Poverty at High Frequency*

Joshua D. Merfeld† Jonathan Morduch‡

May 27, 2022

Abstract

High-frequency household data show that people with low yearly incomes also tend to face within-year economic instability and have limited ability to smooth consumption. We explore implications of the resulting within-year economic variation for the concept and measurement of poverty. We introduce a framework for defining annual poverty that takes into account variation in the extent and intensity of deprivation within the year. The framework expands on conventional conceptions of poverty which, by basing measurement on yearly household income or consumption, do not reflect within-year variation. We use the framework to analyze five years of monthly household data from rural India, a region marked by seasonality and vulnerability to extreme poverty. Because the experience of poverty of "non-poor" households is counted in the framework, measuring the headcount poverty rate by household-months in poverty increases the poverty rate by 26% relative to the conventional statistic based on yearly expenditure; 35% of months-in-poverty are attributable to deprivations among households that would not conventionally be considered poor. (2) The most vulnerable households see the largest increase in measured poverty when accounting for within-year variation. (3) Exiting poverty rarely means immediately leaving poverty behind. Almost half of all individuals who have "exited" poverty according to yearly measures nevertheless experience at least six months of poverty. (4) Policy that gives greater support during particularly challenging months lessens the experience of poverty more effectively than steady monthly transfers. We describe challenges for income and expenditure measurement and implications for public action.

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

---

*We thank Ingrid Leiria and Harry Keehoon Jung for excellent research assistance. We are grateful for comments from seminar participants at the Hong Kong University of Science and Technology, Tilburg University, and Penn State University.

†KDI School of Public Policy and Management and IZA; merfeld@kdis.ac.kr
‡Robert F. Wagner Graduate School of Public Service, New York University; jonathan.morduch@nyu.edu

1 Introduction

The economic definition of poverty simplifies a complex set of economic and social deprivations, boiling down deprivation to the condition of lacking sufficient resources relative to a minimum financial threshold. The time dimension is left unspecified, but one of the sharpest and most useful simplifications has been to measure poverty by yearly income or yearly consumption, a convention used since the late 19th century (Himmelfarb, 1984). For poverty headcounts this means classifying all households as poor or not based solely on total yearly resources.

There are compelling reasons to focus on yearly resources rather than varying monthly, quarterly, or seasonal conditions. One big advantage is the relative availability of annual data. The choice is especially appealing if income is received steadily through the year and households have little difficulty smoothing consumption. Having a low yearly income, however, is often accompanied by vulnerability to swings of income during the year. Research has long documented how illiquidity combined with agricultural seasonality shapes the economic lives of rural residents, for example, who comprise 80% of the world's population living in extreme poverty (e.g., Breza et al. 2021, Bryan et al. 2014, Devereux and Longhurst 2012, Longhurst et al. 1986, Khandker 2012).[1] This fluidity of poverty is not just a feature of agricultural economies. Low-income workers in non-farm settings, including the urban United States, must cope with income swings due to the varying availability of work through the year (Collins et al. 2009, Maag et al. 2017, Morduch and Siwicki 2017, Schneider and Harknett 2020, Storer et al. 2020). During the year, economic life is also shaped by events including changes in household composition, health shocks, and moving to new locations. In the United States, most people who are counted as

[1]Of the population measured as poor by the World Bank's \$1.90 per day per person extreme poverty measure, roughly 80 percent live in rural areas and are subject to seasonality (Castañeda et al., 2018). Globally, of all workers living on \$1.90 or less per day (aged 15 and above), 65 percent work in agriculture. Castañeda et al. (2018) estimate that in 2013, 770 million people lived in extreme poverty, and about 1 billion were moderately poor (living on more than \$1.90 per day but less than \$3.20). Castañeda et al. (2018) find that 76 percent of the people living in "moderate poverty" as defined by the World Bank live in rural areas, and 52 percent of workers who are among the "moderate" poor work in agriculture.

being poor in a given year are not poor for the full year. Others who are counted as non-poor are poor some of the time. Using the 2009-11 US Survey of Income and Program Participation, Edwards (2014) finds that half of all poverty spells lasted less than 7 months, and 44 percent of spells lasted just 2-4 months.[2]

We explore the consequences of bringing these within-year ups and downs into a framework for conceptualizing and measuring poverty. We describe a straightforward generalization of conventional poverty measures that defines annual poverty as the average of monthly poverty measures. For the headcount, for example, if a household has income below the poverty line for 9 months out of 12, its contribution to our poverty headcount rate is 0.75 of a year. The conventional approach, in contrast, counts the household as having experienced a full year of poverty (if yearly earnings are below the annual poverty threshold). We similarly incorporate distributionally-sensitive poverty measures.

Our aim is to capture seasonality and other forms of within-year variation as elements of annual poverty, rather than isolating episodic poverty like seasonal poverty as a concern distinct from poverty measured across the year. Since the rhythm of economic life is typically arranged around yearly cycles, aggregating to the annual level has intuitive appeal and the measures are comparable to other annual figures like GDP or life expectancy. Moreover, by maintaining an annual view, the framework avoids the concern that month-by-month snapshots may be largely driven by transitory phenomena (Atkinson, 2019).

We illustrate the framework with five years of monthly income and expenditure data from rural India. The data show that, in practice, conventional approaches can miss, and sometimes misinterpret, a large part of the experience of poverty. We estimate that in our sample the overall headcount poverty rate is 29% when measured conventionally with yearly consumption. If households experience no income variability and they perfectly

[2]Using data from the 1983-86 U.S. Survey of Income and Program Participation, Ruggles and Williams (1989) found that the median poverty spell lasted only 4-6 months. (Spells are two or more continuous months of poverty.) The picture was similar 25 years later. When Edwards (2014) looks at only 2011, she estimates that 8.3 percent of Americans were poor every month of the year, but about one quarter of Americans spent two or more months below the poverty line.

smooth consumption, the fraction of months in which households experience poverty should also be 29% (since monthly expenditure will be a constant proportion of yearly expenditure). However, we find that the poverty rate increases by 26% (to 37%) when taking into account monthly movements in and out of poverty during the year.

Two opposing forces explain the increase relative to the conventional headcount. The months-in-poverty measure is reduced by the fact that poor households (as classified by yearly consumption) actually spent just 86% of the year below the poverty line on avearge. But the measure is increased by the fact that "non-poor" households in the sample spent 16% of their time below the poverty line.[3] Since non-poor households make up 71% of the sample, their months of poverty add up. Altogether, 35.1% of all months-in-poverty are attributable to people who would not conventionally be considered poor. More than a quarter of non-poor households in our sample experience at least one spell of poverty in any given year.[4]

Incorporating distributionally-sensitive poverty measures into the framework reveals the varying intensity of deprivation, and the gap widens relative to conventional yearly measures. We find an increase of 42 and 48% in measured poverty when adapting the months-in-poverty measure to the Watts (1968) and Foster et al. (1984) squared-gap indices respectively. The relative increase is attributable to sensitivity to the intensity of poverty experienced in particularly low-consumption months. In addition, variability in household consumption for annually poor households leads to increases in these distributionally-sensitive measures, while this is not the case for the headcount.

Evidence Action (2019) describes seasonal poverty as "the biggest development problem

[3]Consistent with our findings, data from Tajikistan show that only 10% of the sample was always poor across 4 quarters while 40% of the sample was sometimes poor during the year (Azevedo and Seitz 2016a). Similarly too, Dercon and Krishnan (2000) explore poverty and seasonality with three waves of data from Ethiopia in 1994-95, finding considerable movement in and out of poverty during the year due to uninsured shocks. Morduch and Schneider (2017) describe the prevalence of being "sometimes poor" in the United States.

[4]Again, spells are defined as two or more months of poverty in a row, defined here by monthly household expenditure.

you have never heard of" and writes that "Seasonal hunger and deprivation are perhaps the biggest obstacles to the reduction of global poverty, yet they've remained largely under the radar." Evidence Action (2019) estimates that seasonal hunger affects around 600 million of the world's rural poor. Similarly, Chambers (1983) argues that the nature of rural poverty has long been "unperceived," even by experts. 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." Vulnerability to the ups and downs of resources within the year is thus both empirically important and often hidden by the aggregation of survey data.

The study builds from studies of poverty dynamics across years (e.g., Bane and Ellwood 1986, Jalan and Ravallion 1998, Baulch and McCulloch 2000, Addison et al. 2009, Christiaensen and Shorrocks 2012). Our main interest is to explore dimensions of poverty, but we also generate results on poverty dynamics at high frequency. The literature on poverty dynamics documents that households regularly move in and out of poverty from year to year, showing that much poverty is transient rather than chronic. Our approach shows that (1) households can experience regular ups and downs of poverty within the year while remaining chronically deprived across years. Thus, counter to common intuition about the nature of poverty, transience and persistence often exist together. As with agricultural seasonality, within-year instability is a stable feature of many people's lives. (2) The approach shows that exits from and entrances to poverty are seldom as sharp as implied by annual snapshots. In our sample, almost half of all individuals who have "exited" poverty according to yearly measures nevertheless experience at least six months of poverty during the year of "exit."

The generalized approach helps integrate insights on poverty measurement with new research on within-year instability and illiquidity, including literatures on seasonality and illiquidity (e.g., Fink et al. 2020, Breza et al. 2021) and seasonal hunger (e.g., Christian and Dillon 2018, Dostie et al. 2002). Our main results use household consumption as the basis for measuring poverty, thus reflecting outcomes after households have smoothed consumption to the extent they can. With data on income, we estimate the degree of consumption smoothing and relate it to poverty measures. Like much of the cross-year literature (e.g., Skoufias and Quisumbing 2005, Townsend 1994), we find substantial, but imperfect, consumption smoothing. If there was 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%. Instead, we estimate that the ratio is just 21%, indicating considerable (but imperfect) smoothing. Poor households (as measured by yearly consumption) have greater difficulty smoothing than non-poor households: When limiting our sample to households with yearly expenditure below the annual poverty line, the ratio of month-to-month consumption variation to month-to-month income variation rises from 21% to 31%.

The rise of experimental trials and the creative use of administrative data provide new evidence on the challenges that low-income households have in smoothing consumption within the year. Using data from a large US financial institution, for example, Ganong et al. (2020) show that households in the United States with low liquid wealth cut their consumption far more sharply than wealthier households when exposed to the same-sized income shocks during the year. The racial wealth gap can explain why Black households, on average, cut their consumption 50 percent more than white households in the face of similar income shocks. In a very different setting, agricultural communities in rural Zambia, Fink et al. (2020) document the prevalence of pre-harvest lean seasons and seasonal hunger, showing that limited liquidity forces poorer households to sell more labor, putting downward pressure on wages and reinforcing inequality.[5]

The framework complements conventional approaches in the way that distributionally-sensitive measures have broadened understandings without replacing the still-popular

[5]This relates to the relationship between wages and labor allocation documented by Jayachandran (2006).

headcount poverty measure. Like Sen (1976), Watts (1968), and Foster et al. (1984), the proposed generalization retains the yearly time frame. Distributionally-sensitive measures were introduced to account for the variability of resources across households in a given period. Our approach allows for variability of resources across sub-periods. Combined, the approaches provide a richer view of the experience of poverty than either alone.

The framework opens new perspectives on policy. Most directly, the framework quantifies how interventions that re-distribute resources between periods (or that make it easier for households to do so) can lessen the experience of poverty by improving consumption smoothing. This is so even when conventional poverty measures based on yearly resources are unchanged or worsening. In considering hypothetical monthly transfers, for example, we show that targeting transfers to the most challenging months (rather than spreading them through the year as in typical basic income and conditional cash transfer programs; Hanna and Olken 2018) can most cost-effectively reduce months-in-poverty (holding yearly transfer size constant). This raises questions about past interventions, like microfinance, that have had relatively small effects on total household consumption or income. If these interventions helped households move money across time, they may have had important effects on household well-being that were missed by focusing only on total consumption.

As high freqency data sets become more common—from surveys collected over multiple waves within years (e.g., Azevedo and Seitz 2016a, Dercon and Krishnan 2000), administrative data (e.g., Ganong et al. 2020), and financial diaries (e.g., Collins et al. 2009)—the framework will have broader application. In the final section, we discuss methodological challenges and ethical questions that arise when using high-frequency data to measure poverty, including accounting for spending on durable goods.

2 Framework

Measuring poverty usually involves answering two questions: Where should we set the poverty line and how should we create a poverty index? Economists have answered the second question by proposing alternative ways to weigh different degrees of deprivation across people (Watts 1968, Sen 1976, Foster et al. 1984). We focus on a third question that rarely gets asked: How should welfare be aggregated across time for individuals?

To the extent that time is considered, the issue is often framed as the choice between measuring poverty month-by-month or over a year (Atkinson, 2019). The high frequency measures give insight into time-specific deprivation, but transitory events can distort the view of overall conditions. The year-long period has the advantage of encompassing more time, but it requires extended recall for survey respondents, which brings its own distortions (Atkinson, 2019).

Our framework shifts the question. We consider poverty over a year – avoiding many of the important concerns discussed by Atkinson (2019) – but we aggregate across experiences of poverty within the year for the same households. We thus allow for within-year variation in the experience of poverty for each household, while keeping an approach based on annual averages.

We begin with the year divided into 12 months.[6] In each month t, household i earns income y_{it} and consumes c_{it}. The household's poverty status is determined in each month by the per-month poverty line z, the household's consumption, and the poverty mapping

[6]An alternative approach would divide the year into quarters, seasons, or other partitions. We choose months to conform to the form of the Indian VDSA data. The data we use – which we discuss in more detail in the next section – is collected monthly. There are detailed questions on expenditures and income for each month. Rather than aggregating these to some higher level, we instead choose to use the monthly set up of the data. This approach is flexible, however, and can easily be adapted to data collected on a different timeframe.

$P(c_{it})$:[7]

$$[P(c_{i1}), P(c_{i2}), P(c_{i3}), ..., P(c_{i12})], \tag{1}$$

where $P(c_{it}) = 0$ if $c_{it} \geq z$. The poverty levels can be aggregated across the time frame in various ways. For comparability with conventional practice, we retain the year-long time frame and define annual poverty as poverty status measured in each month, averaged across the 12 months and N households:

$$\frac{1}{12N} \sum_{t=1}^{12} \sum_{i=1}^{N} P(c_{it}), \tag{2}$$

where the poverty index is decomposable and differentiable. Calvo and Dercon (2009) and Foster (2009) fruitfully use a similar approach when considering the persistence of poverty across years. (The approach can alternatively be applied to quarters, seasons, or other periods.)[8]

We call this *average months in poverty* when the time frame is a year split into months. Equation 2 departs from standard practice by reflecting changes in the incidence and intensity of poverty within the time frame. For example, if a household is poor for 6 months of the year, their contribution to the aggregate poverty headcount is counted as 0.5 of a year of poverty.

The conventional practice of measuring *yearly poverty*, in contrast, focuses only on each household's total consumption over a year, with no accounting for variability within the year. This corresponds to a special case of Equation 2 in which poverty status in each

[7]The per-month poverty line z is assumed to be identical for all people and all periods. Poverty lines can be adjusted across space and time without changing the basic nature of the approach.

[8]The most commonly used poverty measures are decomposable and differentiable, including the headcount measure, the distributionally-sensitive Watts measure (Watts, 1968), and the Foster-Greer-Thorbecke class of measures (Foster et al., 1984). To simplify notation, we ignore population weights and weights for different long periods. Adding weights would be straightforward; for example, except in a leap year, poverty in January would contribute 31/365 to the weighted annual average, poverty in February would contribute 28/365, etc.

period is determined by household i's average monthly consumption, $\bar{c}_i$:

$$\frac{1}{N}\sum_{i=1}^{N} P\left(\frac{1}{12}\sum_{t=1}^{12} c_{it}\right) = \frac{1}{12N}\sum_{t=1}^{12}\sum_{i=1}^{N} P(\bar{c}_i), \tag{3}$$

where $P(\bar{c}_i) = 0$ if $\bar{c}_i \geq z$. In this case, the household that is poor for 6 months would count as being poor for the whole year, or as never being poor, depending on whether $\bar{c}_i < z$ or not. Equation 2 yields the same measure of poverty as equation 3 when each household i faces no instability or no illiquidity such that each and every month they consume the amount $\bar{c}_i$.

The connection between the approaches is seen by adding and subtracting Equation 3 to rewrite Equation 2:

$$\frac{1}{12N}\sum_{t=1}^{12}\sum_{i=1}^{N} P(c_{it}) = \frac{1}{12N}\sum_{t=1}^{12}\sum_{i=1}^{N} P(\bar{c}_i) + [P(c_{it}) - P(\bar{c}_i)] \tag{4}$$

The first term, $P(\bar{c}_i)$, reflects average consumption over the year, the focus of conventional poverty measurement. The second term within square brackets, $P(c_{it}) - P(\bar{c}_i)$, reflects the contribution of the months-in-poverty framework by capturing variation from the yearly average.

The notation allows analysis of how poverty is affected by changes in the economic environment—for example, the impact of the introduction of a cash transfer program, an increase in financial inclusion, or a tightening of labor markets. Taking the derivative of Equation 4 with respect to a change in an environmental factor x_t yields:

$$\frac{1}{12N}\sum_{t=1}^{12}\sum_{i=1}^{N} \frac{\partial P(\bar{c}_i)}{\partial \bar{c}_i}\frac{\partial \bar{c}_i}{\partial c_{it}}\frac{\partial c_{it}}{\partial x_t} + \left[\frac{\partial P(c_{it})}{\partial c_{it}}\frac{\partial c_{it}}{\partial x_t} - \frac{\partial P(\bar{c}_i)}{\partial \bar{c}_i}\frac{\partial \bar{c}_i}{\partial c_{it}}\frac{\partial c_{it}}{\partial x_t}\right] \tag{5}$$

The first term of the sum reflects the impact on poverty of an intervention x in period t. In each period, x may affect that period's consumption level and thus contribute to a change in average consumption, $\bar{c}_i$. An intervention that increases households' liquidity, for example, could spur investment and thereby reduce poverty by driving up average income and consumption during the year. This term captures the conventional focus of poverty analyses on totals and averages across the year. When observers say that microcredit has not reduced poverty(Banerjee et al., 2015), for example, they are implicitly saying that this term cannot be distinguished from zero.

The second term, within the square brackets, captures the impact via changes in the incidence and intensity of poverty from period to period. The term in brackets registers, for example, how increased liquidity may reduce poverty by allowing households to better protect their consumption during lean seasons by shifting resources from other seasons; how microcredit might help buffer health shocks (Berg and Emran 2020, Islam and Maitra 2012); or how saving groups might help smooth consumption (Beaman et al. 2014)—even with no change in total consumption across the year.

This second term represents the shift from the conventional conception of poverty. The notation helps show several implications of the framework. First, the conventional approach to measuring poverty (yearly poverty) reflects conditions that hold in one of two cases: (i) the special case in which there is no instability within the year (earnings, needs, and consumption are unvarying across periods), or (ii) the special case in which households face instability but have ample financial mechanisms to perfectly smooth away within-year instability. Equation 4 makes explicit that the standard approach to measuring poverty with yearly aggregates (reflected in Equation 3) is identical to the more general form in Equation 2 when consumption is completely smooth during the time frame; i.e., when $c_{it} = \bar{c}_i$ for all t. In this case, the term in square brackets is zero in Equation 4.[9] As noted above, prior work supports neither case (i) nor (ii) as a realistic assumption.

Second, and similarly, because the framework registers the impact of imperfect con-

[9]It is mathematically possible that Equation 2 is identical to Equation 3 even without perfect consumption smoothing, but it is unlikely. This is when, for example, the poverty mapping is completed with the headcount measure and there happen to be an identical number of months in poverty experienced by non-poor-on-average households as there are non-poor months experienced by poor-on-average households.

sumption smoothing during the year, Equation 5 shows that interventions that allow for re-distribution of resources between periods may reduce poverty as measured by Equation 2 even when aggregate resources are unchanged (or possibly falling). For example, relaxing liquidity constraints can raise households' consumption in bad months even if $\bar{c}_i$ is constant.

Third, less obviously, increasing liquidity can increase aggregate poverty as measured by Equation 2. This can happen in a particular (but realistic) circumstance in which the form of $P(c_{it})$ is the headcount and $\bar{c}_i < z$. Consider a household that is poor as measured by yearly resources but whose consumption is greater than the poverty line in a peak season. Improving the ability to save may reduce consumption in the peak season but expand resources available in the subsequent lean season. It is possible that the household will then count as being poor in both peak and lean seasons, whereas previously they counted as poor in just the peak season. Still, their revealed preference suggests that their well-being has improved by being able to save and smooth consumption. All is not lost, however, as distributionally-sensitive poverty measures would register the poverty reduction, even though the headcount does not.[10]

Fourth, by relating the experience of poverty to instability and the ability to smooth consumption, researchers can identify parts of the population that face particular de-

[10]To be more explict: The average of per-period poverty headcounts across the year may rise, for example, if resources are transferred out of period t where initially $c_{it} > z$ and afterward $c_{it} < z$. If resources are moved to period j where $c_{jt} < z$ before and after the transfer, the average headcount in periods i and j increases from 0.5 pre-tranfer to 1.0 post-transfer. With smoother consumption, the household's well-being may be improved and the average of distributionally-sensitive measures like those of Watts (1968) and Foster et al. (1984) may fall, but the average headcount will rise in this example. Here, if a household optimally smooths consumption in the sense of Jappelli and Pistaferri (2017), the average of distributionally-sensitive measures in periods i and j always fall given the assumptions, as long as the measures conform to the transfer axiom and sub-group monotonicity. This is true of both the Watts (1968) measure and squared-FGT measure (Foster et al. 1984). For example, suppose the poverty line is 60 USD per month and the household consumes 56 USD in 11 months and 80 USD in the final month. The average expenditure across all 12 months is 58 USD. Imagine the household gained access to a smoothing mechanism that allowed it to consume exactly 58 USD in each month. The measured monthly poverty would actuall increase (from 11 of 12 months to all 12 months). However, the distributionally-sensitive Watts (1968) index would actually decrease by around 64 percent and the squared poverty gap would decrease by almost 75 percent. The poverty gap, which is just the total amount of money the household is short of the poverty line in each poor month, would decrease from 44 USD (4 USD short of the poverty line in 11 months) to just 24 USD (2 USD short in all 12 months).

privation. In the sense of Equation 5, it becomes possible to identify a broader set of anti-poverty interventions. We show that households with the lowest average incomes and lowest average consumption over the year are also the households most exposed to intra-year volatility.[11] The extent of their challenges with poverty (and the implications of better-than-average periods) is unaccounted for in the standard measure of average yearly poverty (Collins et al., 2009). This also allows us to identify other characteristics that are distinct from average consumption but may also affect the ability to smooth consumption. For example, two households with the same average consumption may nonetheless have different consumption-smoothing capabilities depending on sector of employment, education, networks, etc.

Fifth, we assume here that poverty can be measured using data on household expenditures and consumption; in most countries, especially low-income economies, consumption data are generally higher quality than income data (Deaton 1997, Carletto et al. 2021). Those data may not be available, however. It could be that only data on income are available or that the income data are much more reliable than consumption data. Measuring poverty at a higer frequency heightens the distinction between income-based measures and consumption-based measures (Atkinson 2019, Bradbury et al. 2001). Measuring poverty with Equation 2 with income rather than consumption is equivalent to assuming that no consumption smoothing is possible: $c_{it} = y_{it}$ for all t. In contrast, measuring poverty with Equation 3 with income rather than consumption is equivalent to assuming that there is perfect consumption smoothing during the time frame: $c_{it} = \bar{y}_i$ for all t. Neither assumption accords with the imperfect smoothing (neither fully absent nor fully complete) observed in low-income communities, and high-frequency data on income and consumption can be used to quantify tradeoffs. Depending on the case, yearly

[11]For example, in our data, the coefficient of variation of household expenditures for households below the poverty line is 42.6%, while that number is 22.1% for households between the poverty line and two times the poverty line and is 14.5% for households between two and three times the poverty line. In fact, the standard deviation of expenditures for the poorest is actually higher than for households between the poverty line and two times the poverty line, despite the poor having mean expenditures of just 54% of the latter group.

income measures may yield a more accurate picture than monthly income measures, even if monthly consumption measures are preferred. In other words, measuring poverty with Equation 3 using income ($\bar{y}_i$) could yield results closer to those found when measuring poverty with Equation 2 using consumption (c_{it}), relative to measuring poverty with Equation 2 using income (y_{it}).

Sixth, most analysis of poverty dynamics uses multi-year panel data to quantify entry to and exit from poverty across years (Addison et al. 2009, Krishna 2016, Biewen 2014, Christiaensen and Shorrocks 2012, Valletta 2006, Baulch and McCulloch 2000, Bane and Ellwood 1986, Stevens 1999, Ravallion 2016). The multi-year analyses provide a natural extension to the literature on poverty and do not trouble conventional approaches that rely on yearly aggregates. The aim here is complementary. We integrate ways of thinking about poverty dynamics with understandings of (imperfect) intra-year consumption smoothing. The framework, for example, allows us a more granular examination of households' experiences of "entry" to and "exit" from poverty. "Exit" from poverty is typically defined as a transition from $\bar{c}_i < z$ in one year to $\bar{c}_i > z$ in the next year. It is thought of as a discontinuous break marked by the crossing of a threshold from one year to the next, but our high-frequency approach shows that households often continue to experience poverty, even if they are counted as having fully "exited" from poverty by standard definitions (and the converse is true for those who have "entered" poverty).

3 Data and methodology

3.1 Data

The data we use in this paper are from the Longitudinal Village Level Studies of ICRISAT from India between 2010 and 2014. This data collection project is also known as Village Dynamics in South Asia (VDSA). These data are unique because they bring an insight into the social and economic changes faced by households in rural areas of East India and SAT India at the monthly level. The information is separated into modules, each one providing a detailed collection of data on production activities, financial transactions, and household expenditure and consumption at a monthly or yearly level.

Using the different survey modules, we can aggregate expenditures, net income, and wealth of the households over 60 months, from July 2010 to June 2015. This composition allows us to explore the production activities, financial transactions, and expenditures of 1,526 households. These households come from 30 villages across 15 districts in nine different states. Approximately 94% households in the full sample are Hindu and the other 6% are divided between Christians, Muslims, Sarnas, and others.

However, not all households are observed in all 60 months of the survey. In particular, some breaks occurred during the first quarter of 2012 and the first quarter of 2014. Additionally, the households from the state of Telangana contain only yearly records from the 2014 wave, which does not permit looking over more than a year of uninterrupted records of household activities. Since we aim to create a panel data set with the greatest number of households possible – but with a balanced panel within a given year – we only include households for which we have four or five full years of monthly data. Table A1 shows observations across all 60 months.

Column one of Table 1 shows the summary statistics for the households that we drop from the analyses – those with fewer than four full years of data. There are 581 households and 24,713 household-month observations. In the second and third columns, we have balanced panels (four full years in column two and five full years in column three). Column two contains 116 households and column three contains 829 households. The second and third columns do not have any records from the state of Telangana, meaning they cover just eight of the states instead of all nine.

Table 1: 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 education (years)	5.359	2.927	5.210
	(5)	(0)	(5)
Income p.c. (R's)	1205.167	1387.972	1466.794
	(869.173)	(1059.050)	(1025.111)
Expenditures p.c. (R's)	791.014	1366.117	1094.271
	(630.786)	(1026.164)	(860.767)
Wealth p.c. ('000s R's)	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.

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 – ignoring adult equivalence scales – to calculate per capita values for income and expenditures. We deflate the monetary measures to 2011-2012 rupees. The average household in our final sample (columns two and three) has slightly more than six individuals across the six demographic groups, with the most common groups being prime-aged males and females, defined as those between 15 and 59 years of age. The household head is around

50 years of age, with an average of five years of education.

Net income is a combination of production activities and financial transactions. In production activities, we include all the costs and revenues originating from cultivation, employment, and livestock. In financial transactions, we include all the savings, remittances, benefits from the government, loans, and gifts that the households receive or spend on a monthly level. 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, for example, during the agricultural planting season. This prevents us from taking logs and from calculating certain poverty measures for income, which we discuss below.

Expenditures are more straightforward, as they are explicitly asked each month. They come in three separate categories – home produced, purchased, or gifted – but we take a simple sum across these categories. We are also able to disaggregate expenditures into food and non-food expenditures.

Wealth records include all the durables and assets (land, gold, and machinery) that the household owns. It is originally a yearly level variable, but it is expanded to the monthly level and the depreciation rate is calculated for each type of good. The wealth variable is highly right skewed, with a mean much higher than the median.

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 this sample is drawn. Importantly, this is not a random sample of rural households in India. However, the households in the sample are a random sample of the households in each area, stratified by landholdings. In line with the stratification, we overweight landless households, multiplying the household size by 1.5 for the final sample weights.[12]

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

3.2 Poverty measures

We use four separate poverty measures at different points in the paper: the poverty headcount, the poverty gap, the squared poverty gap, and the Watts (1968) index. The first three are specific cases of a more general Foster-Greer-Thorbecke (Foster et al. 1984)) – FGT – family of measures. The FGT family is defined as:

$$FGT(\alpha) = \frac{1}{N}\sum_{i=1}^{N}\left[\left(\frac{z-c_i}{z}\right)^{\alpha} \cdot 1_{c_i<z}\right] \tag{6}$$

where N is the total number of people in the population, z is the poverty line, and c_i is the monetary measure, either consumption or income. The indicator $1_{c_i<z}$ is one when households are poor in the given period and zero otherwise. The parameter α affects the relative weights of different individuals in the population. As α increases, the weight placed on the poorest gets larger.

The most common poverty measure is headcount poverty, which corresponds to $\alpha = 0$ – $FGT(0)$ – and is defined as

$$FGT(0) = \frac{1}{N}\sum_{i=1}^{N}\left[\left(\frac{z-c_i}{z}\right)^{0} \cdot 1_{c_i<z}\right] = \frac{N_{poor}}{N}, \tag{7}$$

where N_{poor} is the number of people in poverty. In other words, the headcount poverty rate is simply the number of people in poverty divided by the entire population.

The second poverty measure, the poverty gap, corresponds to $\alpha = 1$ and can be written as

$$FGT(1) = \frac{1}{N}\sum_{i=1}^{N}\left[\left(\frac{z-c_i}{z}\right) \cdot 1_{c_i<z}\right]. \tag{8}$$

Equation 8 shows the average amount of money – as a proportion of the poverty line – per person in the population needed to raise all households' consumption to the poverty line in time t. Unlike the headcount, the poverty gap registers households' deprivations relative to the poverty line, with the weight on each unit of money below the poverty line being constant. As a result, taking a unit of money from a very poor person and giving it to someone less poor does not change measured poverty.

The third poverty measure is the squared poverty gap, which is defined as $FGT(2)$:

$$FGT(2) = \frac{1}{N}\sum_{i=1}^{N}\left[\left(\frac{z-c_i}{z}\right)^2 \cdot 1_{c_i<z}\right]. \tag{9}$$

The measure is useful ordinally to rank poverty in different samples, but, unlike the headcount or poverty gap, it is not cardinally meaningful. However, it has the key of being distributionally sensitive. Here, taking a unit of money from a very poor person and giving it to someone less poor registers as an increase in measured poverty. Extra weight is placed on interventions that reduce extreme deprivation.

The fourth poverty measure we use is the Watts (1968) index, which is defined as

$$Watts = \frac{1}{N}\sum_{i=1}^{N}\left[\ln(z/c_i) \cdot 1_{c_i<z}\right] \tag{10}$$

Like the squared poverty gap, the Watts (1968) index is distributionally sensitive. This sensitivity increases only slowly at first, as income decreases from the poverty line, but then increases rapidly at the lower end of the distribution.

Figure A1 compares the differences in weights across these four measures. Note that the curves are scaled to allow their display on a single figure. As such, it is the relative shapes that are important, and not the levels, per se.

When we calculate conventional poverty measures based on yearly resources, we compare the monthly poverty line z_{month} to $\bar{c}_i$, the average monthly expenditure (or income) for the year. Using the squared poverty gap as an example and integrating it within our framework, we calculate

$$P(\bar{c}_i) = \frac{1}{12N} \sum_{t=1}^{12} \sum_{i=1}^{N} \left[\left(\frac{z_{month} - \bar{c}_i}{z_{month}} \right)^2 \cdot 1_{\bar{c}_i < z} \right] \tag{11}$$

This is equivalent to constructing a yearly poverty line and taking total yearly expenditure (or income) and redefining the sums appropriately. Throughout the paper, we use $P(\bar{c}_i)$ to refer to these yearly poverty measures.

We compare this to annual average months-in-poverty, which for the case of the squared poverty gap is defined as

$$P(c_{it}) = \frac{1}{12N} \sum_{t=1}^{12} \sum_{i=1}^{N} \left[\left(\frac{z_{month} - c_{it}}{z_{month}} \right)^2 \cdot 1_{c_{it} < z} \right], \tag{12}$$

where c_{it} is expenditure in month t for household i. Throughout the paper, we use $P(c_{it})$ to refer to these monthly poverty measures.

Due to the construction of the measures, we use only headcount poverty for income, but calculate all four measures for expenditures.[13]

3.3 Methodology

Many of the results we present below are descriptive in nature. We discuss methodology for the other main results in turn.

3.3.1 Regressions

Our first set of regressions tests for consumption smoothing because this assumption is key to the difference between traditional, annual poverty measurements and our proposed

[13]Since income can take on negative values in some agricultural seasons, it is not possible to construct the Watts (1968) index with income on a monthly basis. 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).

measure. If households are smoothing consumption perfectly, then income and expenditures should not covary within households. We test this by estimating regressions of the form,

$$c_{it} = \alpha + \beta y_{it} + \varepsilon_{it}, \tag{13}$$

where c_{it} is monthly expenditures, y_{it} is monthly income, ε_{it} is a mean-zero error term, and α is a set of fixed effects. We formally test the hypothesis that, under perfect smoothing, $\beta = 0$.

We vary the fixed effects across specifications. We consider three different fixed effect specifications. First, we include year-month and household fixed effects. Across households, monthly income and expenditures can be correlated, since higher income households also tend to spend more money. However, with the household fixed effects, we restrict identification to only changes in income and expenditures within the household.

Of course, income and expenditures can covary within a household if permanent income changes. Consider, for example, if a household enters into a new type of employment that increases their expected income. Then, their income and expenditures may move together, even if they are smoothing perfectly. This motivates our second fixed effect specification, which replaces the household fixed effects with household-year fixed effects. This decreases the window across which we are identifying coefficients from 60 months to 12 months. As additional robustness checks, we include flexible lags and leads of income to even better capture possible changes in expected income. We present these results in the appendix.

Finally, we replace the year-month fixed effects with village-year-month fixed effects. This addition is inspired by Townsend (1994), who tested for collective risk sharing and insurance in villages in India. Here, however, we are interested in partialling out covariate shocks. In other words, if the totality of the covariance between expenditures and income are driven by village-level covariate shocks, then $\beta = 0$ in this specification even if it did not equal zero in previous specifications. This is not a test for consumption smoothing, per se, but instead is meant to better understand how much covariate shocks may drive deviations from consumption smoothing.

After the consumption smoothing tests, we turn to relationships between different poverty measures:

$$P(c_{it}) = \alpha + \beta P(y_{it}) + \gamma P(\bar{c}_i) + \varepsilon_{it}, \tag{14}$$

where $P(c_{it})$ is the expenditure-based months-in-poverty measure, $P(y_{it})$ is the income-based months-in-poverty measure, and $P(\bar{c}_i)$ is the expenditure-based yearly poverty measure. Our aim here is to show how the inability to smooth consumption translates to poverty measures. These regressions are, effectively, tests for consumption smoothing, except with poverty measures instead of monetary variables. The same intuition around the fixed effects in the previous subsection apply here.

3.3.2 Monetary comparisons of monthly and yearly poverty

One of our goals is to compare yearly poverty measures to monthly poverty measures. This is somewhat straightforward with headcount poverty, since a simple comparison of rates is relatively intuitive. However, especially for the squared poverty gap and the Watts (1968) index, such a comparison is more complicated. To aid intuition of the changes, we calculate monetary comparisons using what we refer to as "implied" yearly income. Consider the monthly squared poverty gap described in Equation 11 and Equation 12. Using the monthly expenditure values, we can calculate the monthly poverty measure, $P(c_{it})$. We then plug this value into the yearly definition (i.e. Equation 11), as such:

$$P(c_{it}) = \frac{1}{12N} \sum_{t=1}^{12} \sum_{i=1}^{N} \left[\left(\frac{z_{month} - \bar{c}_i}{z_{month}} \right)^2 \cdot 1_{\bar{c}_i < z} \right]. \tag{15}$$

Inverting this equation and solving for $\bar{c}_i$ gives us the yearly expenditures/income implied by the monthly poverty measure, $\tilde{c}_i$. We construct a ratio of this income to actual yearly income and use this as a measure of the change when going from yearly to monthly poverty measures:

$$\text{implied income} = \frac{\tilde{c}_i}{\bar{c}_i} \tag{16}$$

3.3.3 Policy experiment

Our final set of results is a policy experiment. One of the implications of our framework is that focusing on yearly measures can miss important changes in consumption smoothing within the year. This policy experiment is meant to bring this result into focus.

We imagine a hypothetical government transfer to households of 960 rupees per capita per annum (80 rupees per capita per month). This is approximately 7.2% of average per capita expenditures in the entire sample, or 14.8% of the average per capita expenditures of the poor. For simplicity, we design this transfer to go only to those households living below the poverty line.[14]

We vary how these 960 rupees per year are allocated across months. We compare the resulting poverty rates from four separate allocation designs: no transfer at all, a transfer of 80 rupees per month across all months, a transfer of 160 rupees per month across six months, and a transfer 320 rupees per month across three months. We assume that the totality of this transfer is consumed in the month of the transfer and then examine how estimated poverty measures change in response to these different designs. Specifically, we transfer these amounts to the lowest relevant months. In other words, for the transfer across six months, we choose the poorest six months. Similarly, for the transfer across

[14] In reality, such a design would present perverse incentives for households living just below the poverty line. Since we do not taper the transfer, those just below the poverty line can actually end up with a higher income than those just above the poverty line. However, we believe the simplicity of this design allows for a more straightforward elaboration of the results.

three months, we choose the poorest three months.

Recall the discussion around Equation 5 of how overall changes in months-in-poverty respond to a change in an external factor can be decomposed into two parts: the effect of the external change on average expenditures and the effect of the external change on the deviation of monthly expenditures around that average. Throughout this exercise, the total amount of the transfer is unchanged. This means that average monthly expenditures are unchanged across the three separate allocation designs and, as such, that the effect on yearly poverty measures are unchanged. We focus on how this transfer allows households to change the distribution of resources throughout the year. For example, while a microfinance initiative may not lead to increases in yearly income, it may allow households to move resources across time, affecting how monthly expenditures vary around the mean. The focus on yearly incomes (or income in a single month) misses this change, despite it being a clear increase in welfare.

4 Results

4.1 Tests for consumption smoothing

We only expect changes in poverty rates when moving from yearly values to monthly averages when households are not able to perfectly smooth consumption across time. As such, we first present basic empirical tests for consumption smoothing before moving on to changes in estimated poverty rates.

Table 2 presents these results. The dependent variable is monthly expenditures and the independent variable is monthly income. Intuitively, if households are perfectly smoothing consumption (proxied here with expenditures), then monthly consumption should not vary with monthly income. Of course, failure to find a (significant) correlation between the two does not necessarily imply that households smooth consumption perfectly. Given

Table 2: Monthly expenditures and income – Tests for consumption smoothing

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4	(5) Model 5
Current income	0.058***	0.034***	0.032***	0.033***	0.033***
	(0.006)	(0.004)	(0.004)	(0.004)	(0.004)
Average exp.				0.209***	
(last 12 months)				(0.057)	
Fixed effects:					
Year-month	X	X	X		
Household		X		X	
Household-year			X		X
Village-year -month				X	X
Observations	55,308	55,308	55,308	43,968	55,308

Notes: The dependent variable in all columns is monthly expenditures. "Current income" is monthly income. All standard errors are clustered at the household level.
* p<0.1 ** p<0.05 *** p<0.01

the measurement error known to exist household income and expenditure calculations, one possibility is that the noise is so large relative to the signal that we cannot reject zero. However, rejections of no correlation are consistent with a failure to smooth consumption across time.

This consumption smoothing result is not necessarily true in the cross section, as households with higher incomes are more likely to have higher expenditures, on average. This result is shown in column one, which includes only year-by-month fixed effects. Indeed, higher monthly incomes are correlated with higher monthly expenditures.

Column two adds household fixed effects, which remove variation across households and instead focus on changes in monthly income and monthly expenditures within households. We strongly reject no correlation between monthly income and monthly expenditures within households over the course of the sample. However, the coefficient in column two is decidedly smaller than the coefficient in column one, as we would expect given the previous conversation.

Even if households are smoothing consumption, we might see a correlation between current income and current expenditures if permanent income changes. We implement two alternative specifications in Table 2 in an attempt to rule out this possibility. First, column three includes household-year fixed effects – instead of household fixed effects. Looking at just 12 months will reduce the probability that household permanent income changes markedly. In fact, consistent with the results in column two, the coefficient barely changes and remains highly significant.

Column four takes a different approach. Instead of including household-year fixed effects, we include household fixed effects and instead add an additional control: average expenditures in the last 12 months. The idea here is again to try and capture changes in permanent income. However, results remain consistent, with the coefficient barely changing and remaining significant.

We present additional robustness checks in Table A2. First, we include 12 lags of expenditures instead of average expenditures over the previous 12 months. This increases the flexibility of the specification. Second, we instead include 12 leads. Since changes in permanent income are in the future, controlling for past expenditures may not be sufficient. Finally, we also include both 12 leads and 12 lags. All results are completely consistent with the results in Table 2; in fact, all three coefficients are either 0.033 or 0.034.

Finally, column five includes village-by-year-by-month fixed effects. This column is not meant to test for consumption smoothing, per se. Instead, given the results in the first four columns, column five instead asks whether the failure of consumption smoothing is driven by the failure of villages to self-insure following covariate shocks as in Townsend (1994) or, more generally, whether failures of consumption smoothing are largely concurrent with other households in a village. However, we reject this explanation, as the coefficient is unchanged.

A key question is whether households with higher social status (a proxy for greater assets and connections) are more able to smooth consumption than others. We test for this possibility in Table 3, which repeats four of the specifications in Table 2, but separately

Table 3: Tests for consumption smoothing, marginal effects by education of head

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4
Current income - less than primary	0.059***	0.044***	0.040***	0.042***
	(0.007)	(0.006)	(0.006)	(0.006)
Current income - primary graduate	0.060***	0.038***	0.037***	0.037***
	(0.010)	(0.009)	(0.008)	(0.008)
Current income - secondary graduate	0.050***	0.022***	0.021***	0.021***
	(0.012)	(0.005)	(0.006)	(0.006)
Tests of equality (p):				
Less than primary = Primary	0.893	0.590	0.707	0.587
Less than primary = Secondary	0.561	0.009	0.017	0.008
Primary = Secondary	0.538	0.114	0.115	0.113
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. Coefficients are marginal effects, not interaction terms. All standard errors are clustered at the household level.
* p<0.1 ** p<0.05 *** p<0.01

across households on a key variable: education of the household head. Specifically, we present marginal effects of income for three different levels of education: less than primary, primary graduate, and secondary graduate (or higher). Column one again presents coefficients using identification across households. Here, we see no significant differences in correlations between current monthly income and current monthly expenditures.

Column two includes household fixed effects and the first differences start to emerge. The results are consistent with the least-educated households being least able to smooth consumption. Specifically, the coefficient for current income for those with less than primary education is twice as large for those who are secondary graduates, and this different is significant ($p<0.01$). Households with heads who ended their education after graduating from primary school are also less able to smooth consumption, with a coefficient more than 70% larger than those who graduated from secondary school, and this difference is marginally significant ($p=0.11$).

Results in columns three and four are consistent with the results in column two. In particular, across both columns, we strongly reject equality of the coefficient for those with less than a primary education and those who graduated secondary, while we just barely do not reject equality for the two higher groups. Moreover, as in Table 2, the coefficients are remarkably stable across all household fixed effects specifications. Finally, in Table A3, we present results based on initial household wealth. Higher wealth is significantly correlated with a better ability to smooth consumption. Overall, these results show (1) households in our sample do not perfectly smooth consumption and (2) disadvantaged households struggle more in this regard than others.

4.2 Measuring months in poverty

Having shown that households do not perfectly smooth consumption, we next show how moving from yearly estimates of poverty to monthly estimates affects conclusions. In particular, the results here show how a yearly focus misses a substantial amount of poverty.

Table 4 presents weighted poverty summary statistics across three columns, based on household expenditure. The first column is the simple average for the entire sample. The second column presents means for households who are defined as poor for the entire year – in other words, using conventional poverty measures – while the third column presents means for households who are not poor for the year.

Poor households comprise approximately 29.2% of person-months in the sample. As such, poor households are the minority of households. Compare this to the mean monthly poverty rate of 36.8%. Focusing on months-in-poverty instead of yearly poverty, the head-count poverty rate increases by more than a quarter. The increase for the distributional sensitive measures, the Watts (1968) index, and the squared poverty gap of Foster et al. (1984) are even larger, at 40% and 48%, respectively.

In addition, the majority of households experiences poverty at least once per year; 62.7% of all households experience at least one month of poverty while 47.3% of *non-poor* households experience at least one month of poverty. When looking at poverty spells – defined as at least two months in a row of poverty – we come to similar conclusions: more than a quarter of non-poor households experience at least one poverty spell in any given year. In other words, focusing on yearly poverty misses a substantial amount of person-months of poverty.

Table 4: Poverty summary statistics

	(1) Everyone	(2) Poor for the year	(3) Not poor for the year
Weighted proportion		0.292	0.708
Mean yearly poverty	0.292	1.000	0.000
Mean monthly poverty	0.368	0.863	0.164
Mean yearly watts	0.089	0.303	0.0
Mean monthly watts	0.125	0.361	0.028
Mean yearly squared poverty gap	0.025	0.087	0.0
Mean monthly squared poverty gap	0.037	0.113	0.006
Poor at least once in year	0.627	1.000	0.473
At least one poverty spell in year	0.514	0.998	0.267
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. All statistics are weighted.

Since more than 70% of people live in households defined as non-poor for the year, poverty experiences for these households add up to a substantial proportion of total poor months across all individuals. Figure 1 breaks down the total number of people in poverty in each month across the sample using expenditures. Towards the beginning of the sample period, (yearly) non-poor households sometimes contribute up to 40% of total person-months of poverty. This proportion is decreasing over the sample period, but is still between 20 and 30% of all poverty by the end of the sample. In other words, focusing on yearly poverty may miss between 20 and 40% of monthly poverty.

Figure 1: Total number of people in poverty,
by month and poverty status

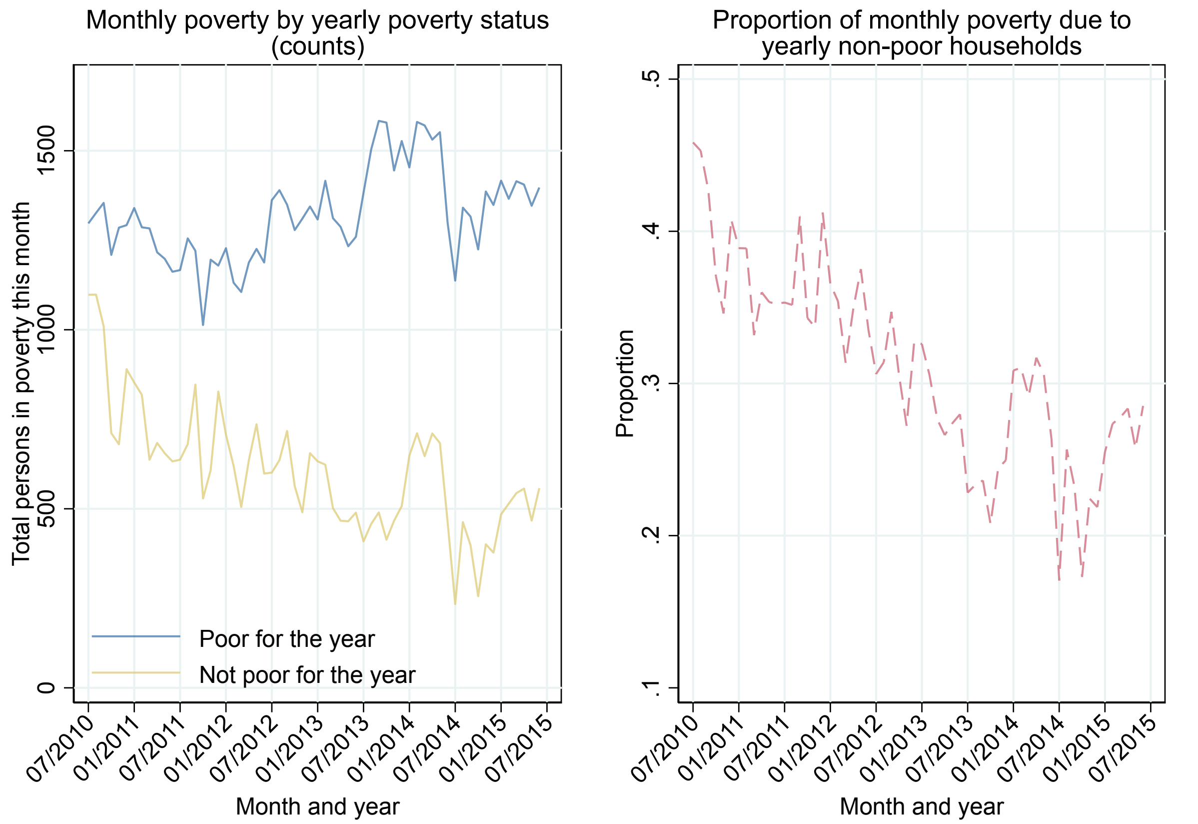

Notes: The left figure disaggregates households into those who are poor for the entire year, using average monthly expenditures across the 12 months and those who are not poor for the entire year. All counts and proportions are weighted.

Up to this point, we have presented statistics using only expenditures. An alternative is to instead use income. However, income is quite variable in our context. Most households are involved in the agricultural sector in some way, meaning that their income is seasonal. Since income is constructed as revenues net of costs, it can also be negative, leading to large swings, even for relatively well off households.

Figure 2 makes this clear. To construct these figures, we take each household's average monetary measure (income for the left panel and expenditures for the right panel) across the entire sample and divide by the poverty line. This ratio is on the x-axis, with a value of one indicating that the household is right at the poverty line; a value of 3 indicates that the household's annual resources are 300% of the poverty line. The y-axis is the proportion of months that a household is in poverty. The red curves are smoothed estimates of the average share of months in poverty for the sample.

If there was no income instability, households would be either poor all year or not poor all year. All the dots would be lined up at 1 (=12 months in poverty) if $y/z < 1$ (the poor part of the sample) and all dots would line up at 0 (no months in poverty) if $y/z > 1$ (the not poor part of the sample). The panels present a very different pattern.

Figure 2: Months in poverty and annual income

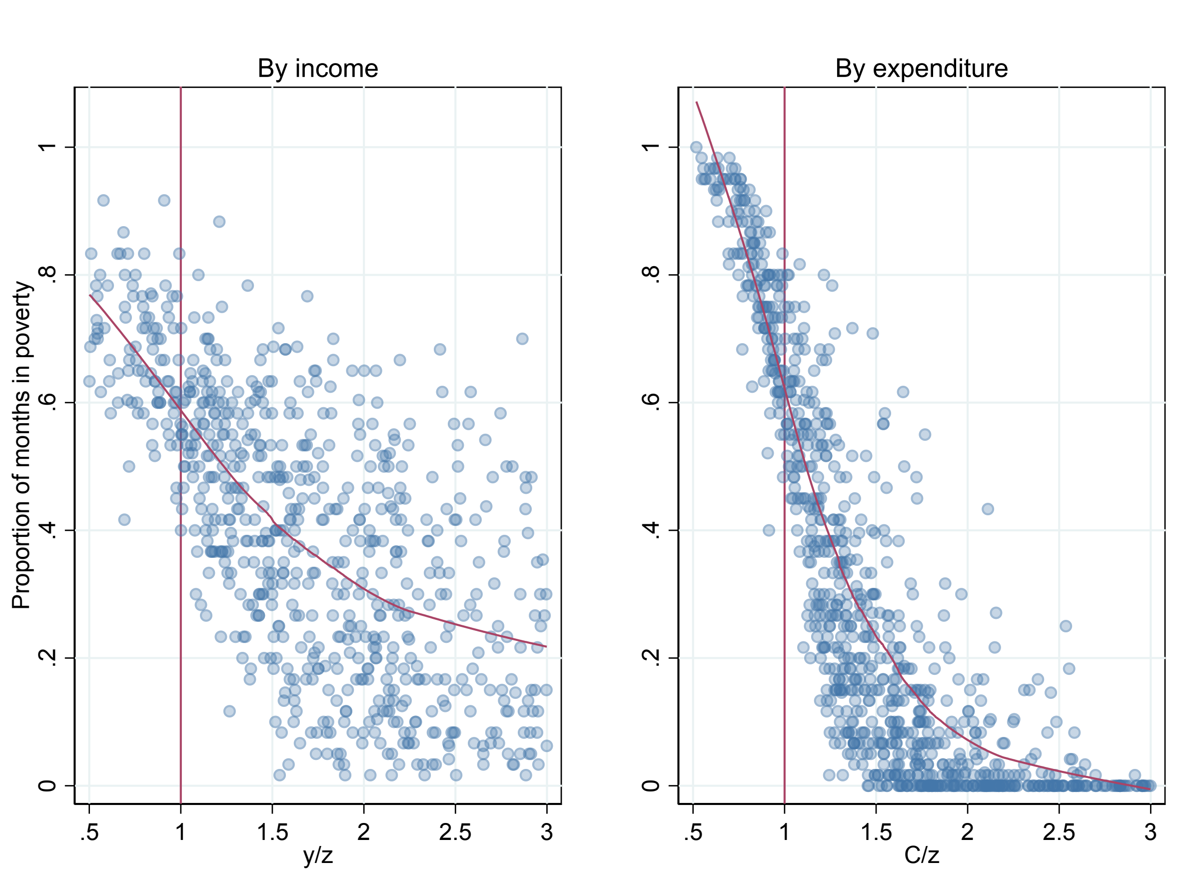

Notes: In both figures, the x-axis is the ratio of the monetary measure (income for the left figure and expenditures for the right figure) to the poverty line, averaged across the entire 60 months of the sample. The y-axis is the proportion of all months, across the entire sample, that a given household is in poverty. For ease of presentation, households below 0.5 are dropped from the figure. For income, these households make up around four percent of all household-month observations. For expenditures, not a single household as an average below 0.5. The red line is a local polynomial regression of y on x.

There are several important features in these figures. First, out of 692 households in the figure, not a single household has zero months of income poverty. Note that the figures are restricted to households with yearly expenditure or income below 300% of the poverty line, and even for households above a ratio of 300% – not shown on the graph – just five households – out of 179 – do not have a single month of income poverty. Moreover, household relatively far from the poverty line – above a ratio of two, for example – still experience a substantial amount of income poverty, upwards of 50 or 60%.

Second, the difference when compared to expenditures is striking. There are a substantial number of households who are never expenditure poor. Similarly, the distribution of poverty for a given ratio is much smaller for expenditures than for income. All the same, many households experience months of poverty even when measured by expenditure. These graphs are consistent with the evidence that households smooth consumption, but imperfectly.

Figure 3: Poverty entrance and exit

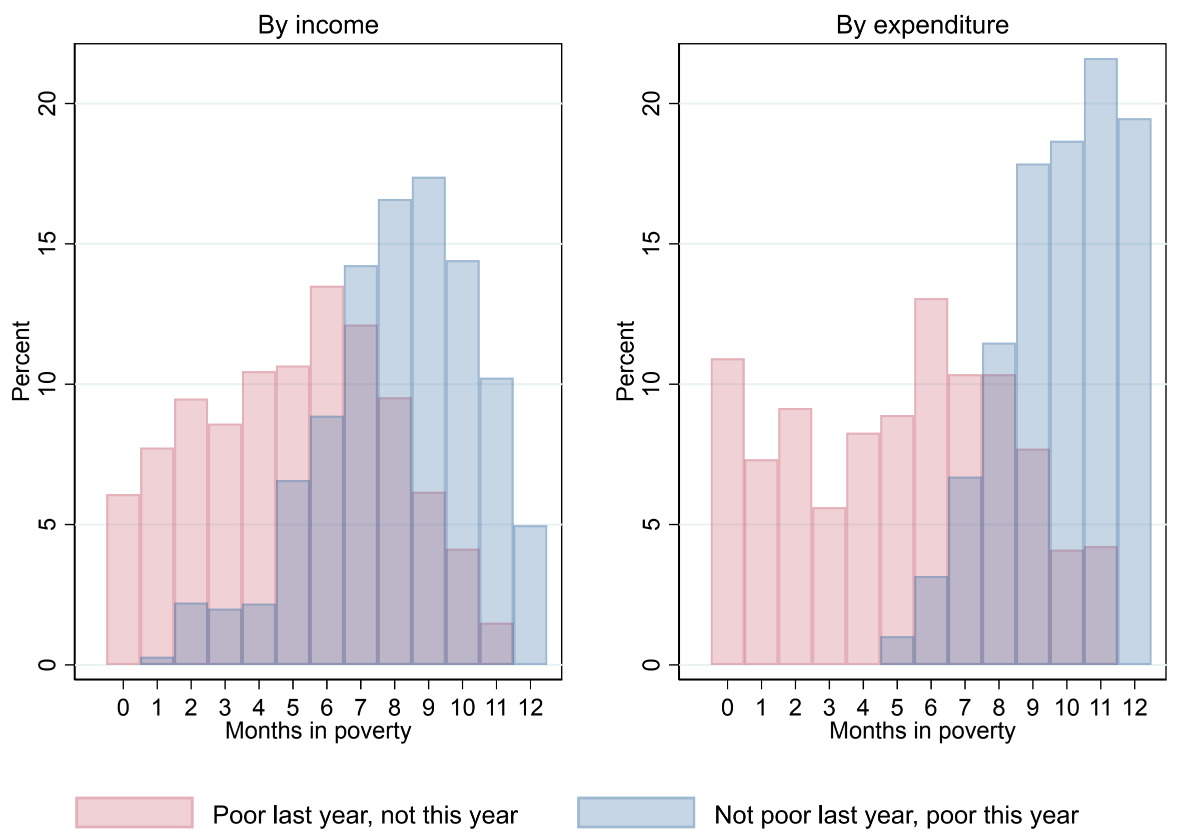

Notes: In both figures, the x-axis is the number of months a household is in poverty in a given year. The sample is split using the previous year's overall poverty classification and this year's poverty classification, with income for the left figure and expenditures for the right figure. All counts are weighted.

Given the amount of poverty experienced by non-poor households, an important question is: what does it mean to exit or enter poverty? We do not have a comprehensive answer to this question, but we can shed some light on poverty dynamics in the data. Figure 3 presents what we might traditionally define as "exit" and "entrance" with respect to poverty. Specifically, we split households into those who were poor last year but are not poor this year ("exit") and those who were not poor last year but are poor this year ("entrance").

The conventional view of poverty would suggest that these are completely different states, but the two panels make clear that the terms are not as well defined as they might seem. A simple expectation would be that people who are poor last year but not this year should experiences zero months of poverty now. Similarly, people not poor last year but poor this year should experience 12 months of poverty. The panels should thus have a single red spike at zero and a single blue spike at 12.

The panels show something very different, with the mode for the red bars at six months in both panels. For households who exit poverty, a substantial proportion continue to experience poverty, regardless of whether we use income or expenditures. With income, almost 95% of all individuals experience at least one month of poverty, while the number with expenditure is almost 90%. In fact, almost half of all individuals experience at least six months of poverty despite having seemingly “exited” poverty, using either income or expenditures.

The story is clearer for households that “enter” poverty, however, especially for expenditures. Almost 60% of individuals who enter expenditure poverty are poor for at least nine months, and not a single person is poor for less than five months. This is especially stark when compared with those who exited expenditure poverty.

Figure 4: Monthly poverty, 12-month poverty, and expenditures,
Monetary comparison

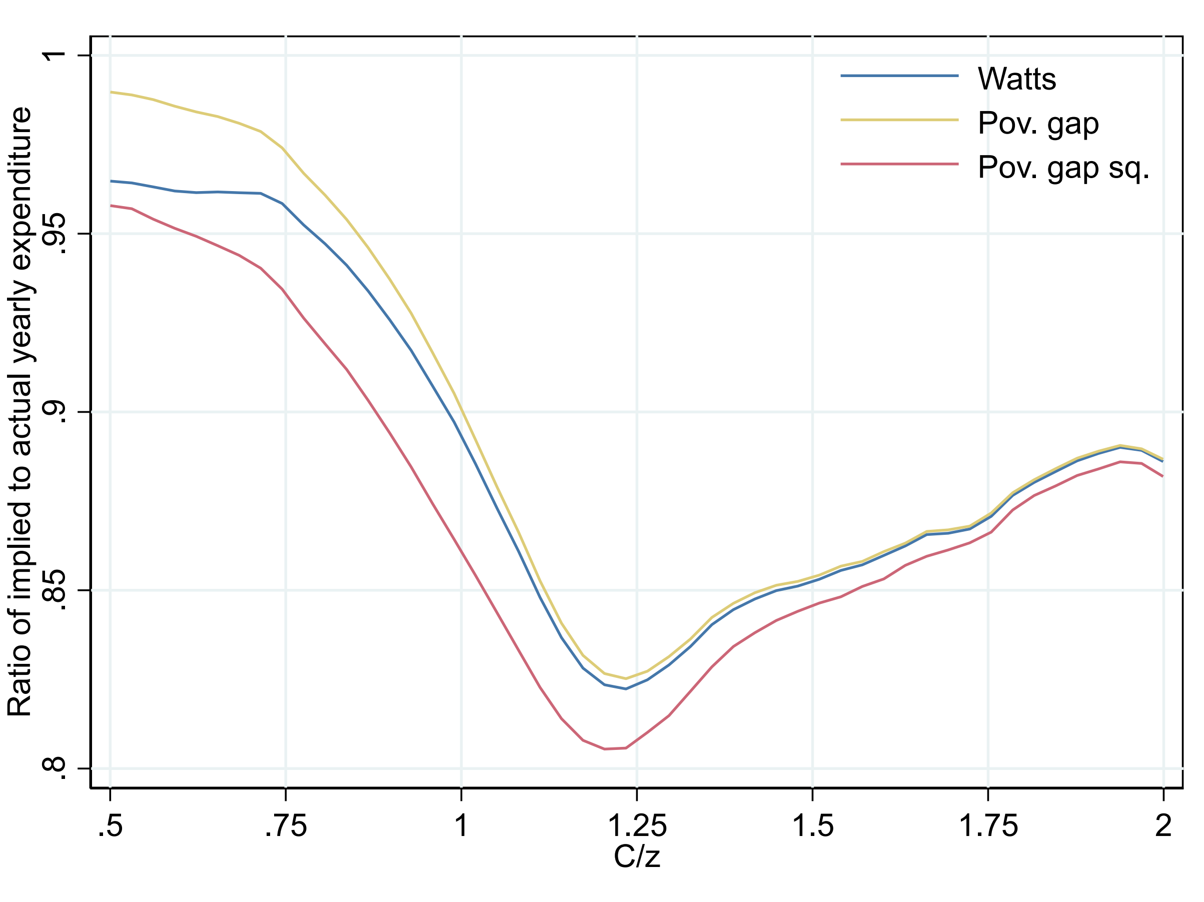

Notes: We calculate the implied yearly expenditure by inverting the poverty measure to find the yearly expenditure that yields the same value on the poverty measure as using the monthly value.

The figures show that allowing month-to-month variation in poverty broadens understandings of headcount poverty. How does the change of temporal focus affect distributionally-sensitive poverty measures? Given the sensitivity to variation in resources below the poverty line of the squared poverty gap and Watts (1968) index, we expect even larger differences when moving from the year to the month.

Figure 4 presents one way to visualize this change. We calculate the average monthly measures – $P(c_{it})$ – which are higher than the corresponding yearly value – $P(\bar{c}_i)$. To calculate the change in welfare resulting from this change of focus, we take $P(c_{it})$ and invert the yearly measure to find the equivalent yearly income that would result in the average monthly poverty measure. We then compare this implied yearly income to actual income and plot the resulting values in Figure 4, with this ratio on the y axis and the consumption over poverty line ratio on the x axis.

From half the poverty line to twice the poverty line, this implied yearly income is always lower than actual income, as seen by a ratio less than one. This difference is largest just above the poverty line – around approximately 1.2 – with a ratio of implied to actual income of between 0.8 and 0.85. The difference is largest for the squared poverty gap and smallest for the poverty gap (which is not distributionally sensitive), with the Watts (1968) index falling in between.

Results in the previous section suggested that education is an important predictor of the ability to smooth consumption. We now turn to a breakdown of monthly and yearly poverty rates by education of the household head.

Table 5 presents simple means of three types of headcount poverty: monthly expenditure poverty, monthly income poverty, and annual expenditure poverty. This yearly poverty measure is defined based on each survey year, motivated by conventional yearly statistics that are defined similarly. We present means for each of these three poverty measures across three groups based on education.

Table 5: Average poverty measures by length and education of head

	(1) Less than primary	(2) Primary	(3) Secondary
P(month, expenditures)	0.416	0.360	0.251
	(0.018)	(0.026)	(0.024)
P(month, income)	0.384	0.422	0.392
	(0.013)	(0.021)	(0.020)
P(annual, expenditures)	0.334	0.307	0.204
	(0.020)	(0.029)	(0.027)
Tests of equality (p):			
P(m,c)=P(m,y)	0.131	0.053	0.000
P(m,c)=P(a,c)	0.000	0.000	0.000
Observations	22,560	11,201	11,152

Notes: Coefficients are simple means. Standard errors are in parentheses. Standard errors are clustered at the household level.

The linkage of poverty and household finance makes three clear predictions regarding consumption smoothing and poverty measures. First, if households smooth perfectly, the measures of expenditure-based months-in-poverty and yearly poverty should be identical. Second, if households do not smooth at all, then expenditure-based months-in-poverty should equal income-based months-in-poverty. Third, if households smooth but do so imperfectly, then expenditure-based months-in-poverty will fall somewhere between income-based months-in-poverty and expenditure-based yearly poverty.

There are several striking patterns in Table 5. First, row two shows that all three groups have similar levels of months-in-poverty as measured by variation in household income. Even for households that are presumably more well off (with higher levels of education), income is still quite variable.

Second, the first row shows that months-in-poverty as measured by expenditure decrease markedly from the first column to the third column. The difference for the group with the most education is that they have managed to reduce their exposure to poverty far more than their less-educated neighbors.

Third, households with heads that have less than a primary education appear particularly exposed to poverty. We are unable to reject the equality of months-in-poverty as measured by expenditure and as measured by income. If anything, expenditure-based months-in-poverty is higher than income-based months-in-poverty (p=0.13). In other words, we are unable to reject that these households do not smooth their consumption at all. (Note that expenditure-based months-in-poverty is also higher than expenditure-based yearly poverty, as expected given the similar expenditure and income means.)

Fourth, if one only had income data, how well would it approximate months-in-poverty as measured by expenditure? For those with less than a primary education, months-in-poverty as measured by income is a close predictor of months-in-poverty as measured by expenditure. But that is not true for better educated household heads, who smooth consumption to a much greater degree. For them, conventional measures that use yearly income provide a closer prediction. Households whose heads have stopped after primary education fall somewhere between these two extremes. Specifically, while we cannot reject that monthly expenditure and income poverty are the same for column one, monthly expenditure poverty is approximately half way between annual expenditure and monthly income poverty for those in column two and monthly expenditure poverty is only around one quarter of the way between annual expenditure and monthly income poverty for those in column three.

4.3 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 6: 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 6 shows that, if anything, 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.

4.4 Expenditures on durables

Durables pose complications when measuring poverty. Specifically, they pose a challenge when we have expenditure data instead of consumption data, which is nearly always the case. Consider a household that purchases a bicycle, for example. The purchase of the bicycle shows up in one month of expenditures and leads to a large "spike" in spending. However, the actual consumption of that bicycle does not take place entirely in that month.

This is actually one rationale for measuring poverty at the yearly level instead of month by month. Our proposal maintains the yearly lens, which helps decrease concerns that variation might be driven partly by these large spikes in expenditures, raising the mean and leading to higher variance. As long as expenditures on durables are consistent across years, then this spike will show up year after year, and any changes in our monthly poverty

Table 7: Percent of expenditures on durable

	(1) Durables	(2) Semi-durables
Top expenditure quartile	2.3%	4.7%
Q3	1.6%	5.2%
Q2	0.9%	4.8%
Bottom expenditure quartile	0.4%	3.5%

Notes: The percentages indicate the percent of average monthly expenditures spent on each type of good. Quartiles are defined using yearly expenditures.

measure will be driven by other behavioral changes. Moreover, if increasing incomes leads to larger spikes like this, overall variation in expenditures might actually be increasing in expenditures, rather than decreasing.

Nonetheless, we believe it is helpful to contextualize how much durables are contributing to our results. Our sample is a group of relatively poor households. As such, it turns out that expenditures on large durables and semi-durables are quite low. Table 7 breaks out the percentage of average monthly expenditures spent on durables (column one) and semi-durables (column two). In addition, it breaks out spending into expenditure quartiles. Recall that approximately 30 percent of our sample is considered poor. As such, the poverty line is at the very low end of the second quartile. Table 7 shows that expenditures on durables and semi-durables in these two quartiles are very low, especially in the bottom quartile. For this quartile, less than one-half of one percent of all expenditures go to durables, while just 3.5 percent of all expenditures go to semi-durables. Overall, we believe these averages are low enough that they are not responsible for the entirety of the result. Nonetheless, the proper treatment of durables could be important in other contexts and, as such, their treatment remains an important measurement issue when discussing poverty.

4.5 Policy experiment

Since households do not smooth consumption perfectly, there may be welfare gains to be had from improving their ability to smooth consumption, even if their average consumption does not increase. For example, McCulloch and Baulch (2000) write: "Anti-poverty programmes often seek to improve their impact by targeting households for assistance according to welfare measures in a single time period. However, a growing literature shows the importance to poor households of fluctuations in their welfare from month to month and year to year." Similarly, Ruggles and Williams (1989) estimate with monthly data from the United States in the 1980s that over than one-third of all poverty spells could have been eliminated if households' financial assets were targeted to alleviating poverty in the most difficult periods.

In this section, we present a final set of results to quantify some of these possibilities. We consider alternative hypothetical transfer programs and analyze how the design affects poverty measures.

More details can be found in the methodology section, but the basic outline is as follows. In each hypothetical intervention, we transfer 960 rupees per capita per year to all households below the poverty line in a given year (as measured by yearly expenditure). However, we vary how we make the transfers: either monthly (80 Rs per month), over six months (160 Rs per month), or over three months (320 Rs per month). In the latter two cases, we make this transfer in the poorest months that the households experience in a given year. We assume the household consumes the entirety of the transfer in the month they receive it.[15]

This relates to Equation 5, which showed that any effects of changes in the policy environment can be broken down into effects on mean consumption and effects on variation around that mean. Since all three designs transfer the same amount of money, any differ-

[15]Given that we were unable to reject no smoothing in Table 5, we do not think this assumption is too extreme, even if it is not perfectly accurate.

Table 8: Policy Experiment Results

	(1) Headcount	(2) Watts
No transfer	0.863	0.361
12 months (80 Rs)	0.743	0.219
Six months (160 Rs)	0.751	0.208
Three months (320 Rs)	0.699	0.234
Households	391	391
Month observations	12,300	12,300

Notes: The same total amount of 960 rupees per person per year is transferred in all designs (except "no transfer"). That money is transferred in 12 equal payments for the "all months" design, in six equal payments for the "six months" design, and in three equal payments for the "three months" design. Transfers are always made in the poorest months.

ences are attributable purely to the latter effect.

Table 8 shows the overall results for both the headcount (column one) and Watts index (column two). The transfer of course has a large impact on overall poverty, relative to the no-transfer baseline. However, despite transferring the same amount of money, the three separate designs have sometimes very different effects. For example, with the headcount, the monthly transfer of 80 Rs decreases poverty by 13.9 percent (12 p.p.). Transferring 320 Rs every four months (or three months per year) decreases poverty by 19.1 percent. This is a 37 percent larger decrease in the poverty rate, despite transferring the same amount of money.

The results for the Watts index are not as stark given how it weights those farther from the poverty line. The monthly transfer decreases poverty by 39.4 percent, while the transfer across six months decreases poverty by 42.4 percent, or 7.6 percent more than the monthly transfer. Interestingly, the transfer across three months performs worse than the other two options here, despite performing best with headcount poverty. This underlines the importance of the choice of poverty measures when evaluating government programs.

We present the results graphically for headcount poverty in Figure 5, with expenditure-based months-in-poverty on the y axis and the (yearly) consumption to poverty line ratio on the x axis. Across nearly all of the range, a transfer focused on the poorest three months performs better at reducing monthly headcount poverty than the other two designs. However, transferring across six months actually performs best for those closest to the poverty line.

Figure 5: Policy experiment - Headcount poverty

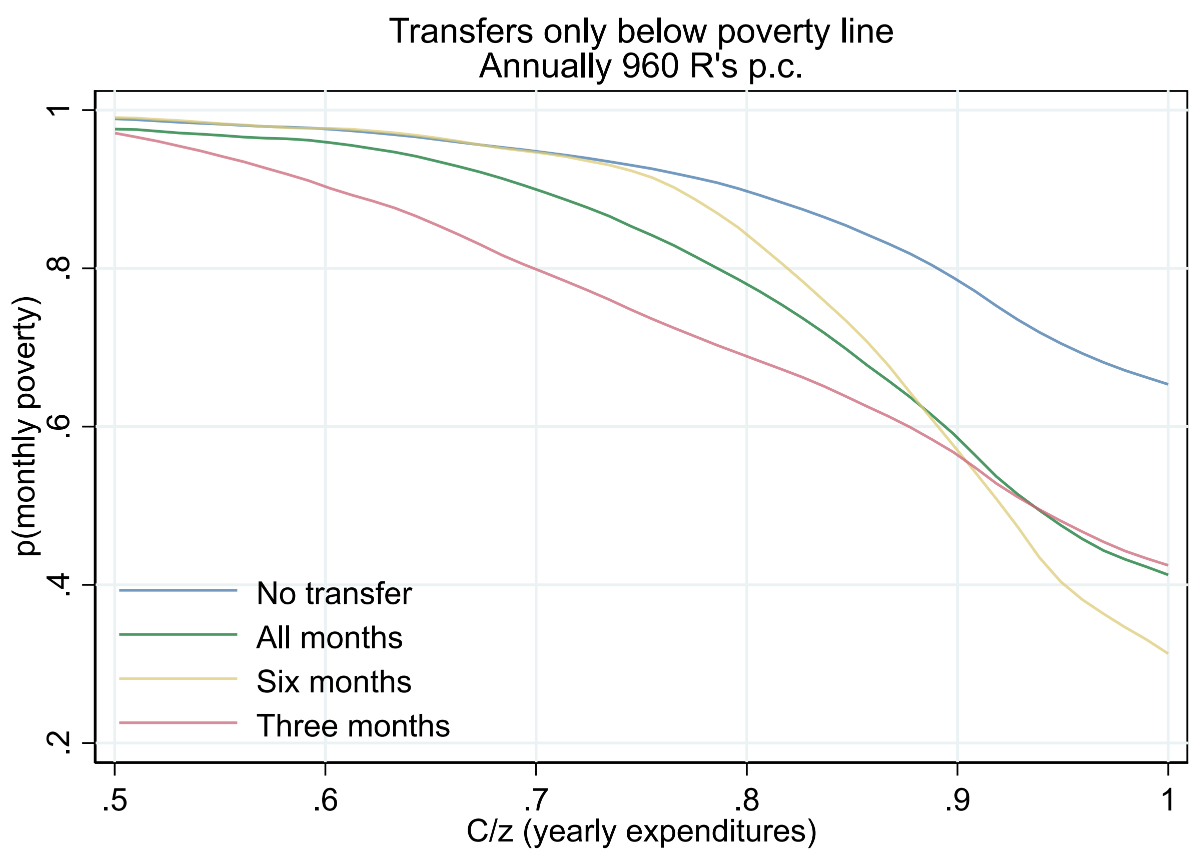

Notes: The same total amount of 960 rupees per person per year is transferred in all designs (except "no transfer"). That money is transferred in 12 equal payments for the "all months" design, in six equal payments for the "six months" design, and in three equal payments for the "three months" design. Transfers are always made in the poorest months.

Figure 6: Policy experiment - Watts (1968) index

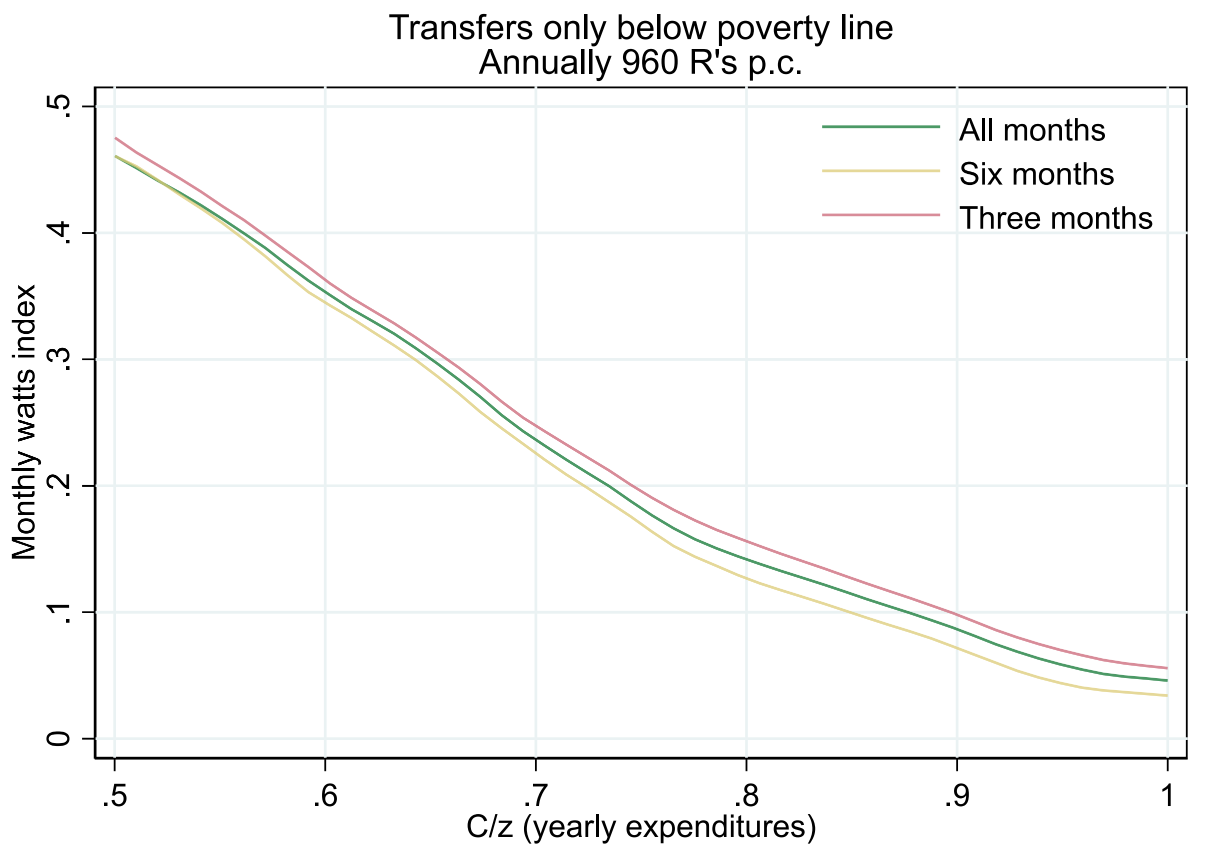

Notes: The same total amount of 960 rupees per person per year is transferred in all designs (except "no transfer"). That money is transferred in 12 equal payments for the "all months" design, in six equal payments for the "six months" design, and in three equal payments for the "three months" design. Transfers are always made in the poorest months.

Figure 6 graphically presents the results using the Watts (1968) index. Here, instead, cost-effective strategies involve targeting the most disadvantaged people in their most difficult times. Overall, all three designs perform much better here than with headcount poverty. This is driven by the fact that increasing the consumption of the poorest by around 80 Rs per month (the size of the 12-month transfer) does not have much of an effect on headcount poverty – since many households are too far from the poverty line to cross it with an 80 Rs/month transfer– whereas with the Watts (1968) index, any transfer registers as a poverty reduction. The six-month transfer performs best across almost the entire range, with aa exception for those households at the very bottom of the distribution (of whom there are relatively few).

The intuition here is that it is helpful to concentrate effort in the poorest months, but that the challenges go beyond the hardest three months. In other words, the challenges go beyond a well-defined period of seasonal poverty.

5 Conclusion

The conventional approach to measuring poverty is convenient and highlights people's overall earning capacity, but it creates a distance between poverty as measured and poverty as experienced. Evidence from household finance shows that the experience of poverty is captured by the interaction of insufficiency, instability, and illiquidity. Insufficiency reflects low overall earnings as seen in annual sums. Instability reflects the variation in those earnings and in needs within the year. Illiquidity reflects households' challenges in coping with instability, leading to spikes and dips of within-year consumption.

Conventionally measured poverty based on annual data highlight overall insufficiency but not the experience of poverty due to within-year instability and illiquidity. In a sense, our approach, focusing on months-in-poverty during the year, opens a window on variance around mean levels of deprivation, not just the mean alone. The aim is not to criticize yearly measures but to show their limits and how they can be addressed.

The data are from agricultural villages in South India. They are not representative of the global poor in a statistical sense, but they represent an important setting for understanding global poverty.

The most important practical limit to our approach is the need for monthly data. The approach can work as well with several waves of data collected within the year (e.g., Azevedo and Seitz 2016b), but even that data is rare. Our hope is that, once the advantages of high-frequency data are recognized, new data collection efforts will follow.

But, even with monthly data, there remain empirical challenges. One is the standard problem of measurement error. A second is the fact, well known to economists, that spending does not equal consumption. A household may buy a motorcycle, say, purchasing it at the start of the year. Consumption of the motorcycle's "mobility services," however, happens throughout the year. Households spend at a particular time, but their consumption often occurs at other moments. To the extent possible, data need to be converted into consumption equivalents which occur over time, rather than simply analysing spending numbers in the moment of purchase. Conceptually, the problem is easy to handle. Practically, it is demanding. We took efforts to convert data to consumption units here, but we often had to rely on reasoned assumptions rather than exact data on consumption. As high frequency poverty gets addressed with high frequency data, this challenge will need to be in mind from the start.

The way that poverty is conceived shapes the way that it is measured. The reverse is also true: the way that poverty is measured can shape the way that poverty is conceived. Much of the paper has demonstrated the empirical relevance of high-frequency poverty measurement in rural India. The larger motivation is to expand what it means to reduce poverty in general, building from people's actual experience of poverty. In our framework, we count experiences of poverty that are not registered as strongly, or at all, in conventional measures; for example, we capture the particular struggles with seasonal poverty. This expansion can help policymakers better target deprivations. In doing so, it raises questions outside the scope of this study: Specifically, should the social weight placed on reducing months-in-poverty be conditioned on the broader temporal context? How should it matter, if at all, if months-in-poverty are experienced by people who would conventionally be considered not poor? Is seasonal poverty deserving of similar concern to other periods poverty? We cannot answer these questions in our framework, but see value in clarifying them.[16]

[16] A parallel question arises for conventionally-measured yearly poverty when viewed across years. Turning to the ethics of conventional poverty measures: are there compelling philosophical defenses, beyond convenience, for measuring poverty with yearly income and consumption when doing so obscures the depths of poverty?

References

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

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

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.

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

Banerjee, A., D. Karlan, and J. Zinman (2015): "Six Randomized Evaluations of Microcredit: Introduction and Further Steps," *American Economic Journal: Applied Economics*, 7, 1–21.

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

Beaman, L., D. Karlan, B. Thuysbaert, , and F. Truffa (2014): "Saving for a (not so) Rainy Day: A Randomized Evaluation of Savings Groups in Mali," Tech. Rep. Northwestern Department of Economics working paper.

Berg, C. and M. S. Emran (2020): "Microfinance and Vulnerability to Seasonal Famine in a Rural Economy: Evidence from Monga in Bangladesh," *B.E. Journal of Economic Analysis Policy*, 20, 1–36.

Biewen, M. (2014): "Poverty persistence and poverty dynamics," *IZA World of Labor*.

Bradbury, B., S. Jenkins, and J. Micklewright (2001): "Conceptual and measurement issues," in *The Dynamics of Child Poverty in Industrialised Countries*, ed. by B. Bradbury, S. Jenkins, and J. Micklewright, Cambridge: Cambridge University Press, 27–61.

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.

Calvo, C. and S. Dercon (2009): "Chronic poverty and all that: The measurement of poverty over time," in *Poverty Dynamics: Interdisciplinary Perspectives*, ed. by T. Addison, D. Hulme, and R. Kanbur, Oxford: Oxford University Press.

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.

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.

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.

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

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

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

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

Edwards, A. (2014): "Dynamics of Economic Well-Being: Poverty, 2009–2011," Current population reports, US Census Bureau, Washington, DC.

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

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.

Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2020): "Wealth, Race, and Consumption Smoothing of Typical Income Shocks," Working Paper 27552, National Bureau of Economic Research.

Hanna, R. and B. Olken (2018): "Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries," *Journal of Economic Perspectives*, 32, 201–226.

Himmelfarb, G. (1984): *The Idea of Poverty: England in the Early Industrial Age*, New York: Knopf.

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.

Jappelli, T. and L. Pistaferri (2017): *The Economics of Consumption*, New York: Oxford University Press.

Jayachandran, S. (2006): "Selling labor low: Wage responses to productivity shocks in developing countries," *Journal of Political Economy*, 114, 538–575.

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

Krishna, A. (2016): "The Dynamics of Poverty," in *The Oxford Handbook of the Social Science of Poverty*, ed. by D. Brady and L. M. Burton.

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

Maag, E., H. E. Peters, A. Hannagan, C. Lou, and J. Siwicki (2017): "Income Volatility: New Research Results with Implications for Income Tax Filing and Liabilities," Tech. rep., Tax Policy Center / Brookings Institution and Urban Institute, Washington, DC.

McCulloch, N. and B. Baulch (2000): "Simulating the impact of policy upon chronic and transitory poverty in rural Pakistan," *Journal of Development Studies*, 36, 100–130.

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

Morduch, J. and J. Siwicki (2017): "In and Out of Poverty: Poverty Spells and Income Volatility in the US Financial Diaries," *Social Service Review*, 91, 390–421.

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

Ruggles, P. and R. Williams (1989): "Longitudinal Measures of Poverty: Accounting for Income and Assets over Time," *Review of Income and Wealth*, 35, 225–243.

Schneider, D. and K. Harknett (2020): "Hard Times: Routine Schedule Unpredictability and Material Hardship among Service Sector Workers," *Social Forces*, 99, 1682–1709.

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

Skoufias, E. and A. R. Quisumbing (2005): "Consumption Insurance and Vulnerability to Poverty: A Synthesis of the Evidence from Bangladesh, Ethiopia, Mali, Mexico and Russia," *European Journal of Development Research*.

Stevens, A. H. (1999): "Climbing out of poverty and falling back in," *Journal of Human Resources*, 34, 557–588.

Storer, A., D. Schneider, and K. Harknett (2020): "What Explains Racial/Ethnic Inequality in Job Quality in the Service Sector?" *American Sociological Review*, 85, 537–572.

Townsend, R. M. (1994): "Risk and insurance in village India," *Econometrica*, 62, 539–591.

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

VALLETTA, R. (2006): "The ins and outs of poverty in advanced economies: Government policy and poverty dynamics in Canada, Germany, Great Britain, and the United States," *Review of Income and Wealth*, 52, 261–284.

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

Appendix

Table A1: Year-month counts

	(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 A1: Relative weights of different poverty measures

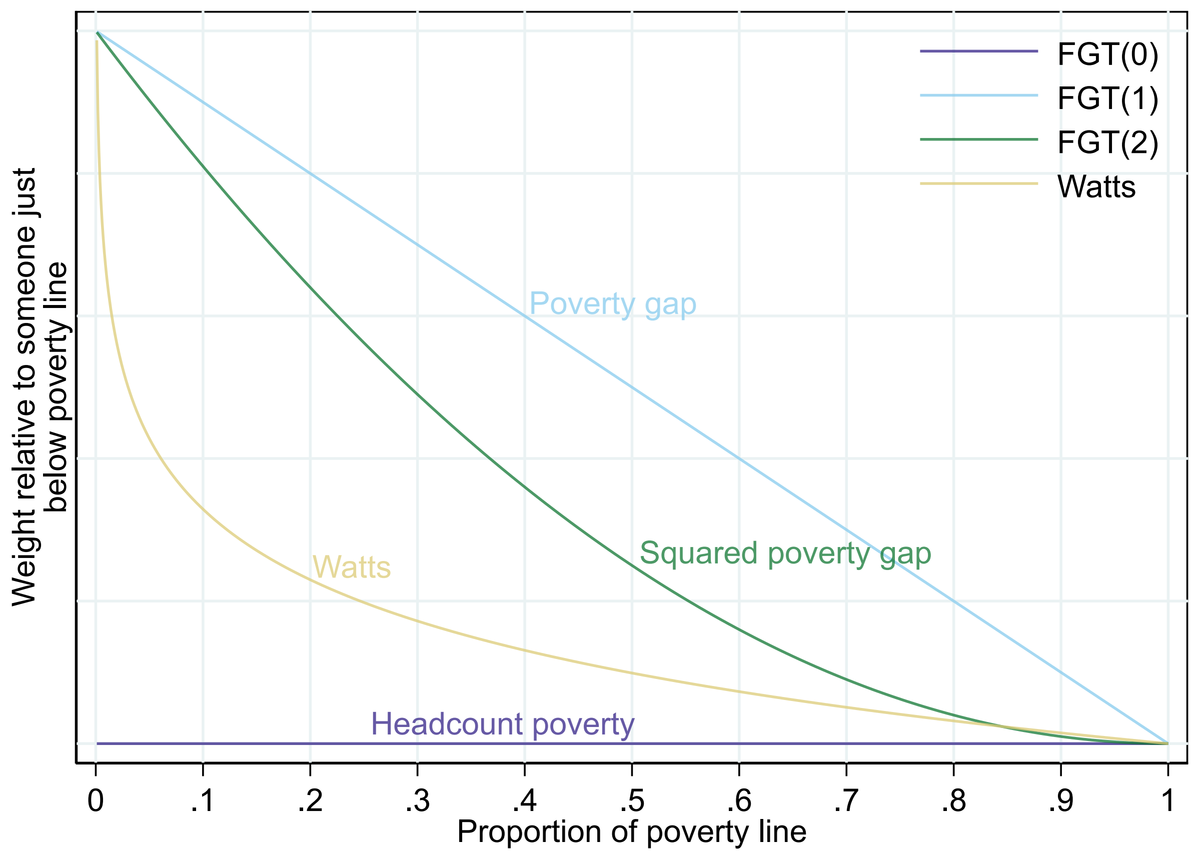

Figure A2: Median income and expenditures

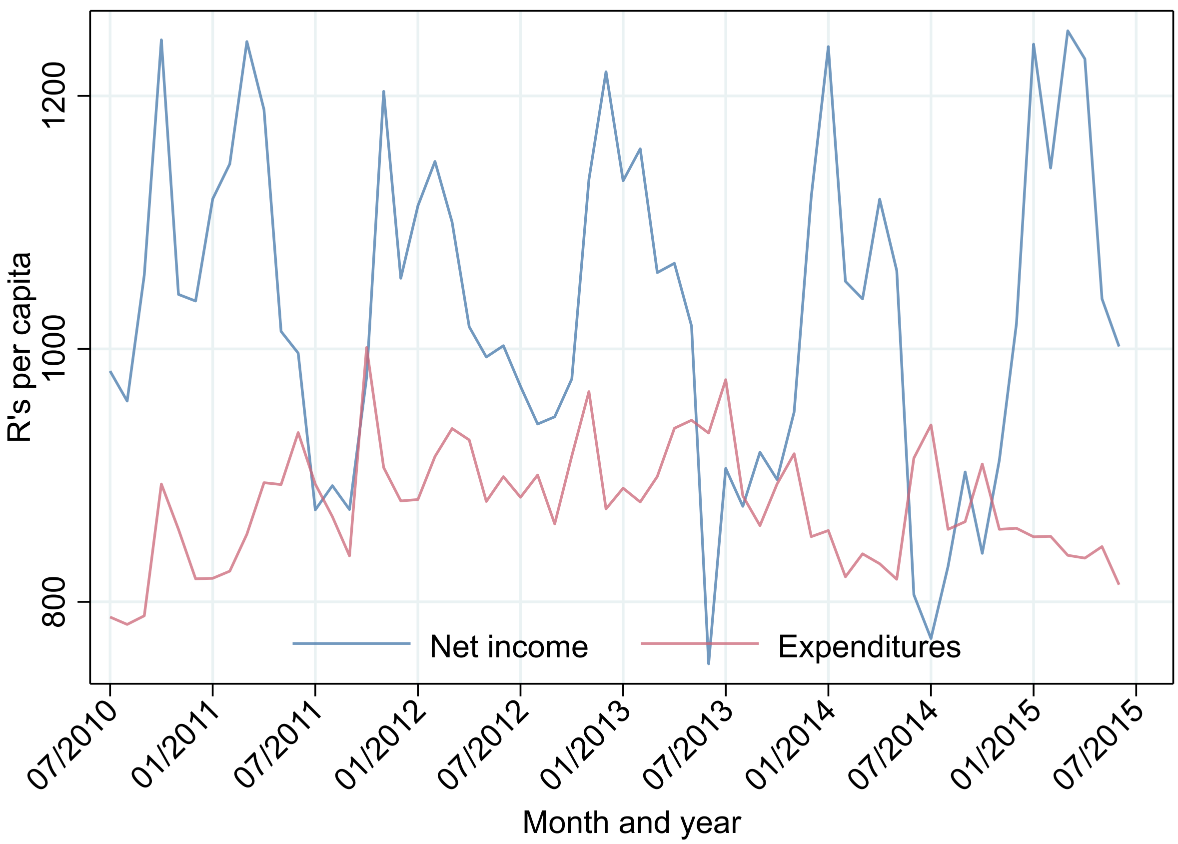

Notes: The lines are simple medians for each month, weighted by sampling weights and household size.

Table A2: Tests for consumption smoothing, flexible lags and leads

	(1) Model 1	(2) Model 2	(3) Model 3
Current income	0.033***	0.034***	0.034***
	(0.004)	(0.005)	(0.005)
Fixed effects:			
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 A3: Tests for consumption smoothing, effects by initial household wealth

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4
Current income	0.062***	0.049***	0.047***	0.049***
	(0.005)	(0.005)	(0.005)	(0.005)
Initial wealth (100,000k rupees)	305.903***			
	(41.697)			
Current income times initial wealth	−0.006**	−0.006***	−0.006***	−0.006***
	(0.002)	(0.002)	(0.002)	(0.002)
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