects both better and worse times not captured by the headcount.[30]

Consistent with these findings, longitudinal quarterly data from Tajikistan show that only 10% of the sample was always poor across all 4 quarters while 40% of the sample was sometimes poor during the year (Azevedo and Seitz, 2016a). In

[27] Table A4 presents the same statistics but with expenditures smoothed for spending on durables, by reallocating durable spending equally across all months. The adjustment aims to bring the measure of expenditure closer to consumption by spreading the value of durable purchases beyond the month in which they were purchased. The adjustment leads to smoother patterns of consumption, but the effect is modest in the data. The overall patterns remain similar when comparing monthly aggregates and yearly aggregates. The timecount without adjusting for durables is 37%, for example, and it falls to 35 percent when adjustments are made. The fraction of households that are poor at least in one month falls from 63 percent (unadjusted) to 57 percent (adjusted). Given that these differences are small, we present the figures unadjusted for durables in the main results below and provide results with adjusted data in the appendix.

[28] More than a quarter of non-poor households in the sample experience at least one spell of poverty in any given year, where spells are defined as two or more months of poverty in a row, defined by monthly household expenditure.

[29] Table A5 in the appendix shows that these results are not driven by rising incomes over the sample period.

[30] In principle, timecounts are not necessarily larger than headcounts, following the logic of Ravallion (1988).

Table 1: Poverty summary statistics

	(1) Everyone	(2) Poor for the year	(3) Not poor for the year
Population share		0.29	0.71
Headcount	0.29	1.00	0.00
Timecount	0.37	0.86	0.16
Watts poverty index	0.09	0.30	0.00
Weighted timecount (Watts)	0.13	0.36	0.03
Squared poverty gap index	0.03	0.09	0.00
Weighted timecount (squared poverty gap)	0.04	0.11	0.01
Poor at least once in year	0.63	1.00	0.47
At least one poverty spell in year	0.51	1.00	0.27
Households	945	391	893
Month observations	55,308	12,300	43,008

Notes: Poverty is based on household expenditure. The first column includes all households. The second column includes only households who are poor for the entire year, using average monthly expenditures across the 12 months. The third column includes only households who are not poor for the entire year. Spells are two or more contiguous months in poverty. All statistics are weighted with population shares. The number of households in columns 2 and 3 sum to more than the total because the same household can show up in both columns.

Jordan, longitudinal data show that only 6% of the sample was always poor across 4 quarters while 33% was sometimes poor (Jolliffe and Serajuddin, 2018). Taking into account all quarters experienced in poverty, 35% were due to "non-poor" households falling below poverty lines for a time. Similarly, Dercon and Krishnan (2000) document movement in and out of poverty during the year due to uninsured shocks in three waves of data from Ethiopia in 1994-95, and Morduch and Schneider (2017) describe the prevalence of being "sometimes poor" in the United States.

5.1.2 Approximating the timecount with 10,000 replications

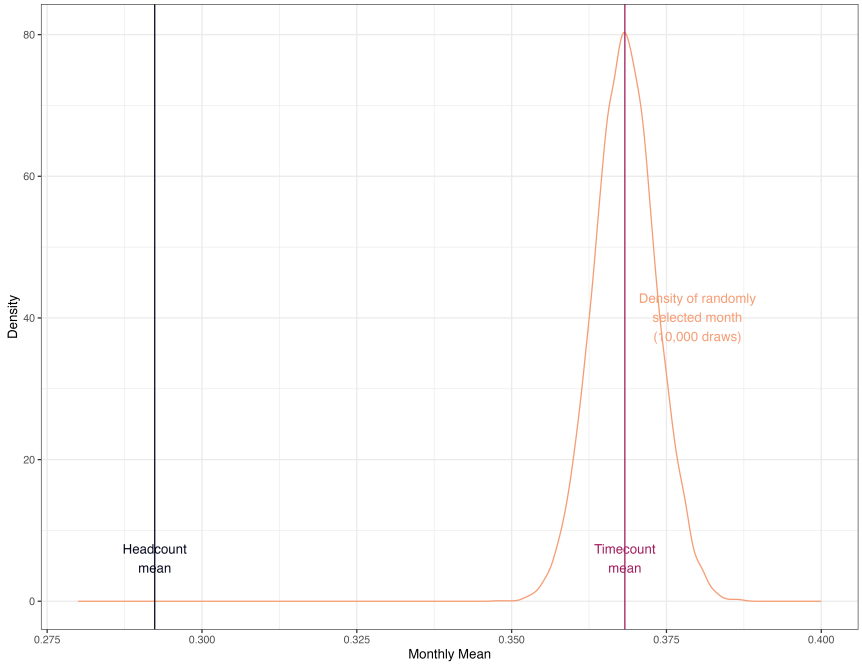
Figure 2 shows the difference between the headcount and timecount graphically. The bar on the left is the headcount measured using average expenditure for the entire year. As above, the headcount is centered just above 29 percent. To the right is the timecount based on monthly data, centered near 37 percent, showing the average number of months in poverty across the sample.

To demonstrate the logic in section 2.3, which shows how headcounts become timecounts, we estimate annual poverty rates using the standard formula for the

headcount in equation (2) but with data from a single, randomly selected month for each household, used as a proxy for the household's annual average consumption (as described in equation 7). The aim is to mimic poverty calculations when using data collected following the six methodologies guidelines in section 3.1. We then form a "headcount poverty rate" for the sample.[31]

The figure shows the density of the "headcounts" for the 10,000 replications. As shown mathematically in equation (10) of section 2.3, the average of the 10,000 replications converges to the timecount; across these 10,000 draws, the average is 0.368, which is identical to the true timecount mean out to three digits. These "noisy" proxies for annual consumption deliver approximate timecounts rather than noisy headcounts.

Figure 2: Comparisons of poverty measures



Notes: The "Headcount" mean is the annual poverty rate when using average expenditure for the entire year. The "Timecount" mean is the annual poverty rate when using the proportion of months in the year that the household is poor. The density estimate is estimated annual poverty rates when using a single, randomly selected month, across 10,000 random draws.

[31] Common practice in data collection is to stratify on quarter rather than month, but here we can take advantage of the ability to select random samples by month. We obtain a similar result when aggregating to quarters.

5.1.3 Distribution-sensitivity and the weighted timecount

With annual household consumption, distribution-sensitive poverty measures put extra weight on the poorest of the poor. Here, their analogues in the weighted timecount in equations (11) and (12) put extra weight on each household's most challenging periods (the poorest of the poor also get extra weight overall). When measured with yearly expenditure, the Watts (1968) index is 0.089 but rises by 40 percent to 0.125 in the weighted timecount version that encompasses within-year variation (equation 11). The Foster et al. (1984) squared poverty gap similarly rises from 0.025 to 0.037 when implementing equation (12), a 48 percent increase. The results show how seasonality particularly affects these poverty measures. While the indices are not cardinally-meaningful, the large changes suggest the possibility of substantial changes in relative rankings of distributions when comparing samples.

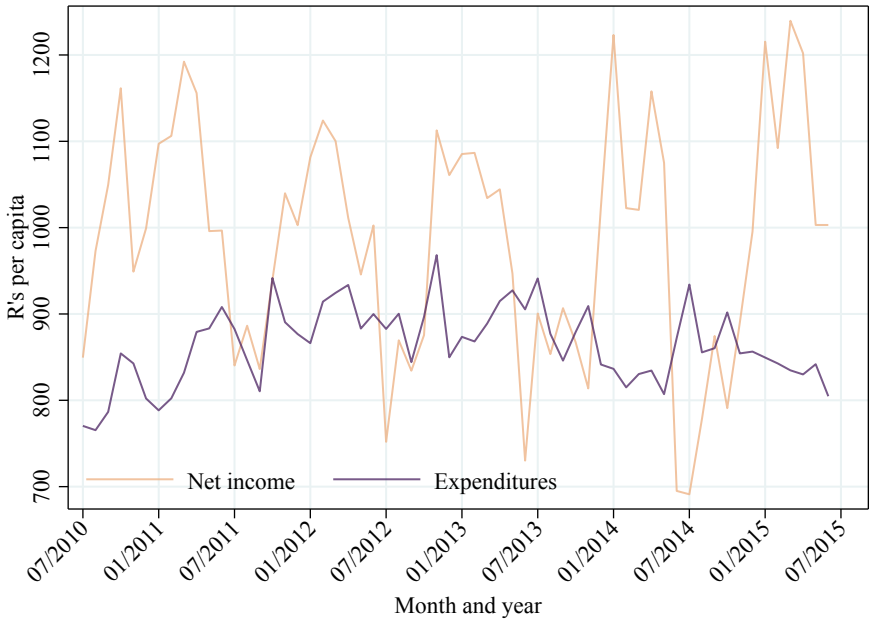
5.2 Income and expenditure variability

The sampling methods described in section 3 and, more generally, the difference between headcounts and timecounts rest on the extent of consumption smoothing. Figure 3 turns to the underlying within-year variability in income and consumption. Figure 3 shows the data on median per capita income and expenditure over time, from 2010 to 2015. Clear seasonal ups and downs mark the income data, which is considerably more variable than the expenditure data. Expenditure data is less variable in relative terms, but it is, nonetheless, absolutely variable.

Coefficients of variation (CV) of income and of expenditure are calculated for each household across the months of the survey in a given year and then averaged across households. The median CV of income is 0.86 and the median CV of expenditure is 0.25. The median ratio of the latter to the former is 27%. If there were no smoothing at all, the variability of month-to-month consumption would be identical to the variability of month-to-month income, and the ratio of their coefficients of variation would be 100%. If instead households could smooth consumption perfectly, the ratio would be 0. The median of the annual ratio is 27% across all households, indicating considerable but imperfect smoothing. The median CV of expenditure is 0.25 and indicates that households still experience considerable variability in expenditure; it is roughly the number generated if a household's

monthly expenditure was one quarter greater than the yearly average for half the year and one quarter less than the yearly average for the other half year.[32]

Figure 3: Median income and expenditures, 2010-2015



Notes: The lines are simple medians for each month, weighted by sampling weights and household size. We restrict the sample to households which show up in all five years so that the sample does not change across years.

5.3 The poverty exposure (PE) curve

Figure 4 shows how the experience of poverty (measured as each household's average months in poverty in a year) compares to their household's poverty status according to yearly resources. This is the variation underlying headcounts versus timecounts. To construct the figures, we take each household's per capita household income (left panel) and per capita household expenditures (right panel) averaged across all months of the survey. We then normalize by dividing by the poverty line. The variable y/z on the horizontal axis of the left panel is normalized income, with a value of one indicating that the household is exactly at the poverty line; similarly,

[32] An additional way to quantify variability is by regression of monthly expenditure on monthly income. Table A6 and Table A7 show the same patterns: household smooth imperfectly.

a value of 3, say, indicates that the household's per capita income is 300% of the poverty line. The variable C/z in the right panel is the equivalent for expenditures. The headcount is the share of the sample to the left of zero.

The y-axis gives the proportion of months that a household is in poverty, and it is the basis of timecounts. We call this the household poverty exposure (PE) rate. The scattered points give PE rates for individual households indexed by their average resources. The downward-sloping curves amid the points are "PE curves", calculated as local polynomial regressions of PE rates on total yearly resources. In terms of Figure 4, the headcount is the share of the population below 1 on the x-axis. The timecount is the average of y-axis values summed across the sample.

The PE curve provides another way to see that poor households are not always poor and non-poor households are sometimes poor. In the special case in which monthly expenditures are completely smoothed (i.e., monthly expenditures are always 1/12 of yearly expenditures)—the PE curve is a flat line at 1 on the y-axis (poor households are always poor) until it hits 1 on the x-axis (i.e., the poverty line), after which it drops to zero and proceeds as a flat line (non-poor households are never poor). With complete smoothing of expenditures:

$$PE\left(C/z\right) = \begin{cases} 1 & \text{if } C/z < 1 \\ 0 & \text{if } C/z \geq 1 \end{cases} \tag{13}$$

Figure 4 shows how far the data are from the pattern that would arise if equation 13 held in practice. The PE curve for income is downward-sloping, indicating that poverty exposure falls as households earn more overall, but earning more does not guarantee escape from exposure to poverty even when income is greater than twice the poverty line.[33]

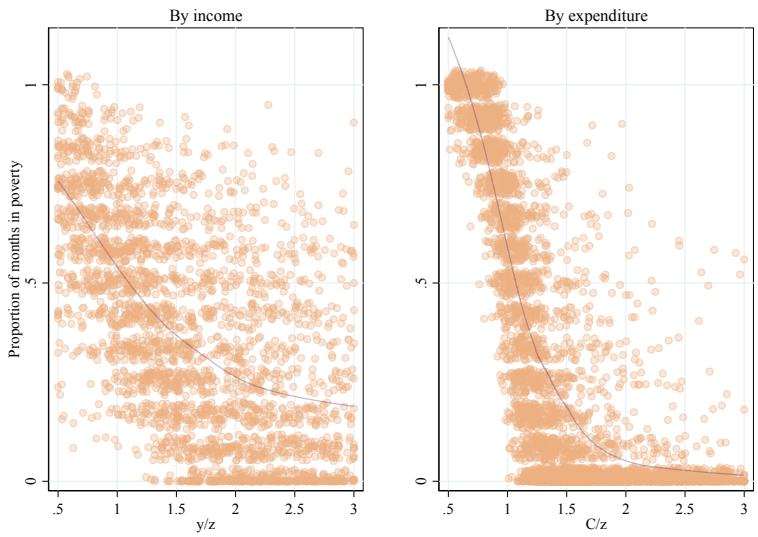
We quantify this with the smoothed value of poverty exposure at different ratios in the left panel. Right at the poverty line,[34] the smoothed value is 0.59, indicating that households near the line spend roughly 60 percent of the year in poverty based on their income. That decreases to 0.42 at 1.5 times the poverty line and 0.31 at twice the poverty line. In other words, poverty exposure stays quite high even

[33]The figures are restricted to households with yearly expenditure or income below 300% of the poverty line, but the PE curve is estimated for the full sample.

[34]Since there are no values exactly equal to one, we take the mean between 0.98 and 1.02. We use identical widths for the other values in this section.

when income is large relative to the poverty line. Because of seasonality, on average households with annual income double the poverty line still spend roughly one third of the year in poverty.

Figure 4: Months in poverty and annual income/expenditures



Notes: In both figures, the x-axis is the ratio of the monetary measure (income for the left figure and expenditures for the right figure) to the poverty line, averaged across the entire 60 months of the sample. The y-axis is the proportion of all months, across the entire sample, that a given household is in poverty. For ease of presentation, households below 0.5 and above 3 are dropped from the figure. The PE curve, shown by the smoothed curve through the middle of the scattered points, is a local polynomial regression of y on x.

The right panel shows poverty exposure by expenditures. By this measure, 53 percent of households are never expenditure-poor across the five years, and the data are distributed more compactly. The PE curve now slopes more steeply downward but still contrasts with the shape expected with perfect smoothing in equation 13.

Households are, on average, poor for slightly more than 60 percent of months right at the poverty line, but that number decreases to 23.6 percent at 1.5 times the poverty line and just 7.24 percent at twice the poverty line.

5.4 Which "poverty" best predicts health outcomes?

We considered the timecount first as a normative concept, but within-year variability also has predictive power to explain household outcomes of practical interest.

Table 2: Selecting the best predictors of anthropometrics through LASSO

	Weight		Height	
	All	lasso	All	lasso
Current year				
In poverty (headcount)	0.017		0.074	
In poverty (random month)	0.012		−0.003	
In poverty (timecount)	−0.020		−0.380	−0.004
Mean pov. spell length	0.005		0.039	
Log expenditures p.c.	0.071	0.029	0.108	
Previous year				
In poverty (headcount)	−0.002		−0.078	
In poverty (random month)	−0.010		−0.138	
In poverty (timecount)	−0.008		−0.578	−0.176
Mean pov. spell length	0.010		0.094	
Log expenditures p.c.	0.097	0.066	0.021	

Notes: Anthropometrics is only collected once each year in July. Each survey "wave" is from July to June of the following year. As such, poverty in the "current" year is actually in the future when considering anthropometrics. For this reason, the "current" poverty measure is for the previous 12 months, while the "lagged" poverty measure is for the 12 months prior to those months.

Using a least absolute shrinkage and selection operator (LASSO) procedure, we show that variation in the timecount is a stronger predictor of weight (for all individuals) and of height (for children under 20) relative to the predictive power of poverty status as defined by the headcount. The finding follows earlier studies that draw connections between seasonality and health outcomes (e.g., Christian and Dillon 2018) and evidence that even short term shocks experienced by pregnant mothers can have long-term consequences for children as they age (Barker et al., 2002).

The VDSA data include anthropometrics – weight and height – once per year for each household member.[35] Weight can change in relatively short time periods, for both children and adults. Height, on the other hand, takes longer to show changes due to changes in nutritional status and is generally applicable only to children. As such, we explore correlations of poverty measures for the previous 12 months ("current") as well as the 12 months prior to those ("previous"). Appendix Table A8 shows that the timecount and weighted timecount are more negatively correlated with both weight and height than their corresponding headcount.

[35]We note, however, that there are quite a few missing values.

Another way to see this is to use a least absolute shrinkage and selection operator (LASSO/lasso), a method designed to choose only the most predictive covariates. In Table 2, we include a range of covariates and let the lasso regression select the most predictive. In addition to the poverty measures in the previous table, we also include a poverty variable related to the lengths of poverty spells, which are defined as at least two contiguous months of poverty.

We present results for weight and height in Table 2. For both weight and height-for-age, the lasso regression selects exactly two variables. Interestingly, it picks current and previous expenditures for weight but current and previous timecount poverty for height. Since height is standardized by age, the coefficients can be interpreted in standard deviations. When the current timecount poverty measure increases from no months of poverty to 12 months of current poverty (zero to one on the indicator), predicted height is slightly lower; previous timecount poverty, however, shows a stronger predictive ability. These are not causal estimates, but they show an association between health and exposure to poverty as captured by the within-year variation in the timecount.

In other words, the evidence from the lasso regression aligns with our argument that measuring poverty at higher frequency reflects the experience of poverty in dimensions that are meaningfully different from poverty measured with yearly aggregates.[36] We note, however, that overall annual expenditures are predictive.

Table A9 in the Appendix restricts estimation to just children under 15 years of age. The model selects nothing for weight but lagged timecount poverty for height. Table A10 uses smoothed expenditures; the model selects the same variables as Table 2.

6 Evidence from China, Jordan, and Nigeria

Gibson (2001), Gibson et al. (2003), Gibson and Alimi (2020), and Jolliffe and Serajuddin (2018) provide important earlier insights on related issues. They take advantage of the few large household surveys with multiple observations on the consumption of the same households during the year, sharing the concern that

[36]The finding that the monthly poverty variable is quite predictive also suggests that – at least in our context – measurement error is not driving the main results of differences across months.

different methodological choices about data collection and extrapolation can lead to "noncomparable" poverty measurement (Jolliffe and Serajuddin, 2018). Each study shows deviations from the conventional poverty headcount that would result if an observation from just one randomly-chosen period of the year was used for each household, an idea described in section 2 and connected to the expected headcount. We reinforce their insights that (1) the timing and frequency of data collection can substantially affect poverty rates, even when the mathematical forms of poverty measures are standard across countries, (2) using data from one-time snapshots during the year introduces sensitivity to the timing of the surveys, and (3) as a result, poverty measures can be sensitive to the within-year shocks faced by households and their ability to smooth consumption.

We extend their work by introducing the timecount and weighted timecount as ethically and economically distinct concepts of poverty. This permits interpretation of what is *de facto* being measured by statistical agencies and reported as official poverty rates. That is, under specific conditions, *de facto* measures are not simply noisy headcounts but are approximate headcounts. In turn, we re-interpret the evidence in Gibson et al. (2003), Gibson and Alimi (2020), and Jolliffe and Serajuddin (2018) in terms of the timecount concept.

Gibson et al. (2003) use China's Household Income Expenditure Survey which is based on daily diaries for a full year, aggregated to form monthly sums. They use 1997 data from urban Hebei and Sichuan Provinces and exclude consumption of durables. They find an annual poverty rate of 30 percent when using all 12 months of data to calculate a headcount based on annual consumption. Gibson et al. (2003) then consider various approximations to annual consumption, including choosing data from a single randomly-chosen month and multiplying by 12 to predict annual consumption. This process is akin to equation (7) above, which we show leads to an approximation of the timecount as in equation (10). The Chinese poverty rate increases to 47 percent, a 53 percent increase in measured poverty. Under the same scenario but applied to the squared poverty gap of Foster et al. (1984) (which leads to an approximation of the weighted timecount in equation 12), measured poverty increases 3.5 times. An implication of the results is that, even in urban China, poor households appear to face substantial within-year consumption variability.[37]

[37]Gibson et al. (2003) speculate that the increases in poverty in the extrapolated data are driven by measurement error, but it is impossible to distinguish error from underlying consumption

Gibson and Alimi (2020) use a similar approach with data from Nigeria in 2012/2013, where data were collected twice during the year from the same households. Using both rounds of data, they estimate that the conventional headcount was 18 percent nationally (24 percent in rural Nigeria and 4.3 percent in urban). But using just one of the rounds and extrapolating as done in China by Gibson et al. (2003), they find large increases: now they estimate poverty of 37 percent nationally (47 percent in rural Nigeria and 15 percent in urban). As above, the latter rates can be seen as a version of the timecount, and they show an increase of 54 percent. Like the data from China, the numbers show substantial increases in urban areas. [38]

Jolliffe and Serajuddin (2018) use data from the 2010 round of Jordan's Household Expenditure and Income Survey, the source of official poverty statistics, which includes quarterly observations of household consumption. The data show that one third of the population was exposed to poverty in at least one quarter, even though the official poverty rate was 14.4 percent. Conditional on being "poor" based on annual data, Jolliffe and Serajuddin (2018) show that only 45 percent were poor for all four quarters. Another 37 percent were poor in all but one quarter, and 17 percent of "poor" households were poor in just two quarters. In parallel, two-thirds of the total population were "non-poor", but 22 percent of them were poor for at least one quarter.[39]

The official poverty rate of 14.4 percent is calculated as a conventional headcount by summing all four quarters of data to construct annual consumption. But when Jolliffe and Serajuddin (2018) use just one of the quarters (as in our exercise in Figure 2 and in Gibson et al. 2003 and Gibson and Alimi 2020), the rate—which can be interpreted as an approximate timecount—increases to 18.2 percent, a 26 percent increase. In their rural sample, the rate increases by 31 percent and in urban areas by 25 percent. When turning to the squared gap of Foster et al. (1984), the overall increase (which can be interpreted in the context of the weighted

variation.

[38] Gibson and Alimi (2020) recognize the policy implications: "In terms of policy, anti-poverty interventions that offer new consumption smoothing possibilities—such as the micro-lending component of the National Social Investment Program (N-SIP)—should matter more to the urban poor than the rural poor because the transient component of poverty is larger in urban areas. Such interventions may be needed because transfers through informal networks may not fully insure against shocks." (p. 103)

[39] Jolliffe and Serajuddin (2018) suggest that one alternative welfare measure could be a Rawlsian metric that registers welfare during households' most challenging periods during the year.

timecount) is 62 percent. As Jolliffe and Serajuddin (2018) note, these numbers show households' limited ability to smooth consumption during the year and that it can affect measured poverty.

7 Conclusions and future directions

7.1 Complementary concepts of poverty

Discussion of poverty statistics largely overlooks the within-year temporal dimension of deprivation. As Deaton (1997) writes, however: "In theory, we need to decide the reference period for welfare measurement, whether someone is poor if they go without adequate consumption or income for a week, a month, or a year." Even without a precise determination, the evidence above shows that in practice poverty measures in India and elsewhere already reflect timecounts—and thus already reflect short-term episodes of poverty within the year.

Considering poverty at a higher frequency than a year shows that the historical headcount of poverty based on consumption across a yearly reference period can be seen as an idealized version of welfare. The headcount based on annual spending leaves as a separate question how easily people can get hold of their money when needed. If poverty is defined as the condition that exists "when people lack the means to satisfy their basic needs," however, the concern with being able to earn sufficiently (the basis of the headcount) might reasonably be joined by concerns about the ability to save, borrow, insure, and spend sufficiently (as seen in the timecount). A question raised for policy and ethics is thus whether lacking "means" should also encompass lacking mechanisms to aggregate and move resources through time?[40] If the first UN Sustainable Development Goal is to "End poverty in all its forms everywhere" (United Nations, 2025), should this mean setting the global timecount to zero, rather than the global headcount (Merfeld and Morduch, 2026b)?

The research on year-to-year poverty dynamics highlights the role of insurance and safety nets in reducing poverty (Addison et al., 2009). The framework here

[40] Definition of poverty from Brittanica, https://www.britannica.com/topic/poverty. Written and fact-checked by the Editors of the Encyclopedia Brittanica. Accessed June 5, 2025.

goes further by showing how the specific timing of interventions during the year matters to the experience of poverty. For example, interventions targeted to lean seasons (e.g., Bryan et al. 2014; Imbert and Papp 2015) can have a greater impact on poverty timecounts than the same interventions implemented at other times of year. Similarly, interventions like microfinance that have been found to have relatively small average impacts on average household consumption (Cai et al., 2025) can reduce timecounts if households are enabled to smooth consumption within the year (Amin et al. 2003, Islam and Maitra 2012, Somville and Vandewalle 2023). Resources targeted to "non-poor" households can also reduce timecounts by targeting months of particular deprivation.

Timecounts and headcounts, taken together, reveal more than either alone, and they have different limits. The headcount measured with annual household consumption can show an improvement even when households' struggles are in fact worsening during part of the year—as long as the struggles are offset by periods of relative plenty.[41] Conversely, the timecount can show an improvement even when households are doing worse on average during the year.

While the timecount and weighted timecount recognize varied experiences of poverty, they fail to differentiate between persistent deprivation and isolated, short-term downturns. Moreover, they fail to distinguish between consumption variability that my be deliberately chosen (e.g., spending more in periods with festivals; Paxson 1993) versus variability driven by illiquidity. Analogues to these concerns have been raised in the context of temporal patterns of poverty across years (e.g., Foster 2009, Bossert et al. 2012, Dutta et al. 2013, Hoy and Zheng 2011), and they provide ways to systematically consider ethical concerns connected to durations of deprivation.

7.2 Practical steps

The analysis highlights that the meaning of poverty measures is defined not just by the mathematical forms of the measures but also by the data that is fed into them. A next practical step is to better catalogue global data collection practices, following the lead of Smith et al. (2014) in creating a survey of household surveys.

[41] This concern parallels Foster (2009)'s concern with perfect substitutability assumed in approaches to defining chronic poverty across years.

The step would bring together the theoretical and normative principles of poverty measurement with analyses of household survey methods. Elements to catalogue include the choices described in section 3.1: the variety of recall periods, the timing of surveys during the year, and the details of sampling and stratification. This step would permit a clearer assessment of what is being labeled and described as "poverty" globally and in individual countries.

A second step is to expand adherence to the expert guidelines for data collection already established in FAO and World Bank (2018) and Mancini and Vecchi (2022), practices already used in India and in surveys connected to the World Bank through its Living Standards Measurement Study program. This would require following the consensus to collect data throughout the year, draw nationally-representative samples in each period, and collect data with short-term recall, especially for food. Following the logic in section 2, this step would generate a clearer set of approximations to the timecount, provide data useful for extapolation and interpolation, and would eliminate measures that have no ready interpretation despite their labels.

A third step is to develop new approaches to calculating headcounts—which require accurate measures of each household's annual income or consumption, measures that are difficult to collect in the face of within-year instability and imperfect recall by households (Deaton and Grosh, 2000). One path is to take a modeling approach and develop statistical models that predict households' annual consumption (Merfeld and Morduch, 2026a), preferably based on higher-frequency longitudinal data for at least a subset of households (Scott, 1992; Gibson, 2001) along with data collected using the guidelines in section 3.1. Another path is to expand budgets for data collection so that all households can be surveyed multiple times during the year to create more accurate measures of each household's annual consumption (Deaton and Grosh, 2000).

Even with these steps, there remain empirical challenges that are generic to poverty measurement but which can have a larger impact in higher frequency data. One is measurement error which can exaggerate evidence of within-year volatility (Deaton and Grosh, 2000). The seasonal patterns in the all-India NSS data (Figure 1) and the finding in the rural data that within-year poverty predicts child health (section 5.4) shows that within-year variation is not just measurement error, but the broad concern with measurement error remains. A second challenge is the fact, well known to economists, that spending does not equal consumption, discussed

in section A3. Purchases of durables and semi-durables and some expenses on health, for example, create a wedge between the timing of spending and the timing of consumption, leading spending volatility to exceed consumption volatility. Bulk purchases and non-purchased consumption also deserve attention, especially the rental-value of housing and land (which, when included in household consumption, should attenuate overall volatility).

A connected problem involves the variability of needs. We have taken the annual poverty line and applied it as the threshold for minimal consumption throughout the year. This is the standard practice used by statistical agencies, and it is reasonable insofar as the fundamental material needs of life—-food, shelter, healthcare—are steady across time, but there may be cases in which needs vary meaningfully across seasons and time-varying poverty lines need to be considered.

Since the timecount is already the *de facto* poverty measure in India and other countries, these concerns are not hypothetical. They remain for future research and application.

8 Acknowledgements

We thank Kathleen Beegle, Dan Bjorkegren, Jonathan Conning, Indranil Dutta, John Gibson, Himanshu, Stephen Jenkins, Dean Jolliffe, Peter Lanjouw, and David Newhouse for discussion and feedback and are grateful for comments from seminar participants at the Indian Ministry of Statistics and Programme Implementation, the Delhi School of Economics, Indian Statistical Insitute-Delhi, NITI Aayog, the World Bank, BREAD, the Hong Kong University of Science and Technology, Tilburg University, Penn State University, University of Illinois, the 2022 Pacific Update Conference, the 5th IZA Labor Statistics Workshop, the National University of Singapore, Singapore Management University, IIM Calcutta, ANU Crawford, UNSW, UNU-WIDER, New York University, the 2022 NEUDC Conference, the 2023 IPA researcher gathering, and the 2024 Conference on Applied Development Economics in Lahore. Ingrid Leiria and Harry Keehoon Jung provided excellent research assistance. We thank the Mastercard Impact Fund, in collaboration with the Mastercard Center for Inclusive Growth, and the KDI School for research funding. We alone are responsible for all views and any errors.

References

Addison, T., D. Hulme, and R. Kanbur, eds. (2009): *Poverty Dynamics: Interdisciplinary Perspectives*, Oxford: Oxford University Press.

Amin, S., A. S. Rai, and G. Topa (2003): "Does microcredit reach the poor and vulnerable? Evidence from northern Bangladesh," *Journal of Development Economics*, 70, 59–82.

Atkinson, A. B. (2019): *Measuring Poverty Around the World*, Princeton, NJ: Princeton University Press.

Augenblick, N., B. Jack, S. Kaur, F. Masiye, and N. Swanson (2024): "Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Hunger," Working paper.

Azevedo, J. P. and W. Seitz (2016a): "How Persistent is Poverty in the Short Run?" World Bank Let's Talk Development Blog.

——— (2016b): "Is Poverty Seasonal?" World Bank Let's Talk Development Blog.

Balboni, C., O. Bandiera, R. Burgess, M. Ghatak, and A. Heil (2022): "Why do people stay poor?" *The Quarterly Journal of Economics*, 137, 785–844.

Bane, M. J. and D. Ellwood (1986): "Slipping Into and Out of Poverty: The Dynamics of Spells," *Journal of Human Resources*, 21, 1–23.

Bardhan, P., T. Srinivasan, and A. S. Bali (2017): "Poverty and Inequality in India: An Overview," in *Poverty and Income Distribution in India*, ed. by A. Banerjee, P. Bardhan, R. Somanathan, and T. Srinivasan, New Delhi: Juggernaut, 567–603.

Barker, D., J. Eriksson, T. Forsén, and C. Osmond (2002): "Fetal origins of adult disease: Strength of effects and biological basis," *International Journal of Epidemiology*, 31, 1235–1239.

Baulch, B. and J. Hoddinott (2000): "Economic mobility and poverty dynamics in developing countries," *Journal of Development Studies*, 36, 1–24.

Beegle, K., J. D. Weerdt, J. Friedman, and J. Gibson (2012): "Methods of household consumption measurement through surveys: Experimental results from Tanzania," *Journal of Development Economics*, 98, 3–18.

Bossert, W., S. R. Chakravarty, and C. D'Ambrosio (2012): "Poverty and Time," *Journal of Economic Inequality*, 10, 145–162.

Breza, E., S. Kaur, and Y. Shamdasani (2021): "Labor Rationing," *American Economic Review*, 111, 3184–3224.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1758.

Cai, J., M. Meki, S. Quinn, E. Field, C. Kinnan, J. Morduch, J. de Quidt, and F. Said (2025): "Microfinance," *VoxDevLit*, 3.

Carletto, G., M. Tiberti, and A. Zezza (2021): "Measure for Measure: Comparing Survey Based Estimates of Income and Consumption for Rural Households," *The World Bank Research Observer*, 37, 1–38.

Carter, M. R. and T. J. Lybbert (2012): "Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso," *Journal of Development Economics*, 99, 255–264.

Casaburi, L. and J. Willis (2018): "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya," *American Economic Review*, 108, 3778–3813.

Castañeda, A., D. Doan, D. Newhouse, M. C. Nguyen, H. Uematsu, and J. P. Azevedo (2018): "A New Profile of the Global Poor," *World Development*, 101, 250–267.

Chambers, R. (1983): *Rural Development: Putting the Last First*, Harlow, UK: Pearson/Prentice-Hall.

Christiaensen, L. and A. F. Shorrocks (2012): "Measuring Poverty Over Time [Special issue]," *Journal of Economic Inequality*, 10, 160.

Christian, P. and B. Dillon (2018): "Growing and Learning When Consumption Is Seasonal: Long-Term Evidence From Tanzania," *Demography*, 55, 1091–1118.

Coibion, O., Y. Gorodnichenko, and D. Koustas (2021): "Consumption Inequality and the Frequency of Purchases," *American Economic Journal: Macroeconomics*, 13, 449–82.

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2009): *Portfolios of the Poor: How the World's Poor Live on $2 a Day*, Princeton, NJ: Princeton University Press.

De Weerdt, J., J. Gibson, and K. Beegle (2020): "What Can We Learn from Experimenting with Survey Methods?" *Annual Review of Resource Economics*, 12, 431–47.

Deaton, A. (1997): *The Analysis of Household Surveys*, Baltimore, MD: Johns Hopkins University Press.

Deaton, A. and J. Dreze (2002): "Poverty and Inequality in India, A Re-Examination," *Economic and Political Weekly*, 3729–48.

Deaton, A. and M. Grosh (2000): "Consumption," in *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*, ed. by M. Grosh and P. Glewwe, Washington, DC: World Bank.

Deaton, A. and S. Zaidi (2002): "Guidelines for Constructing Consumption Aggregates For Welfare Analysis," Living standards measurement study working paper no. 135, World Bank.

Dercon, S. (2002): "Income Risk, Coping Strategies, and Safety Nets," *World Bank Research Observer*, 17, 141–166.

Dercon, S. and P. Krishnan (2000): "Vulnerability, seasonality and poverty in Ethiopia," *Journal of Development Studies*, 36, 25–53.

Devereux, S., R. Sabates-Wheeler, and R. Longhurst, eds. (2012): *Seasonality, Rural Livelihoods and Development*, London and New York: Earthscan/Routledge.

Dostie, S., S. Haggblade, and J. Randriamamonjy (2002): "Seasonal poverty in Madagascar: magnitude and solutions," *Food Policy*, 27, 493–518.

Dutta, I., L. Roope, and H. Zank (2013): "On Intertemporal Poverty Measures: The Role of Affluence and Want," *Social Choice and Welfare*, 41, 741–762.

Dwivedi, L. K., M. Bhatia, A. Bansal, R. Mishra, Shirisha P., S. Jana, S. Subramanian, and S. Unisa (2023): "Role of seasonality variation in prevalence and trend of childhood wasting in India: an empirical analysis using National Family Health Surveys, 2005–2021," *Health Science Reports*, 6, e1093.

Evidence Action (2019): "Seasonal Poverty: The Biggest Development Problem You Have Never Heard Of," Evidence Action blog.

FAO and World Bank (2018): "Food Data Collection in Household Consumption and Expenditure Surveys: Guidelines for Low- and Middle-Income Countries," LSMS guidebook, World Bank and FAO, Rome.

Fink, G., B. K. Jack, and F. Masiye (2020): "Seasonal Liquidity, Rural Labor Markets, and Agricultural Production," *American Economic Review*, 110, 3351–92.

Foster, J., J. Greer, and E. Thorbecke (1984): "A Class of Decomposable Poverty Measures," *Econometrica*, 52, 761–766.

Foster, J. E. (2009): "A Class of Chronic Poverty Measures," in *Poverty Dynamics: Interdisciplinary Perspectives*, ed. by T. Addison, D. Hulme, and R. Kanbur, Oxford: Oxford University Press.

Ghatak, M. (2022): "Introduction to e-Symposium: Estimation of poverty in India," *Ideas for India*, 10 October.

Gibson, J. (2001): "Measuring chronic poverty without a panel," *Journal of Development Economics*, 65, 243–266.

Gibson, J. and O. Alimi (2020): "Measuring poverty with noisy and corrected estimates of annual consumption: Evidence from Nigeria," *African Development Review*, 32, 96–107.

Gibson, J., J. Huang, and S. Rozelle (2003): "Improving Estimates of Inequality and Poverty From Urban China's Household Income and Expenditure Survey," *Review of Income and Wealth*, 49, 53–68.

Gilbert, C. L., L. Christiaensen, and J. Kaminski (2016): "Price Seasonality in Africa: Measurement and Extent," Policy Research Working Paper 7539, World Bank.

Hoddinott, J. (2006): "Shocks and their Consequences Across and Within Households in Rural Zimbabwe," *Journal of Development Studies*, 42, 301–321.

Hoy, M. and B. Zheng (2011): "Measuring Lifetime Poverty," *Journal of Economic Theory*, 146, 2544–2562.

Imbert, C. and J. Papp (2015): "Labor market effects of social programs: Evidence from India's employment guarantee," *American Economic Journal: Applied Economics*, 7, 233–263.

Islam, A. and P. Maitra (2012): "Health Shocks and Consumption Smoothing in Rural Households: Does Microcredit Have a Role to Play?" *Journal of Development Economics*, 97, 232–243.

Jalan, J. and M. Ravallion (1998): "Transient Poverty in Postreform Rural China," *Journal of Comparative Economics*, 26, 338–357.

Johnston, R., G. Dhamija, M. Kapoor, P. Agrawal, and A. de Wagt (2021): "Methods for assessing seasonal and annual trends in wasting in Indian surveys (NFHS-3, 4, RSOC CNNS)," *PLoS ONE*, 16, e0260301.

Jolliffe, D. and U. Serajuddin (2018): "Noncomparable Poverty Comparisons," *Journal of Development Studies*, 54, 533–536.

Khandker, S. (2012): "Seasonality of income and poverty in Bangladesh," *Journal of Development Economics*, 97, 244–256.

Longhurst, R., R. Chambers, and J. Swift (1986): "Seasonality and Poverty: Implications for Policy and Research," *IDS Bulletin*, 17, 67–70.

Lora, E., M. Benítez, and D. Gutiérrez (2024): "Annualizing labor market, inequality, and poverty indicators," *The Journal of Economic Inequality*, 22, 131–164.

Lønborg, J. H., M. Viveros, R. A. C. Aguilar, C. Lakner, G. L. Ibarra, and M. C. Nguyen (2026): "March 2026 global poverty update from the World Bank: New data and updated poverty numbers," Data blog, World Bank, Washington, D.C.

Mancini, G. and G. Vecchi (2022): "On the Construction of a Consumption Aggregate for Inequality and Poverty Analysis," Report, World Bank.

Merfeld, J. and J. Morduch (2026a): "Inequality and Volatility," working paper, University of Queensland and New York University.

——— (2026b): "What will it mean to end poverty?" *World Development*, 200, 107308.

Morduch, J. (1994): "Poverty and Vulnerability," *American Economic Review (AEA Papers and Proceedings)*, 84, 221–225.

Morduch, J. and R. Schneider (2017): *The Financial Diaries: How American Families Cope in a World of Uncertainty*, Princeton, NJ: Princeton University Press.

Mullainathan, S. and E. Shafir (2013): *Scarcity: Why Having Too Little Means so Much*, New York: Times Books.

National Sample Survey Organisation (2001): "Concepts and Definitions Used in NSS. (Golden Jubilee Publication)," Tech. rep., Ministry of Statistics and Programme Implementation, Government of India, Delhi.

Paxson, C. (1993): "Consumption and Income Seasonality in Thailand," *Journal of Political Economy*, 101, 39–72.

Pomeranz, D. and F. Kast (2024): "Savings Accounts to Borrow Less: Experimental Evidence from Chile," *Journal of Human Resources*, 59, 70–81.

Ravallion, M. (1988): "Expected Poverty under Risk-Induced Welfare Variability," *Economic Journal*, 98, 1171–82.

——— (2016): *The Economics of Poverty: History, Measurement, and Policy*, New York, NY: Oxford University Press.

Rodgers, J. R. and J. L. Rodgers (1993): "Chronic Poverty in the United States," *Journal of Human Resources*, 28, 25–54.

Roy, S. and R. Van der Weide (2025): "Estimating poverty for India after 2011 using private-sector survey data," *Journal of Development Economics*, 172, 103386.

Scott, C. (1992): "Estimation of Annual Expenditure from One-month Cross-sectional Data in a Household Survey," *Inter-Stat*, 8, 57–65.

Sen, A. (1976): "Poverty: An Ordinal Approach to Measurement," *Econometrica*, 44, 219–31.

Smith, L., O. Dupriez, and N. Troubat (2014): "Assessment of the reliability and relevance of the food data collected in national household consumption and expenditure surveys," IHSN working paper no. 8, International Household Survey Network.

Somville, V. and L. Vandewalle (2023): "Access to Banking, Savings and Consumption Smoothing in Rural India," *Journal of Public Economics*, 223, 104900.

Subramanian, S. (2012): *The Poverty Line*, New Delhi: Oxford University Press.

United Nations (2025): *The Sustainable Development Goals Report 2025*, New York: United Nations.

Vaitla, B., S. Devereux, and S. H. Swan (2009): "Seasonal Hunger: A Neglected Problem with Proven Solutions," *PLoS Med*, 6, e1000101.

Watts, H. W. (1968): "An Economic Definition of Poverty," in *On Understanding Poverty*, ed. by D. P. Moynihan, New York: Basic Books.

World Bank (2021): *Demographic Trends and Urbanization*, Washington, DC: World Bank.

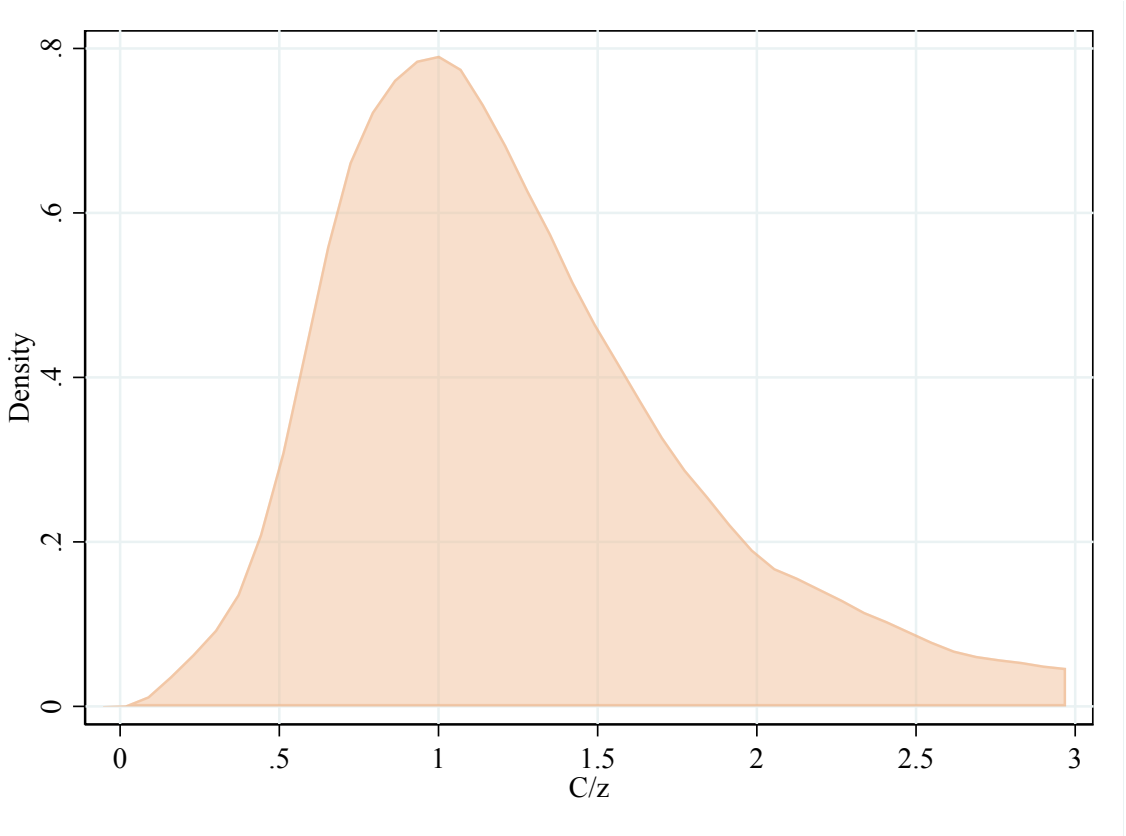
——— (2024): *Poverty, Prosperity, and Planet Report 2024: Pathways Out of the Polycrisis*, Washington, D. C.: World Bank.

Online Appendix

A1 Density of income

Figure A1 shows the estimated density of per capita expenditure for households observed for four years or more. The horizontal axis is annual expenditure per capita of households normalized by the annual poverty line, so households at 1 are exactly at the poverty line. Those to the left, below 1, are poor according to the headcount when using annual expenditure to assess poverty. Those above 1 are not poor by this measure.

Figure A1: Density



Notes: Kernel density estimate of per capita expenditure. The unit of observation is a household-month, so a single household appears multiple times in the data. The total number of household-month observations is approximately 55,000. The horizontal axis is annual expenditure per capita of households normalized by the annual poverty line. Households below 1 are poor according to annual data. The vertical axis is the probability density function.

Table A1: Summary statistics

	(1) Fewer than four years of data mean/(median)	(2) Four full years mean/(median)	(3) Five full years mean/(median)
Prime-aged females	2.031	1.493	1.900
	(2)	(1)	(2)
Prime-aged males	2.116	1.632	2.044
	(2)	(2)	(2)
Elderly females	0.256	0.246	0.293
	(0)	(0)	(0)
Elderly males	0.290	0.201	0.366
	(0)	(0)	(0)
Girls	1.046	0.581	0.880
	(1)	(0)	(1)
Boys	1.032	0.618	0.970
	(1)	(0)	(1)
Head is male (yes==1)	0.946	0.837	0.946
	(1)	(1)	(1)
Head age	48.488	48.124	51.351
	(47)	(47)	(50)
Head graduated primary	0.278	0.153	0.251
	(0)	(0)	(0)
Head graduated lower secondary	0.236	0.159	0.254
	(0.000)	(0.000)	(0.000)
Income p.c. (Rupees)	1128.197	1057.489	1190.137
	(844.083)	(1136.160)	(1004.460)
Expenditures p.c. (Rupees)	791.014	1366.117	1094.271
	(630.786)	(1026.164)	(860.767)
Wealth p.c. ('000s Rupees)	60.380	114.549	104.400
	(37.389)	(101.104)	(68.183)
Households	581	116	829
Month observations	24,713	5,568	49,740

Notes: Means and medians correspond to household-month observations. Households in the first column are dropped from subsequent analyses. Households in the second and third columns are included in all subsequent analyses. Households in the second column have four full years of observations, while households in the third column have five full years of observations. Data are from 2010-2015.

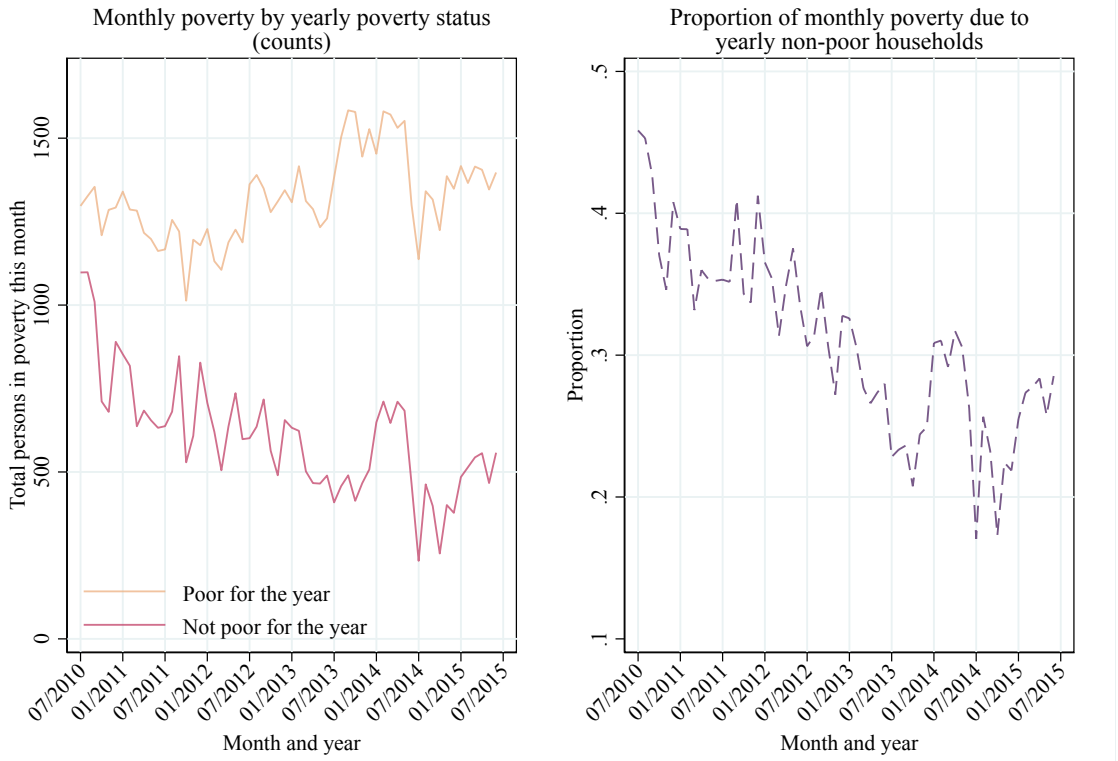
A2 Sample sizes

Table A2: Year-month sample sizes

	(1) 2010-2011	(2) 2011-2012	(3) 2012-2013	(4) 2013-2014	(5) 2014-2015
July	936	945	945	945	838
August	936	945	945	945	838
September	936	945	945	945	838
October	936	945	945	945	838
November	936	945	945	945	838
December	936	945	945	945	838
January	936	945	945	945	838
February	936	945	945	945	838
March	936	945	945	945	838
April	936	945	945	945	838
May	936	945	945	945	838
June	936	945	945	945	838

Notes: A "year" is defined as July to June of the following year. For example, column one is for 2010-2011 and include July-December of 2010 and January-June of 2011.

Figure A2: Total number of people in poverty, by month and poverty status



Notes: The left panel disaggregates households into those who are poor for the year versus those who are not, based on yearly expenditures. The right panel presents the share of household-months spent in poverty that are attributable to households that are not poor based on their yearly expenditure. All counts and proportions are weighted by population shares.

A3 Accounting for expenditures on durables and unusually large expenses

Durables and unusually large expenses pose complications when measuring poverty but especially at higher frequency than a year. Consider a household that purchases a bicycle, for example. Spending on the bicycle shows up in the data in the month it was purchased and leads to a large spike in spending. However, the actual consumption of the services of that bicycle may take place over the next several years. When measuring poverty, the ultimate interest is in consumption rather than

spending, but most surveys focus on spending (Coibion et al., 2021). A related issue involves big, unusual expenses like weddings. As Mancini and Vecchi (2022) write: "If a household spends a fortune on a special celebration during the survey period, such as a marriage, the resulting spike in measured consumption is genuine enough, but unrepresentative of typical living standards for that household" (p. 24).

The conventional method to address this problem is by excluding these kinds of lumpy and infrequent expenditures and restricting household consumption to more "regular" purchases.[A1] Table A3 shows that expenditures on large durables and semi-durables are in fact relatively low in the sample. We define "semi-durables" as clothing (clothes, shoes, and socks) and any item classified in the VDSA as "household articles and small durables (<2 years life)" in the survey. The table breaks out the percentage of monthly expenditures spent on durables (column 1) and semi-durables (column 2) by expenditure quartiles. (Since the headcount poverty rate with annual expenditures is 29 percent in the sample, the poverty line is close to the bottom of the second quartile.) Table A3 shows that in the bottom quartile, the median month includes no spending on durables or semi-durables. Even at the 90th percentile, there is no spending on durables and less than 10 percent of total spending is on semi-durables. The second quartile shows broadly similar expenditure patterns.

As additional evidence that spending on durables and semi-durables does not drive our results, Figure A3 shows the distribution of expenditures per capita when

[A1] Mancini and Vecchi (2022) note the problem with this procedure: "The choice of excluding lumpy expenditures is not, however, entirely uncontroversial. In all likelihood, exceptional expenditures will displace other spending (that is, a household will probably cut back on some of its other expenses in order to afford the big payment). The displacement will be greater for households that are unable to draw on savings or borrow, that is, poorer families and families having to shoulder large expenses that they have not had the chance to prepare for (as in the case of a catastrophic shock). The question, then, is whether spending net of the lumpy components is more typical than total spending. Arguably, if there is displacement, neither of the two measures—net or total—is representative of long run consumption; in fact, both are noisy proxies of it. Ultimately, because we do not observe long-run consumption, and we have no way to ascertain the size of the displacement of current expenditure, we cannot know for sure which of the two proxies is, in fact, the noisiest. A pragmatic strategy is to continue to exclude the shortlist of expenditures that are usually considered lumpy (e.g., weddings, funerals, purchase of durable goods), because they are typically very large with respect to the total budget of the household (and of the likely displacement they may cause), and that, at least to some extent, they were expected. The more a certain expenditure can be anticipated or planned for, the better is the case for its exclusion, as the observed consumption pattern discounts the occurrence of that expenditure."

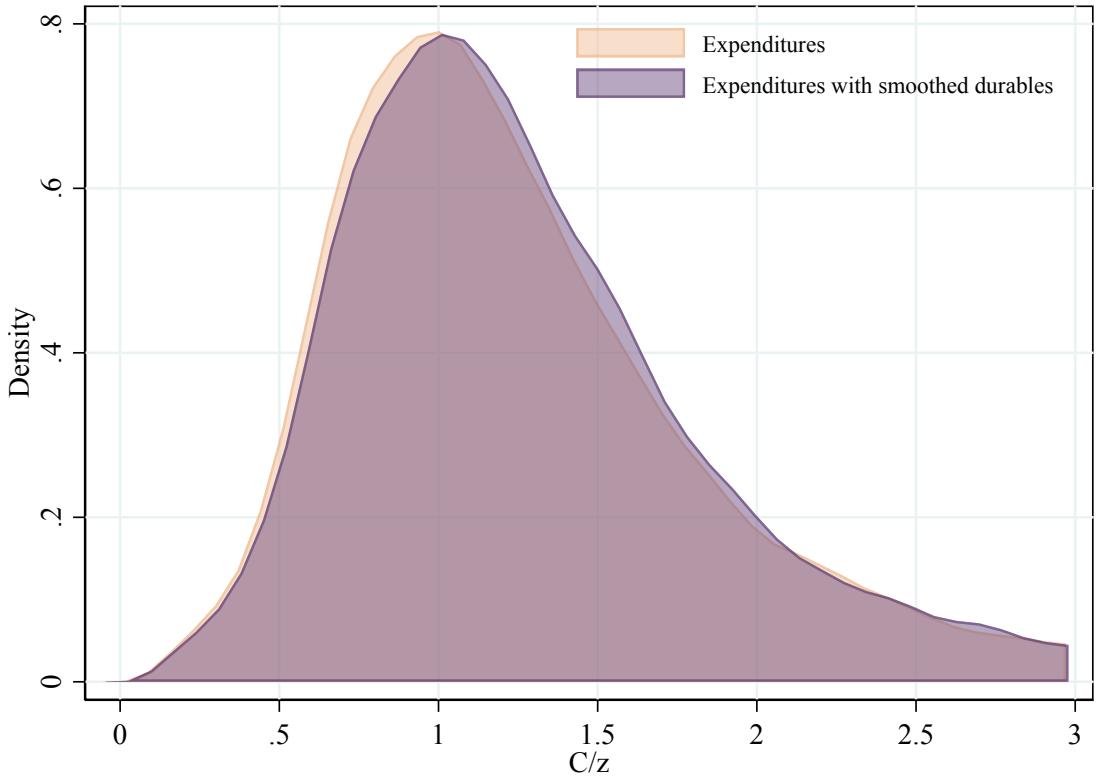
Table A3: Percent of expenditures on durable goods

	(1) Durables	(2) Semi-durables
Median		
Top expenditure quartile	0%	1.0%
Third quartile	0%	1.2%
Second quartile	0%	1.2%
Bottom expenditure quartile	0%	0%
90th percentile		
Top expenditure quartile	0%	14.0%
Third quartile	0%	13.8%
Second quartile	0%	12.7%
Bottom expenditure quartile	0%	9.6%

Notes: The percentages indicate the percent of monthly expenditures spent on each type of good. Expenditure quartiles divide the sample into four groups based on their total yearly expenditure per capita. Within each quartile, the columns give the median (top panel) and 90th percentile (bottom panel) of spending on durables and semi-durables (as a percent of monthly expenditures).

we smooth durable spending across an entire year. To create the figure, we subtract actual durable and semi-durable expenditure from total expenditures in each month and add one-twelfth of total durables/semi-durables expenditure for the subsequent 11 months plus the current month. The distribution almost completely overlaps the original expenditure distribution in Figure A1, and monthly poverty rates with the smoothed durables/semi-durables are still 19.3 percent higher than poverty measured at the yearly level. As such, spending on durables and semi-durables does not create large differences in estimated poverty rates in our context, though this type of spending may be important in other contexts. It is less clear whether there should be similar adjustments for spending on health, for which we do not have as detailed data.

Figure A3: Density with smoothed durables



Notes: Kernel density estimate of per capita expenditure. The unit of observation is a household-month, so a single household appears multiple times in the data. The total number of household-month observations is approximately 55,000. The horizontal axis is annual expenditure per capita of households normalized by the annual poverty line. Households below 1 are poor according to annual data. The vertical axis is the probability density function.

Table A4: Poverty summary statistics, expenditures smoothed for durables

	(1) Everyone	(2) Poor for the year	(3) Not poor for the year
Panel A: Large and small durables			
Timecount	0.347	0.873	0.129
Average of poverties (Watts)	0.116	0.349	0.019
Average of poverties (squared poverty gap)	0.034	0.108	0.004
Poor at least once in year	0.570	1.000	0.392
Panel B: Large durables only			
Timecount	0.359	0.864	0.150
Average of poverties (Watts)	0.122	0.359	0.024
Average of poverties (squared poverty gap)	0.036	0.112	0.005
Poor at least once in year	0.605	1.000	0.442
Households	945	391	893
Month observations	55,308	12,300	43,008

Notes: Poverty is based on household expenditure. The first column includes all households. The second column includes only households who are poor for the entire year, using average monthly expenditures across the 12 months. The third column includes only households who are not poor for the entire year. In the first panel, expenditures on large and small durables are allocated evenly across all months in the year. In the second panel, expenditures are smoothed for large durables only. All statistics are weighted.

A4 Expenditure growth or variable expenditure?

One possible explanation for the higher variance of monthly poverty for certain households is that their expenditures are simply growing. This would complicate the story we tell here. One way to see if growth is responsible for some of our results is to change the way we calculate the "annual" poverty measure. Instead of assuming that expenditures are identical in each month of the year, we can fit household-level trends and use the predicted values from these trends as the annual measure. We can then compare these results to the monthly expenditure results. If expenditure growth explains a large proportion of what we see here, then these new predicted poverty rates should be similar to the current monthly

results.

Table A5: Expenditure growth and predicted poverty rates

	(1) Headcount	(2) Pov gap	(3) Pov gap sq.	(4) Watts
Monthly measure	0.037	0.096	0.037	0.125
Trend measure	0.021	0.058	0.021	0.076
Annual measure	0.025	0.068	0.025	0.089

Notes: The trend measure is calculated by fitting a monthly trend, separately for each household, and using the predicted values from that trend as the poverty measure.

Table A5 shows that the trend poverty measure results in lower poverty than the current annual measure we use. Our concern was that income growth could explain the higher values we see, which would lead to similar poverty rates using the trend or the monthly poverty measure. While this does not seem to be a concern in the present context, we believe our method of comparison here is one that could prove fruitful elsewhere.

A5 Co-movement of monthly expenditure and income

Table A6: Co-movement of monthly expenditures and income, flexible lags and leads

	(1) Model 1	(2) Model 2	(3) Model 3
Current income	0.031***	0.032***	0.031***
	(0.004)	(0.004)	(0.005)
Household	X	X	X
Village-year-month	X	X	X
12 lags	X		X
12 leads		X	X
Observations	43,968	43,968	32,628

Notes: The dependent variable in all columns is monthly expenditures. "Current income" is monthly income. Lags and leads are for expenditures, not income. All standard errors are clustered at the household level.
* p<0.1 ** p<0.05 *** p<0.01

Table A7: Co-movement of monthly expenditures and income, by initial household wealth

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4
Current income	0.062***	0.045***	0.044***	0.045***
	(0.005)	(0.004)	(0.004)	(0.004)
Initial wealth (100,000k rupees)	306.336***			
	(44.349)			
Current income times initial wealth	−0.006**	−0.005***	−0.006***	−0.006***
	(0.002)	(0.001)	(0.001)	(0.001)
Fixed effects:				
Year-month	X	X	X	
Household		X		
Household-year			X	X
Village-year-month				X
Observations	55,308	55,308	55,308	55,308

Notes: The dependent variable in all columns is monthly expenditures. "Current income" is monthly income. Initial wealth is defined using the first wave of the survey and, as such, drops out of the regression when household fixed effects are included. All standard errors are clustered at the household level.
* p<0.1 ** p<0.05 *** p<0.01

A6 Anthropometrics and poverty

Table A8: Anthropometrics and poverty measures - Correlation matrix

	Headcount		Watts	
	Weight	Height	Weight	Height
Headcount (current)	-0.162	-0.113	-0.125	-0.051
Headcount (previous)	-0.175	-0.134	-0.130	-0.053
Timecount (current)	-0.199	-0.150	-0.143	-0.066
Timecount (previous)	-0.207	-0.166	-0.152	-0.078

Notes: Anthropometric data are only collected once each year at the start of the wave of data collection in July. (Each survey wave starts in July and ends in June of the following year.) As a result, the current year's values of income and expenditure cover a period after the anthropometric measurement, so poverty in prior years is most relevant for explaining anthropometric outcomes (so we consider lagged poverty measures only).

Table A9: Selecting the best predictors of anthropometrics through LASSO, only children (< 15)

	Weight		Height	
	All	lasso	All	lasso
Current year				
In poverty (headcount)	0.030		0.010	
In poverty (random month)	−0.003		−0.004	
In poverty (timecount)	−0.061		−0.291	
Mean pov. spell length	0.010		0.032	
Log expenditures p.c.	0.022		0.071	
Previous year				
In poverty (headcount)	0.002		−0.109	
In poverty (random month)	−0.012		−0.099	
In poverty (timecount)	−0.103		−0.815	−0.096
Mean pov. spell length	0.018		0.098	
Log expenditures p.c.			−0.210	

Notes: Anthropometrics is only collected once each year in July. Each survey "wave" is from July to June of the following year. As such, poverty in the "current" year is actually in the future when considering anthropometrics. For this reason, the "current" poverty measure is for the previous 12 months, while the "lagged" poverty measure is for the 12 months prior to those months.

Table A10: Anthropometrics with smoothed expenditures

	Weight		Height	
	All	lasso	All	lasso
Current year				
In poverty (headcount)	0.016		0.033	
In poverty (random month)	0.013		−0.029	
In poverty (timecount)	−0.033		−0.271	−0.055
Mean pov. spell length	0.006		0.025	
Log expenditures p.c.	0.067	0.029	0.105	
Previous year				
In poverty (headcount)	−0.001		−0.063	
In poverty (random month)	−0.010		−0.144	
In poverty (timecount)	−0.005		−0.622	−0.179
Mean pov. spell length	0.010		0.097	
Log expenditures p.c.	0.099	0.066		

Notes: Anthropometrics is only collected once each year in July. Each survey "wave" is from July to June of the following year. As such, poverty in the "current" year is actually in the future when considering anthropometrics. For this reason, the "current" poverty measure is for the previous 12 months, while the "lagged" poverty measure is for the 12 months prior to those months.