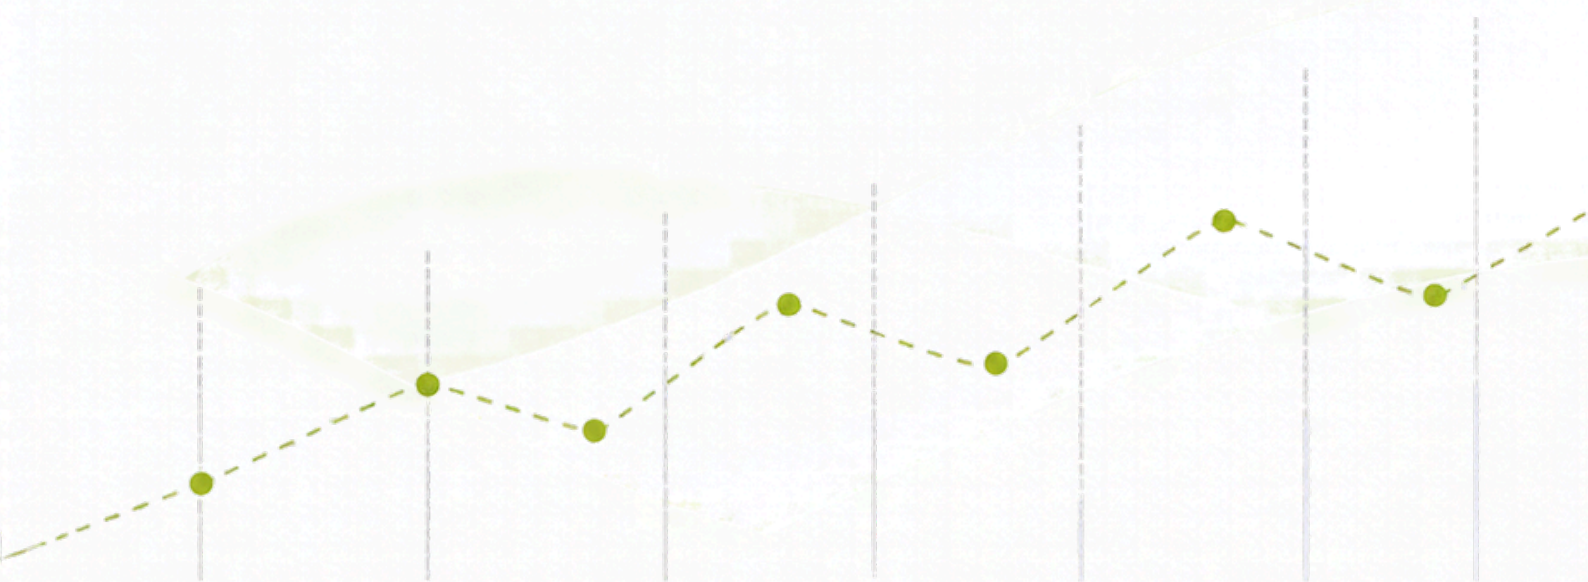
DVARA Research

Neither Always nor Never:

A Study of Poverty at Higher Frequency using CPHS data

By Abhishek Mukherjee



Neither Always nor Never:
A Study of Poverty at Higher Frequency using CPHS data

By Abhishek Mukherjee[1]

Abstract

In traditional policy discourse, poverty is framed as an annual headcount, computed as the share of households, whose welfare, measured by expenditure or income, falls below the poverty line in a given year. Most low and middle income countries, including India, measure this using consumption, while income is often used as a proxy where consumption data is unavailable. While such measures capture who is poor, they fail to provide meaningful insight into the diverse ways in which poverty is experienced. This limitation arises from the implicit assumption that households can perfectly and equally smooth their consumption year-round, an assumption that rarely holds in practice. As a result, annual measures classify households as either *always* poor or *never* poor, obscuring substantial intra-year and inter-regional variability in the poverty experience. This paper argues that a nuanced understanding of poverty demands an inquiry that goes beyond 'who' is poor and looks at 'when' they are poor. Building on the framework of 'Poverty at Higher Frequency' introduced by Merfeld and Morduch (M&M 2022, 2025), I use household-level data from the Consumer Pyramids Household Survey (CPHS) to show that: (i) poverty is more widespread than suggested by yearly headcount measures, (ii) households' ability to smooth consumption varies with socioeconomic status and directly shapes their poverty experience, and (iii) transitions in and out of poverty are not as clear-cut as implied by annual measures. I extend the framework by constructing a dynamic, region and state specific monthly poverty line. It inflates the 2011-12 poverty lines, estimated by the Planning Commission (2013) using the Tendulkar methodology (2009), with the state-level Consumer Price Index published by the Ministry of Statistics and Programme Implementation (MoSPI, 2024). The findings, drawn from recent, nationally representative data, are consistent with those of M&M, and show that the divergence between monthly and annual poverty measures is deeply heterogeneous and consequential for the design of anti-poverty interventions.

[1] The author would like to thank Misha Sharma (Lead at Dvara Research Foundation), Indradeep Ghosh (Ex-Executive Director at Dvara Research Foundation), Jonathan Morduch (Professor at NYU), and Nishanth Kumar (Postdoctoral Researcher at London School of Economics) for their generous time and helpful feedback on this paper. All remaining errors are the author's own.

1. Introduction

Low-income households (LIHs) regularly manage their finances under conditions of insufficiency, instability, and illiquidity. Insufficiency refers to persistently low incomes that are inadequate to cover basic needs. Instability refers to irregular or unpredictable cash flows. These often stem from seasonal employment, informal work, or volatile earnings. Illiquidity refers to limited access to financial tools such as formal credit, savings, and insurance that could help households smooth consumption over time. Collins et al. (2009), in *Portfolios of the Poor*, demonstrate that LIHs engage in complex financial management to cope with these conditions, using informal savings, short-term borrowing, and extensive money-management strategies to handle shocks and meet daily needs. Their findings highlight that the main challenge for LIHs is not just low income but rather the interaction of insufficiency, instability, and illiquidity, which creates an environment where uncertainty proliferates.

A directly observable and quantifiable manifestation of uncertainty in the lives of LIHs is poverty. Such poverty manifests at fluctuating levels and over varying periods of time. Traditional policy discourse, however, usually documents poverty as an annual headcount metric, computed as the share of households, whose welfare, measured by expenditure or income, falls below the poverty line in a given year (Booth, 1893; Rowntree, 1901; Himmelfarb, 1985; Gillie, 1996). Most low and middle income countries, including India, measure this using consumption, while income is often used as a proxy where consumption data is unavailable. Consequently, all national poverty rates and anti-poverty policies are constructed and implemented, respectively, based on this understanding of poverty (Ravallion 2016, Atkinson 2019, and United Nations 2025). Since annual headcount measures of poverty downplay the uncertainty inherent in the experience of poverty among LIHs, they fail to provide meaningful insight into households' disparate and shifting experiences of poverty throughout the year (Chambers, 1983; Dercon and Krishnan, 2000; Collins et al., 2009; Azevedo and Seitz, 2016b). This limitation arises from the implicit assumption that households are able to smooth their consumption equally and perfectly all year round- an assumption that rarely holds in practice, particularly among LIHs (Deaton, 1991; Fink et al., 2020; Augenblick et al., 2023). As a result, conventional interpretations of annual measures classify households as either always poor or never poor, thus obscuring substantial intra-year variability in the poverty experience.

In the paper titled *Poverty at Higher Frequency* (2025) by Joshua Merfeld and Jonathan Morduch[2] (M&M hereon), the authors show that such variability is already embedded in poverty statistics but remains analytically invisible without a deliberate high-frequency lens. They argue that if households were, in fact, perfectly smoothing their consumption, they could acquire and utilise money as and when they need it. They cite Breza et al. (2021) and Bryan et al. (2014), who provide strong evidence in the form of Randomised Controlled Trials (RCTs) to disprove the assumption of perfect consumption smoothing and bring attention to earlier literature (such as Devereux et al., 2012) that acknowledges variability in household consumption behaviour. Ultimately, M&M propose an alternative framework that emphasises the temporal dimensions of poverty. They do so by taking a monthly approach to annual poverty headcount measurement, which allows for a more nuanced understanding of the volatilities and deprivations these households face. They refer

[2] Their work on this topic spans multiple iterations of the same paper, each of which offers a somewhat different framing of the same underlying idea. An earlier version of the paper (Merfeld and Morduch, 2022) introduces the high-frequency framework using the terminology of high and low-frequency poverty, while the most recent version (Merfeld and Morduch, 2025) formalises this as the timecount. This paper does not take up the task of comparing or compiling these iterations. Rather, it draws on M&M's framework as a methodological foundation to reproduce their core empirical results using more recent data. This discussion continues in Section 3.1.

to this type of measurement as 'high-frequency' poverty- a concept that corresponds to what Merfeld and Morduch (2025) later formalised as the time count (defined in Section 3.1).

The paper applies the high frequency poverty framework developed by M&M to India and makes two contributions. The first is methodological. M&M measure poverty against a static line, held constant across all households and the entire study period. I construct a dynamic poverty line that varies by state, rural or urban region, and time. It inflates the 2011-12 Tendulkar-methodology poverty lines with MoSPI's state-level Consumer Price Index, so that the threshold reflects differences in the cost of living across geographies and changes in prices over time. The second contribution is new data. M&M draw on the Village Dynamics in South Asia (VDSA) survey, which covers 945 rural households between 2010 and 2015. I use the Consumer Pyramids Household Survey (CPHS), a nationally representative monthly panel of rural and urban households, covering January 2018 to December 2023. This allows the framework to be tested on a larger and more diverse population, and over a period that includes the COVID-19 shock. Based on these modifications, and consistent with M&M, I find that poverty is more widespread than annual measures suggest, that households' ability to smooth consumption varies with socioeconomic status, and that transitions into and out of poverty are not as clear-cut as annual measures imply. Furthermore, I show that the monthly poverty headcount is twice the annual one, that households remain exposed to months of poverty until their average expenditure reaches roughly twice the poverty line and that the least-educated households, unlike others, pass income shocks almost directly to consumption. Lastly, similar to M&M, using a cash transfer experiment, I show that the most effective timing of a cash transfer depends on how far below the poverty line a household sits.

This paper is organised as follows. Section 2 describes the data used for the analysis. Details of the methodology adapted from the M&M framework, along with the modifications, are presented in Section 3 of the paper. Section 4 presents the results of the analysis. It lays out the exploratory and analytical tools used to build evidence in support of M&M's study. Then, in Section 5, I present a discussion of the results and compare them with those from the original study. Section 6 concludes the paper by laying out the various considerations and implications of using a high-frequency approach to measuring poverty and invites the reader to further build on the modifications made to the framework.

2. Data

This analysis draws on three primary datasets, each serving a distinct role in the construction of poverty measures for the analysis period (January 2018- December 2023):

2.1. CPHS (January 2018- December 2023)

CPHS, conducted by the Centre for Monitoring Indian Economy (CMIE), is a large-scale, nationally representative longitudinal survey designed to provide high-frequency, micro-level data on socio-economic indicators in India. Initiated in 2014, it fills the gap created by the infrequent release of official datasets such as the National Sample Survey (NSS) and the Census. CPHS is the only household survey in India that provides continuous monthly observations at the household level. It forms the primary dataset for this study, replacing the Village Dynamics in South Asia (VDSA) dataset used in M&M's analysis. This exercise uses the following variables from CPHS:

(i) Demographics: Household ID, month, state, district, region (rural/urban), education group, household size group.

(ii) Survey design variables: Homogeneous Region (HR), Stratum, Primary Sampling Unit (PSU) ID, household-level weights for the monthly series (country and state-level), and non-response adjustment factors (country- and state-level).
(iii) Socio-economic indicators: Nominal total household income and nominal total household expenditure.

The household is treated as the primary unit of analysis. Its poverty status is determined by comparing each household's expenditure against the poverty line.

One important contextual consideration is that the study period spans January 2018 to December 2023, which includes the COVID-19 pandemic and the associated nationwide lockdowns in 2020 and 2021. This represents a significant structural break in household income and expenditure patterns, which is unlikely to reflect the kind of 'routine' uncertainty that the high-frequency framework is designed to capture.

2.2. RBI Handbook of Statistics (January- December 2012)

The RBI has published the region and state-specific poverty lines estimated by the Tendulkar Committee (2009) for 2011-12, based on Mixed-Reference Period (MRP) consumption[3]. These figures provide an estimate of the average per capita monthly consumption expenditure required to meet basic needs in rural and urban regions of each state. For this exercise, the 2011-12 poverty lines (expressed in INR) are treated as the baseline for each region-state.

2.3. MoSPI General Consumer Price Index (CPI) (January 2013- December 2023)

MoSPI's General CPI series measures changes in the price of a representative basket of goods and services at the state- region level[4], with 2011-12 as the base year. This exercise leverages the monthly region-state CPI values to adjust the 2012 baseline poverty lines for regionally asymmetric inflation, producing region-state-specific poverty lines for each month from January 2018 to December 2023.

Taken together, these datasets facilitate the construction of a consistent poverty measurement framework over the entire study period: the Tendulkar poverty line (RBI, 2012) provides the baseline, MoSPI CPI inflates it to the relevant month and region-state, and CPHS supplies the household-level demographic and expenditure data needed to conduct a low-frequency and high-frequency assessment of poverty.

[3] Until the 50th round (1993-94), consumption data collected by the National Sample Survey Organisation (NSSO) was based on the Uniform Reference Period (URP), under which all items were recorded on a uniform 30-day recall. Alternative schedules were tested in rounds 51-54 (held between 1994-98), and from the 55th round (1999-2000) the NSSO began collecting consumption on a Mixed Reference Period (MRP) basis alongside the URP. Under the MRP, six low-frequency item groups such as clothing and bedding, footwear, education, institutional medical care, and durable goods, are recorded over a 365-day recall, while all remaining items are recorded over a 30-day recall. The latter comprises food, pan, tobacco and intoxicants, fuel and light, and miscellaneous goods and services, including non-institutional medical care, conveyance, rent, and taxes. Official poverty estimation shifted to MRP-based consumption following the Tendulkar Committee (2009), whose 2011-12 poverty lines form the baseline for this study. Notably, a Modified Mixed Reference Period (MMRP), which shifts several food groups to a 7-day recall, underlies the Rangarajan Committee's estimates, and which came in 2014, i.e., after the Tendulkar Committee's recommendations. Nonetheless, the MRP-based Tendulkar lines are used here for consistency.
[4] Here, region refers to rural and urban areas of the state.

3. Methodology

3.1. Clarification regarding the use of 'high-frequency' and 'low-frequency'

M&M's work on 'Poverty at Higher Frequency' spans multiple iterations of the same paper, each of which offers a somewhat different framing of the same underlying idea. An earlier version of the paper (Merfeld and Morduch, 2022) introduces the high-frequency framework using the terminology of high and low-frequency poverty, while the most recent version (Merfeld and Morduch, 2025) formalises this as the timecount.

Due to this being part of their ongoing work, it is worth clarifying the relationship between the framework adopted in this paper and the updated terminology in Merfeld and Morduch (2025). In their most recent working paper, the authors introduce the concept of the timecount, defined as the average fraction of the year in which households are poor, measured with monthly consumption data. Based on this, they argue that national poverty statistics in India and other low and middle-income countries have already unintentionally become timecounts rather than conventional headcounts, due to the way household surveys are designed and data is collected. The framework in this paper is conceptually aligned with the timecount: what is referred to here as high-frequency poverty is an approximation to the timecount, and what is referred to as low-frequency poverty corresponds to the conventional headcount based on annual average consumption. The terminology of high and low-frequency poverty is retained throughout this paper for clarity and consistency with the earlier version of M&M's framework (2022), on which this paper is based, but the reader should understand these terms as being broadly equivalent to the timecount and headcount respectively, as defined in Merfeld and Morduch (2025). It should also be noted that household surveys that follow certain data-collection guidelines, including those used in this paper, may already produce aggregate poverty estimates that approximate timecounts. However, this does not mean that the within-year dynamics of poverty are transparent or well understood. Therefore, the contribution of the high-frequency framework is not to produce a different aggregate statistic but to make the temporal structure of poverty explicit by showing not just that the timecount exceeds the headcount, but when, where, and for whom poverty accrues within the year.

3.2. Construction of the poverty line

This section outlines the methodological framework for constructing high and low-frequency poverty measures. The underlying analytical logic here is that poverty is not a static condition experienced uniformly over space and time, but a dynamic outcome shaped by insufficiency, instability, and illiquidity in the financial lives of LIHs. This implies that, across different geographies, fluctuations in income and expenditure, combined with an uneven capacity to smooth consumption, result in deprivations that vary across region-states and can occur intermittently within a year, even when annual headcounts may suggest otherwise. A meaningful measurement of poverty must therefore be sensitive to both spatial differences and temporal variation in the cost of living.

The approach adopted in this paper operationalises this logic by aligning household-level expenditure data with a poverty line that varies by month, region, and state, and evolves over time with changes in prices. Rather than treating the poverty line as a fixed annual threshold, it is constructed as a monthly, region-state-specific benchmark that reflects the household's purchasing power at a given point in space and time. As a result, this approach allows poverty to be studied at a frequency that is closer to how LIHs experience and manage their uncertainties.

Within this framework, poverty is measured at two distinct but related frequencies. High-frequency poverty is defined at the monthly level by comparing household expenditure to an inflation-adjusted household poverty line specific to that month, region, and state. Low-frequency poverty, in contrast, is derived by aggregating these monthly thresholds into an annual benchmark, yielding a yearly measure that is comparable to conventional poverty statistics. Together, these measures allow the analysis to capture movements in and out of deprivation over the course of a year.

The remainder of this section formalises this framework mathematically. It describes how baseline poverty lines are updated using consumer price indices, scaled from per capita to household levels, and aggregated across months to generate consistent high-frequency and low-frequency poverty measures.

Let $s \in S$ denote a state in India, and $r \in R_s$ a region (rural/urban) within the state s such that (r, s) denotes a unique point in space. Let $m \in M_y$ denote the month m in the year y such that (m, y) jointly denotes a unique point in time.

The starting point is the region-state specific per capita monthly poverty line in the base year 2011-12, as stated in the RBI handbook of statistics. It is denoted by:

$$PCMPL_{r,s}^{2012} \quad \dots (1)$$

To account for regional differences and price changes over time, these values are updated using the CPI provided by MoSPI. The CPI for the month m of the year y in region r of the state s is denoted by:

$$CPI_{m,y,r,s}^{2012} \quad \dots (2)$$

It should be noted that the Tendulkar Committee's per capita poverty lines are expressed in 2011-12 prices. Meanwhile, MoSPI's CPI series, published with base 2011-12, is available at the month-year-region-state level from January 2013 onward. Therefore, the 2011-12 per-capita poverty line for region r, sector s, to month m of year y is constructed as:

$$PCmPL_{m,y,r,s} = PCMPL_{r,s}^{2012} \; X \; \frac{CPI_{m,y,r,s}^{2012}}{100} \quad \dots (3)$$

This expression carries the 2011-12 per-capita poverty baseline forward into 2013 (and subsequent months) using the reported CPI levels.

Since CPHS is a household-level survey, the per capita threshold must be scaled to the household level. Let the average household size in year y for region r in state s, assumed to be constant in a given year[5], be denoted by:

$$AHS_{y,r,s} \quad \dots (4)$$

[5] By assuming AHS to be constant in a given year, all households surveyed in that year are assigned the same household size scalar regardless of when during the year they were surveyed. This is a simplifying approximation, as household composition may change within a year, and yet, is necessary since it allows for easier computation and interpretability of the high and low-frequency poverty lines.

Then, the household-level monthly poverty line is defined as:

$$HHmPL_{m,y,r,s} = PCmPL_{m,y,r,s}\ X\ AHS_{y,r,s} \quad \ldots (5)$$

This household-level threshold is then applied to each household's reported expenditure in CPHS to determine its poverty status in month m, year y, region r and state s. This is also known as the high-frequency measure of poverty (HFP).

Finally, the annual household-level poverty line is obtained by averaging the household monthly poverty lines over all 12 months of a given year, i.e.:

$$HHyPL_{y,r,s} = \frac{1}{12}\sum_{m=1}^{12} HHmPL_{m,y,r,s} \quad \ldots (6)$$

This represents the low-frequency measure of poverty (LFP). It denotes the monthly average expenditure (or income) threshold in a given year to determine the household's annual poverty status.

3.3. Sample and population estimates

The CMIE CPHS survey is conducted for a sample of ~1,80,000 households every year, in four waves. Between 2018-2023, there were 1,87,963 unique households in the sample. However, not all of these households are successfully surveyed in every wave[6]. A snapshot of the monthly response rates is presented in the table below:

Table 1: Snapshot of monthly response rates in CMIE CPHS

Month-Year	Sampled households	Surveyed households	Response rate(%)
01/01/18	170183	145091	85.26
01/02/18	170615	146974	86.14
01/03/18	171375	148604	86.71
01/04/18	172365	149160	86.54
01/05/18	172573	148241	85.90
01/06/18	172781	147739	85.51
			… and so on…

This is an important consideration since a high-frequency approach requires that the same households be captured in each month. This attrition makes it challenging to develop high-frequency population estimates using just the sampled households[7]. To resolve this, CMIE has

[6] Each household is surveyed only once during a wave (a single wave lasts for 4 months in a year), where they're asked to recall information in the last 4 months. The household-level monthly data for a single year is produced after incorporating the responses based on this recall period.

[7] Merfeld and Morduch (2025) show that when statistical agencies follow certain guidelines for data collection, specifically, surveying different households across the year while stratifying by period, the resulting standard headcount calculation is not a headcount as conventionally understood but an approximation to the timecount. This is because the sampling process delivers a single randomly-chosen observation of consumption for each household rather than

provided various weights in their monthly data series along with a non-response adjustment factor. The final weight can be computed as:

$$Final\ Weight = Country/State\ level\ weight\ for\ the\ monthly\ series \\ X\ Non - Response\ Adjustment\ Factor$$

After filtering for households that were successfully surveyed, the number of unique households in each year, along with the population represented by these households, is shown in the table below:

Table 2: Surveyed households and their representative population in CMIE CPHS

Year	Surveyed households	Population (in Millions)
2018	168344	281.06 mn
2019	168875	287.18 mn
2020	157674	293.33 mn
2021	162323	299.66 mn
2022	166330	306.02 mn
2023	172533	312.47 mn

The various poverty indicators (discussed in Section 3.5.) can only be calculated for the surveyed households. However, due to high attrition rates in the survey, it was not possible to follow the same household throughout the entire study period. Therefore, to resolve this, where the analysis required tracking the same household over consecutive months, the sample was reduced from ~170,000 surveyed households to ~46,000 households by including only those with data for at least four continuous years during the six-year study period. This implies that there may be a household (or multiple households) who was surveyed for the first time only in the first wave of 2019, rather than at the beginning of the study period in 2018 and continues to be surveyed for at least 48 consecutive months. Therefore, analyses based on this 'sub-sample' are unweighted, i.e., they do not represent the population, since directly applying weights to this sub-sample would yield an incorrect population estimate.

A related concern is whether this sub-sample is systematically different from the complete sample of surveyed households. Households that are successfully surveyed for at least 48 consecutive months are likely to be less residentially mobile and more economically stable than those that are dropped from the survey[8]. This introduces a potential selection bias: the sub-sample used in the analyses may under-represent the variety of households (including LIHs) whose poverty experience the high-frequency framework is designed to capture.

their true annual average, and when aggregated across households, these observations converge to the timecount rather than the headcount. While the CPHS similarly surveys households continuously across months, this paper takes a different approach: rather than relying on the cross-sectional sampling process to approximate the timecount, it constructs high-frequency poverty measures deliberately using longitudinal household-level data, tracking the same households across consecutive months. This allows for a more direct examination of within-year poverty dynamics than what the cross-sectional approximation permits.

[8] According to CMIE, households are typically dropped from the survey due to factors that indicate high residential mobility. This includes the residence being locked and unoccupied for consecutive waves and migration of household members or the entire family, more often than not for work, thereby suggesting low economic stability.

3.4. Measuring poverty using income versus expenditure

Measuring poverty at a higher frequency heightens the distinction between income-based and expenditure-based measures (Atkinson, 2019; Bradbury et al. 2001). While poverty is conventionally measured using yearly sums of either income or expenditure, according to M&M, the high-frequency framework reveals that these two approaches offer very different insights into the lived experience of deprivation. Nonetheless, the measurement of poverty requires making a choice between them. This trade-off is motivated by how they account for uncertainty and the household's ability to smooth consumption over time. Broadly, there are two questions to consider when making this tradeoff: (i) what does the measure intend to capture? and (ii) what kind of data is available, and how reliable is it? These are discussed below in further detail.

The choice between income and expenditure-based measures depends fundamentally on what each measure intends to capture. This matters because each measure carries an implicit assumption about a household's ability to move resources across time. An annual income measure implicitly assumes perfect smoothing: it treats the yearly income total as a proxy for measuring welfare and, in doing so, presumes that a household can freely reallocate income across the months of the year to meet needs as they arise. Meanwhile, a monthly income measure classifies a household as poor whenever its income in a given month falls short of the poverty line. It presumes that income must be exhaustively consumed in each month and that no resources can be carried over across periods. Regardless of whether income is measured at annual or monthly frequencies, as an indicator of welfare, it reflects a perfect state of the world where either uncertainties don't manifest or are perfectly understood and addressed by all households. On the other hand, expenditure as the relevant indicator of welfare lies in the space between these two scenarios: it measures what households actually managed to consume in a given period and therefore reflects the outcome of the partial and imperfect smoothing they achieve through various coping strategies[9]. This discussion also carries forward to a later section (Section 4.3), where the hypothesis regarding household consumption smoothing is tested using two separate econometric models.

The divergence between income and expenditure measures is thus not just an artefact of measurement, but a substantive result in itself (see Figures 1 and 2). The former reflects high volatility, especially among LIHs, whose earnings are sporadic and unpredictable. Meanwhile, the latter is typically less variable because it reflects households' ability to smooth consumption over time.

[9] Mas and Murthy (2017) argue that the financial lives of the LIHs are intertwined with their social lives, and their decisions are motivated by an intuition that has organically developed over a long history of shared experiences. This intuition relies on hierarchies and routines to manage uncertainties. Hierarchies help establish priorities, while spending and saving routines help manage the balance between inflows and outflows. Understood this way, financial decisions of LIHs reflect not just preferences or constraints but also an ongoing struggle for control amid uncertainty. As a result, LIHs have developed three coping strategies to deal with uncertainty: scheduling liquidity to manage volatile income flows, farming liquidity by nurturing social and business relationships, and bottling liquidity to protect it from unwarranted use.

Figure 1: (Unweighted) Median per-capita Income and Expenditure over time

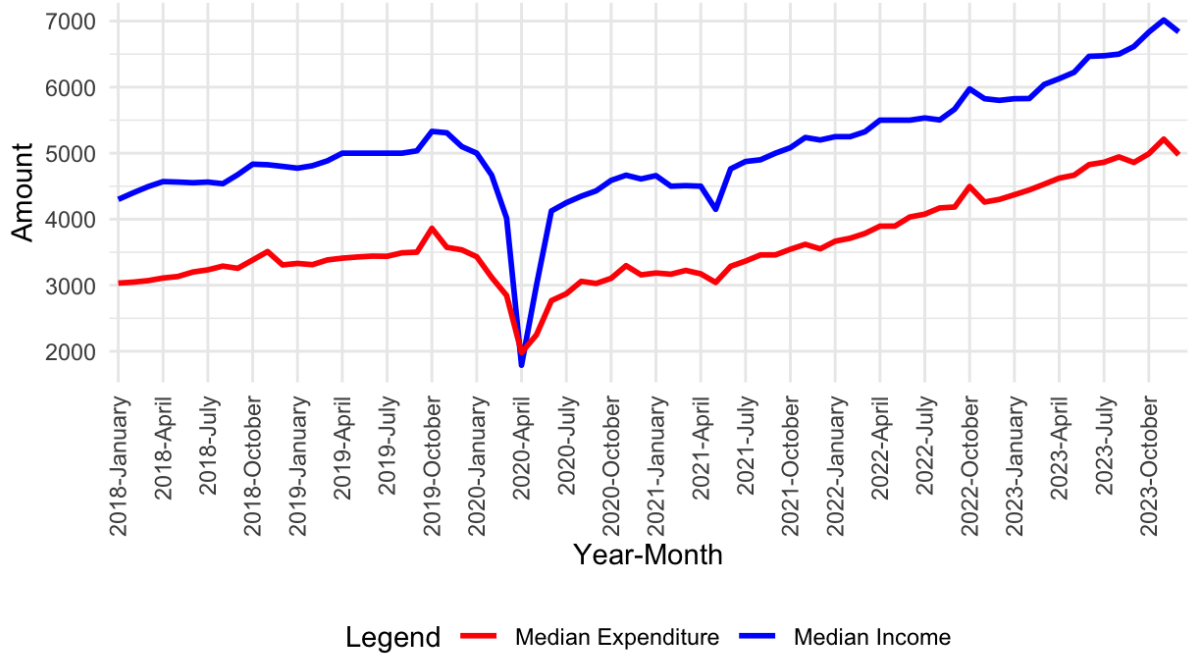
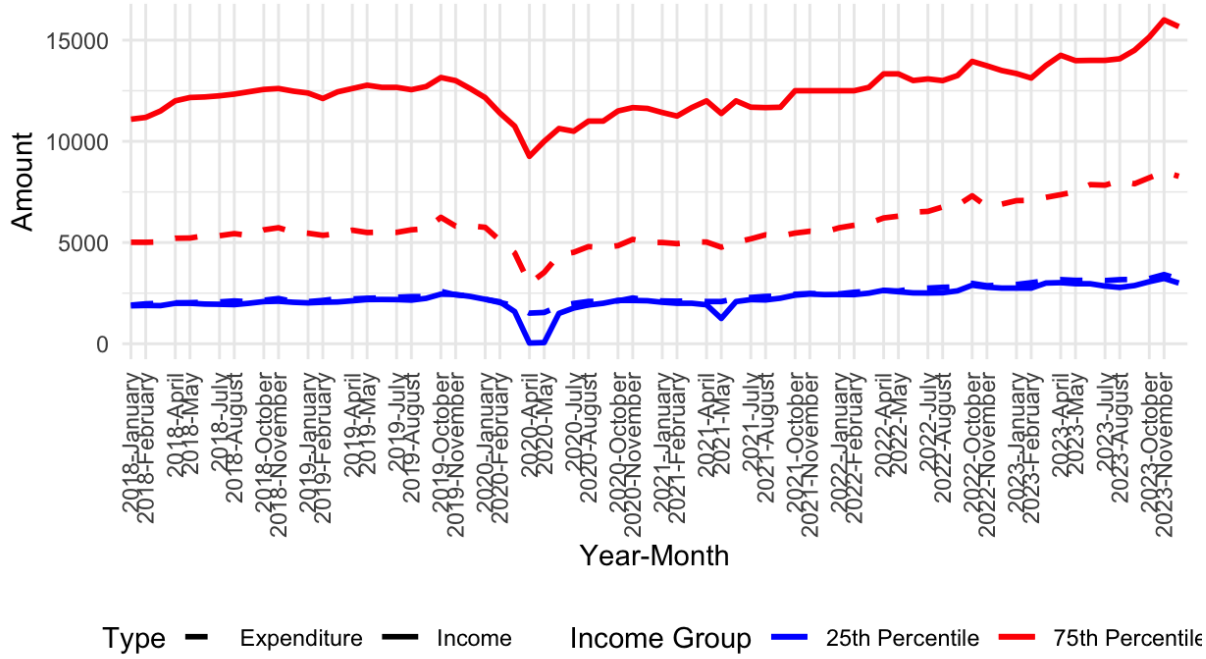


Figure 1 plots the unweighted median per-capita income and expenditure of 46,382 households in the sub-sample. As stated earlier, this sub-sample consists of households that were successfully surveyed for at least 48 consecutive months between 2018 and 2023. This is to facilitate the calculation of a stable median per-capita income and expenditure of households across the entire study period. The relative stability of the number of surveyed households in this sub-sample enables calculation of the coefficient of variation (CV). The median CV based on income is estimated at 36.56, and the median CV based on expenditure is estimated at 22.80. A lower median CV based on expenditure implies that households experience lower variability in expenditure when compared to their incomes. This suggests that households are indeed smoothing their consumption in the presence of uncertain incomes. Additionally, the median per capita income and expenditure of households in the 25th and 75th percentile of income is shown in Figure 2.

Figure 2: (Unweighted) Median per capita income & expenditure (for households in the 25th and 75th Income Percentiles)



Here, the median CV based on income for households in the 25th and 75th income percentiles is 32.06 and 15.77, respectively. Meanwhile, the median CV based on expenditure is 17.74 and 20.36, respectively. This supports the hypothesis that LIHs (in the 25th percentile) experience significantly higher income volatility than their counterparts at the 75th percentile. Interestingly, however, the ratio of the CV of expenditure to the CV of income tells a different story about consumption smoothing. Households in the 25th percentile are able to reduce their income volatility considerably, given that their CV of expenditure is roughly half of their CV of income, thereby suggesting active consumption smoothing in the face of high volatility in income. In contrast, households in the 75th percentile exhibit a CV of expenditure that only marginally exceeds their CV of income, possibly revealing a discretionary spending pattern rather than a need to smooth their consumption.

The second consideration is the availability and reliability of the data itself. Most low- and middle-income countries measure poverty using expenditure rather than income (Mancini and Vecchi, 2022). In contexts such as India, where a large share of households derive their livelihoods from agriculture, casual labour, and self-employment, income streams are diversified and highly irregular. Moreover, households don't usually maintain separate accounts for business and personal use, and net income from own-farm or own-enterprise activity requires imputing the value of home production and netting out input costs, both of which are susceptible to error (Deaton, 1997; Deaton and Zaidi, 2002). Even where household surveys are systematically designed to capture income, Carletto et al. (2021) find that such income is often underreported. Ultimately, there is a general consensus among researchers and policymakers that expenditure is the more stable series and is therefore widely considered as the better proxy to understand households' financial behaviour with respect to their resources (Deaton, 1997; Ravallion, 2016).

In India, the consensus that expenditure is the preferred measurement of welfare is reflected in the NSSO's Household Consumer Expenditure Survey. The Survey collects information regarding household consumption but not income, and both the Tendulkar (2009) and Rangarajan (2014)

poverty lines are defined in consumption terms. This directly impacts this exercise since the poverty line used here is derived from the Tendulkar Committee's MRP consumption baseline, and therefore, household expenditure is the only quantity that can be compared against it consistently. Applying an expenditure-derived threshold to an income series would amount to a measurement error and would make it impossible to derive any meaningful conclusions from the results. Household expenditure is therefore used as the primary measure of poverty throughout this paper.

3.5. Poverty Indicators

In M&M's paper, the authors make the use of four decomposable measures of poverty, namely, poverty headcount, poverty gap, the squared poverty gap of Foster et al. (1984), and the Watts Index. To ensure consistency, this paper also calculates these four poverty indicators using total household reported expenditure, at both high and low frequencies.

3.5.1. Poverty headcount

Poverty headcount is the poverty measure that dominates traditional policy discourse. Most countries determine poverty using this measure, given its simple computation and easy interpretability. It is described as the proportion of poor households in a given population and serves as a viable baseline for comparison to other indicators of poverty. Its simplicity, however, comes with a natural boundary: the headcount treats all poor households as identical, regardless of how far below the poverty line they fall. As a result, a household spending just one unit less than the poverty line is counted the same as one spending half of it. This is not a limitation so much as a reflection of what the headcount is designed to do, i.e., it counts the number of households lying below the poverty line rather than measuring the degree of deprivation for such households. As a result, questions about the depth or severity of poverty, or about inequality among the poor, fall outside its scope by design. In a high-frequency setting, this boundary becomes particularly relevant, since the framework is specifically designed to detect degrees in deprivation across months and geographies that the headcount does not attempt to capture. Nonetheless, it serves as a useful baseline for capturing the incidence of poverty, and the indicators that follow are introduced precisely to address the dimensions that lie beyond its scope.

For each household i in month m, year y, state s, and region r, the high frequency poverty status is determined by:

$$Poverty\ Status^{hf}_{i_{(m,y,r,s)}} = \begin{cases} 1, if\ Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s} \\ 0, otherwise \end{cases} \quad \dots (7)$$

If there are N households, each observed for up to 12 months, then the high-frequency poverty headcount in each year across all households is:

$$Poverty\ Headcount^{hf} = \frac{1}{N\ X\ 12} \sum_{i=1}^{N} \sum_{m=1}^{12} Poverty\ Status^{hf}_{i_{(m,y,r,s)}} \quad \dots (8)$$

Similarly, for each household i in year y, state s, and region r, let the (annual) average monthly expenditure be donated by:

$$(Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}} = \frac{1}{12}\sum_{m=1}^{12} Monthly\ Expenditure_{i_{(m,y,r,s)}} \quad \dots (9)$$

This averaging is necessary to ensure symmetry with the annual household poverty line in equation (6), which is itself expressed as a monthly average. Comparing average monthly expenditure against an average monthly threshold ensures that the low-frequency measure is internally consistent and comparable to its high-frequency counterpart. Then, the low-frequency poverty status for each household is determined by:

$$Poverty\ Status^{lf}_{i_{(y,r,s)}} = \begin{cases} 1, if\ (Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}} \leq HHyPL_{y,r,s} \\ 0, otherwise \end{cases} \quad \dots (10)$$

Aggregating across N households gives the low-frequency poverty headcount in each year:

$$Poverty\ Headcount^{lf} = \frac{1}{N}\sum_{i=1}^{N} Poverty\ Status^{lf}_{i_{(y,r,s)}} \quad \dots (11)$$

3.5.2. Poverty Gap

The (normalised) poverty gap addresses the headcount's insensitivity to depth by measuring the average shortfall of each household's expenditure from the poverty line, expressed as a proportion of the poverty line. Intuitively, it answers the question of how poor the poor are, rather than simply how many are poor. It registers each household's degree of deprivation relative to the poverty line, with the weight of each unit of money below the threshold being constant. In a high-frequency setting, this is particularly useful: a household may remain below the poverty line in consecutive months, but the poverty gap will capture whether its situation is deteriorating or improving within that period. However, because the weight attached to each unit of shortfall is constant regardless of how far below the poverty line a household falls, a transfer of resources from the poorest household to one that is marginally poor but still below the threshold will increase the former's shortfall and reduce the latter's by exactly the same amount, leaving the aggregate poverty gap index unchanged. Therefore, to this extent, the poverty gap is distributionally insensitive[10], since it

[10] Amartya Sen's Transfer Axiom (Sen, 1976) formalises distributional sensitivity as an intuitive principle: if a unit of money is transferred from a poorer household to a less poor one, and both remain below the poverty line, then measured poverty should increase or, at the very least, not decrease. The axiom captures the idea that the welfare loss suffered by the household at the bottom outweighs the gain accruing to the one closer to the poverty line. Measures that satisfy this axiom are said to be distributionally sensitive, because they respond to the shape of the distribution of consumption among the poor, not just its count. The poverty headcount fails this test entirely: since it only counts the number of households below the threshold, a transfer between two poor households leaves it completely unchanged, regardless of how extreme the underlying inequality among the poor might be. The poverty gap also fails: since it weights each unit of shortfall equally, a transfer from a very poor household to a marginally poor one increases the former's shortfall and reduces the latter's by exactly the same amount, leaving the aggregate gap index unchanged. A measure that is truly distributionally sensitive should register this as a deterioration of a manner wherein the very

captures the total depth of deprivation but cannot detect whether that deprivation is concentrated among the most vulnerable or spread more evenly below the threshold. It is this limitation that motivates the next two indicators.

For each household i in month m, year y, state s, and region r, the high-frequency poverty gap is determined by:

$$Poverty\ Gap^{hf}_{i_{(m,y,r,s)}} = \left(\frac{HHmPL_{m,y,r,s} - Monthly\ Expenditure_{i_{(m,y,r,s)}}}{HHmPL_{m,y,r,s}}\right) \times 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} \quad \ldots (12)$$

$$where\ 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Aggregating across all household months, the high-frequency poverty gap index in a given year becomes:

$$Poverty\ Gap\ Index^{hf} = \frac{1}{N\ X\ 12}\sum_{i=1}^{N}\sum_{m=1}^{12} Poverty\ Gap^{hf}_{i_{(m,y,r,s)}} \quad \ldots (13)$$

Similarly, for each household i in year y, state s, and region r, the low-frequency poverty gap is:

$$Poverty\ Gap^{lf}_{i_{(y,r,s)}} = \left(\frac{HHyPL_{y,r,s} - (Annual)Average\ Monthly\ Expenditure_{i_{,(y,r,s)}}}{HHyPL_{y,r,s}}\right) \times 1_{((Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}} \leq HHyPL_{y,r,s})} \quad \ldots (14)$$

$$where\ 1_{((Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}} \leq HHyPL_{y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Aggregating across N households gives the low frequency poverty gap index in each year as:

$$Poverty\ Gap\ Index^{lf} = \frac{1}{N}\sum_{i=1}^{N} Poverty\ Gap^{lf}_{i_{(y,r,s)}} \quad \ldots (15)$$

3.5.3. Foster's Squared Poverty Gap

poor household has become poorer. In other words, neither the headcount or the poverty gap can distinguish between a situation where the poor are clustered just below the poverty line and one where some are in extreme deprivation while others are barely poor. This distinction is particularly significant in a high-frequency setting where such differences can emerge and dissolve within a single year.

The squared poverty gap index addresses the distributional insensitivity of both the headcount and the poverty gap by giving greater weight to households that are further below the poverty line. It does so by squaring the normalised shortfall, which results in disproportionate penalisation of the larger expenditure gaps. Intuitively, this means that a household consuming at half the poverty line contributes far more to the index than one consuming just below it. As a result, the squared poverty gap is sensitive not only to the incidence and depth of poverty but also to the inequality of consumption among the poor. In a high-frequency setting, this property is particularly valuable because it can detect whether the most deprived households are becoming more or less concentrated at the bottom of the distribution across months, even when aggregate headcounts remain stable.

For each household i in month m, year y, state s, and region r, the high frequency Foster's Squared poverty gap is determined by:

$$Foster's\ Squared\ Poverty\ Gap^{hf}_{i_{(m,y,r,s)}} = \left(\frac{HHmPL_{m,y,r,s} - Monthly\ Expenditure_{i_{(m,y,r,s)}}}{HHmPL_{m,y,r,s}}\right)^2 \times 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} \quad \ldots (16)$$

$$where\ 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Aggregating across all household months, the high-frequency Foster's Squared poverty gap index in a given year becomes:

$$Foster's\ Squared\ Poverty\ Gap\ Index^{hf} = \frac{1}{N\ X\ 12}\sum_{i=1}^{N}\sum_{m=1}^{12} Foster's\ Squared\ Poverty\ Gap^{hf}_{i_{(m,y,r,s)}} \quad \ldots (17)$$

Similarly, for each household i in year y, state s, and region r, the low frequency Foster's Squared poverty gap is:

$$Foster's\ Squared\ Poverty\ Gap^{lf}_{i_{(y,r,s)}} = \left(\frac{HHyPL_{y,r,s} - (Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}}}{HHyPL_{y,r,s}}\right)^2 \times 1_{((Annual)Average\ Monthly\ Expenditure_{i_{y,r,s}} \leq HHyPL_{y,r,s})} \quad \ldots (18)$$

$$where\ 1_{((Annual)Average\ Monthly\ Expenditure_{i_{y,r,s}} \leq HHyPL_{y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Accordingly, the low frequency Foster's Squared poverty gap index becomes:

$$Foster's\ Squared\ Poverty\ Gap\ Index^{lf} = \frac{1}{N}\sum_{i=1}^{N} Foster's\ Squared\ Poverty\ Gap^{lf}_{i_{(y,r,s)}} \quad \ldots (19)$$

3.5.4. Watts Index

The Watts index, like the squared poverty gap, is sensitive to the depth and severity of poverty, but it approaches this differently by measuring the logarithmic proportional shortfall from the poverty line. Because the logarithm is a concave function, it assigns progressively greater weight to shortfalls as a household moves further below the poverty line, meaning that the index is more sensitive to the poorest of the poor than the squared poverty gap. In practical terms, similar to the squared poverty gap, the Watts index penalises extreme deprivation more heavily than moderate deprivation, making it a useful complement to the squared poverty gap when the concern is specifically with the welfare of the most vulnerable households. In a high-frequency setting, this sensitivity makes it a particularly informative indicator of whether the worst periods of deprivation within a year are becoming more or less severe over time.

For each household i in month m, year y, state s, and region r, the high frequency Watts Index is determined by:

$$Watts^{hf}_{i_{(m,y,r,s)}} = \ln\left(\frac{HHmPL_{m,y,r,s}}{Monthly\ Expenditure_{i_{(m,y,r,s)}}}\right) \times 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} \quad \ldots (20)$$

$$where\ 1_{(Monthly\ Expenditure_{i_{(m,y,r,s)}} \leq HHmPL_{m,y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Aggregating across all household months, the high-frequency Watts index in a given year becomes:

$$Watts\ Index^{hf} = \frac{1}{N\ X\ 12}\sum_{i=1}^{N}\sum_{m=1}^{12} Watts^{hf}_{i_{(m,y,r,s)}} \quad \ldots (21)$$

Similarly, for each household i in year y, state s, and region r, the low-frequency Watts index is:

$$Watts^{lf}_{i_{(y,r,s)}} = \ln\left(\frac{HyPL_{y,r,s}}{(Annual)Average\ Monthly\ Expenditure_{i_{(y,r,s)}}}\right) \times 1_{((Annual)Average\ Monthly\ Expenditure_{i_{y,r,s}} \leq HHyPL_{y,r,s})} \quad \ldots (22)$$

$$where\ 1_{((Annual)Average\ Monthly\ Expenditure_{i_{y,r,s}} \leq HHyPL_{y,r,s})} = \begin{cases} 1, Poor\ household \\ 0, otherwise \end{cases}$$

Accordingly, the low frequency Watts Index becomes:

$$Watts\ Index^{lf} = \frac{1}{N}\sum_{i=1}^{N} Watts^{lf}_{i_{(y,r,s)}} \quad \ldots (23)$$

4. Results

This section is divided into four parts. Section 4.1 is an exploration of the data, which establishes the empirical context before any formal analysis is undertaken. Section 4.2 discusses the temporal dimension and provides evidence to support M&M's findings directly. Section 4.3 tests the behaviour of households with respect to consumption smoothing in the face of uncertainty. Finally, Section 4.4. reproduces M&M's policy experiment to ascertain the impact of a hypothetical cash transfer scheme for poverty alleviation. Together, these four parts build a cumulative case for the high-frequency framework as a more faithful representation of how LIHs experience poverty.

4.1. Data exploration

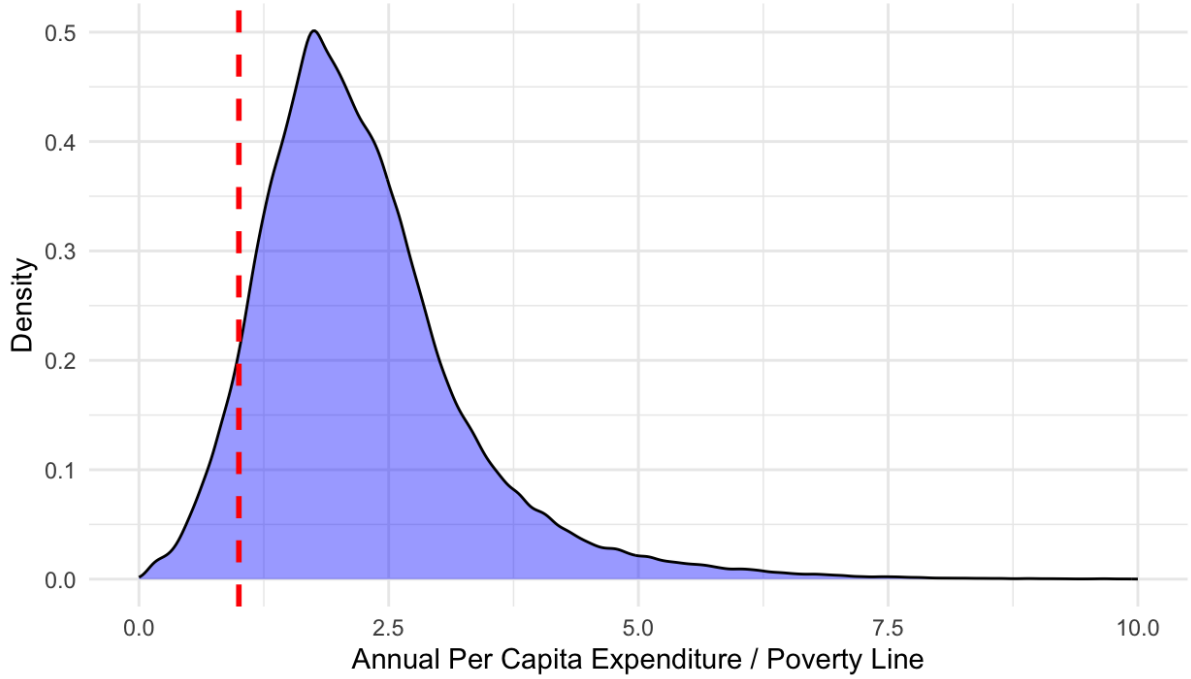
This section provides an exploratory view of the data to establish preliminary observations on the financial lives of LIHs and how uncertainty manifests in their lives. This is to aid and contextualise the analyses that follow. Section 4.1.1 examines the shape of the distribution of household expenditures relative to the poverty line, revealing that a substantial share households sustain themselves in close proximity to the annual poverty threshold. This is followed by Section 4.1.2, which presents temporal poverty summary statistics at both low and high frequency, generating initial evidence that annual measures systematically understate the incidence and depth of deprivation. Together, the two subsections establish the empirical foundation for the analysis that follows and motivate the central argument of this paper, which is that when poverty is measured only at the low frequency or annual level, it obscures a more complex reality that the high-frequency framework is designed to highlight.

4.1.1. Kernel density estimate of annual per-capita expenditure

To understand the shape of the distribution of household expenditures relative to the poverty line, figure 3 plots the kernel density of normalised annual per capita expenditure, defined as each household's annual per capita expenditure divided by its corresponding annual per capita poverty line, pooled across all six years of the study period[11]. Values to the left of 1 indicate households whose annual per-capita expenditure falls below the poverty line and are therefore classified as poor under the low-frequency measure. Values to the right of 1 indicate households whose expenditure exceeds the poverty line and are classified as non-poor.

[11] Each household's annual per capita expenditure is computed as per equation 9, which is then normalized using the poverty line. Then, this is further pooled for all six years for each household. Since these are normalized values, the density of the kernel is unaffected by the number of household-level data points.

Figure 3: (Unweighted) k-density estimate of Annual per-capita expenditure



The resulting distribution is unimodal and right-skewed, with its mode falling slightly to the right of the poverty line at approximately 1.2 times the poverty line. The height of the curve at this point, approaching 0.5, indicates that a large share of the sample is clustered just above the poverty line. The right tail of the distribution is long but thin, indicating a smaller proportion of households with expenditures that significantly exceed the poverty line. The area to the left of the dashed red line at 1, which represents the share of household-year observations classified as poor under the low-frequency measure, is critical, though it constitutes a minority of the total area under the curve.

Two features of this k-density distribution are worth highlighting. First, the concentration of households just above the poverty line suggests that a substantial share of the sample, while not classified as poor on an annual basis, lives in close proximity to the threshold. For such households, even a temporary uncertainty could push them below the poverty line, even if their annual average expenditure keeps them above it. This is precisely the kind of vulnerability that the low-frequency headcount obscures and that the high-frequency framework is designed to surface. Second, the right skew of the distribution implies that mean expenditure lies above the median, driven by a relatively small number of households with very high normalised expenditure.

4.1.2. Poverty Summary Statistics

Table 3 presents summary statistics for three poverty indicators, computed at both low and high frequency, for the full weighted sample and separately for households classified as poor and non-poor under the low-frequency measure.

Table 3: (Weighted) Temporal Poverty Summary Statistics

	Everyone	Poor for the year	Not poor for the year
Weighted proportion		0.02	0.98
Mean yearly poverty Headcount	0.02	1.00	0.00
Mean months-in-poverty Headcount	0.04	0.78	0.02
Mean Yearly Watts	0.004	0.192	0.000
Mean months-in-poverty Watts	0.010	0.267	0.005
Mean yearly squared poverty gap	0.0009	0.045	0.000
Mean months-in-poverty squared poverty gap	0.0028	0.078	0.001

The three columns divide the households by low-frequency poverty status. The first, 'Everyone', reports each statistic for the full weighted sample. The second, 'Poor for the year', is restricted to households whose average monthly expenditure over the twelve months falls below the annual poverty line. Finally, the third column, 'Not poor for the year', restricts to those household for whom the average monthly expenditure over the twelve months is above the annual poverty line.

The first metric is weighted proportion, which gives the share of the population in each group, i.e., 0.02, or 2 percent of households, on average, are poor in a given year, while the remaining 98 percent are not poor. This is the benchmark against which every other metric should be read.

The remaining rows in the table denote each indicator in two related forms. The yearly version applies the poverty measure to a household's average monthly expenditure over the year, i.e., the low-frequency, or annual measure. The months-in-poverty version applies the same measure for twelve months of any given year, i.e., the high-frequency measure, and it can be read directly as the fraction of the year (in months) the average household spends in poverty.

The mean yearly poverty headcount for the full sample is 0.02, or 2 percent, which by construction equals the weighted proportion of low-frequency poor households. In other words, on average, 2 percent of households are found to be poor in any given year of the study period. The mean months-in-poverty headcount, however, is 0.04. In terms of percentages, this says that 4 percent of all household-months are spent in poverty. Alternatively, the average household spends 0.04 X 12 ≈ 0.48 months of the year below the poverty line. The high-frequency headcount is thus, significantly larger than its low-frequency counterpart. This is the first and most direct piece of

evidence that the annual measure understates the incidence of poverty by obscuring spells of deprivation that occur within the year.

The disaggregation by low-frequency poverty status reinforces this further. Among households classified as poor for the year, the mean months-in-poverty headcount is 0.78, implying that, on average, these households are poor in roughly 9 out of 12 months (0.78 X 12 ≈ 9.4). This figure falling short of 1.0 implies that they are not poor in all twelve months, which suggests that even among the chronically poor, there is considerable within-year variation in poverty status that the annual measure does not capture. Among households classified as non-poor for the year, the mean months-in-poverty headcount is 0.02 or 2 percent, which, while small in absolute terms, is non-zero. This confirms that on average, these households spend about 0.24 months (0.02 X 12) below the poverty line despite their annual expenditure placing them above it.

The same pattern holds for all remaining indicators. The Watts index and the Squared Poverty Gap differ from the headcount in that they register not only whether a household is poor but how far below the line it falls, giving greater weight to deeper shortfalls. As shown in the table, the yearly Watts index of 0.004 rises to 0.010 at monthly frequency, and the yearly Squared Poverty Gap of 0.0009 increases to 0.0028. In proportional terms these increases are larger than the difference observed for the headcount. This is because the more sensitive a measure is to the depth of deprivation, the more within-year variation it uncovers. The same patterns emerge within each subgroup: among the annually poor, the Watts index rises from 0.192 to 0.267 and the squared poverty gap from 0.045 to 0.078. Taken together, the table establishes that measuring poverty at higher frequency highlights not only its measured incidence but also its depth and severity.

4.2. Poverty of Time

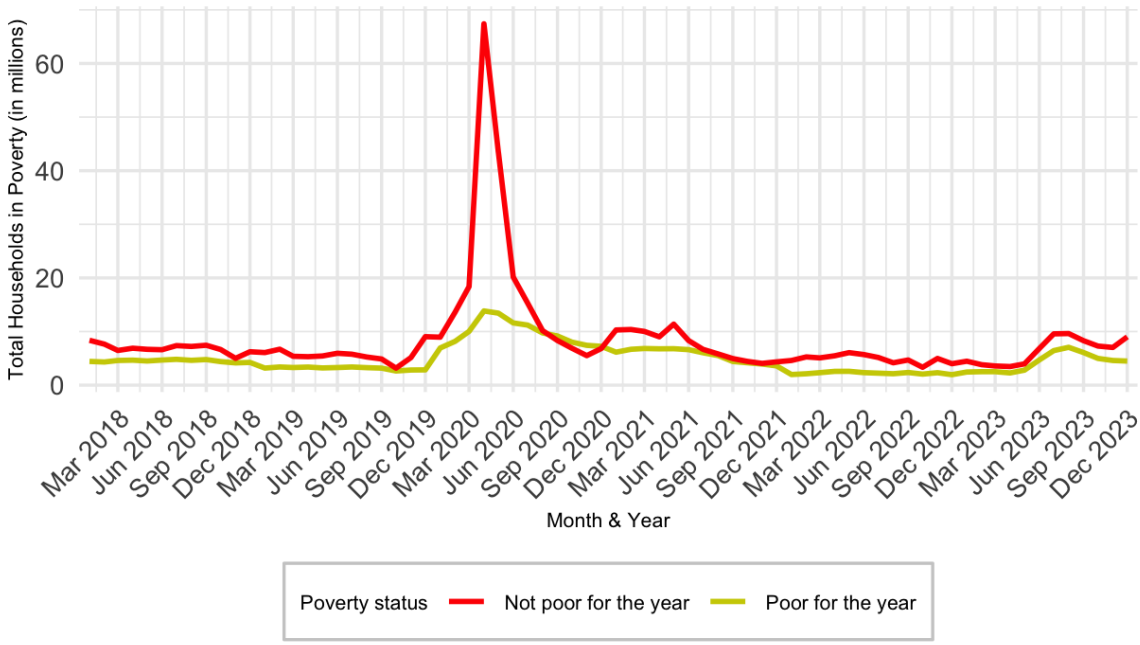
This section analyses the temporal dimension of poverty, and in doing so, closely follows the analytical framework of M&M. The results presented here are drawn directly from M&M's original analysis (2022), and this section should be read as a reproduction of their work using CPHS data and a dynamic poverty line. Notably, where the results diverge from M&M's findings, these differences are discussed in Section 5.

Section 4.2.1 decomposes monthly poverty by low-frequency poverty status to assert that poverty is more widespread than annual measures suggest. Section 4.2.2 examines the household's exposure to monthly poverty and shows that being classified as non-poor does not imply immunity to deprivation. Section 4.2.3 extends the comparison between low and high frequency poverty to distributionally sensitive indicators. Section 4.2.4 examines households at the cusp of annual poverty classification. Together, these build a case for the high-frequency framework as a more faithful representation of how households experience poverty while navigating uncertainty in their financial lives.

4.2.1. Monthly poverty

The summary statistics in Table 3 established that the high-frequency poverty headcount is significantly larger than its low-frequency counterpart, and that households classified as non-poor on an annual basis nonetheless register a non-zero months-in-poverty headcount. Figures 4 and 5 dig deeper into this finding by examining the weighted number of households contributing to the high-frequency poverty headcount in each month, across both low-frequency poor and non-poor households.

Figure 4: (Weighted) Monthly poverty by yearly status (counts)



To construct these figures, households are first classified as low-frequency poor or non-poor in each year based on whether their average monthly expenditure falls below the annual household poverty line defined in equation (6). Within each month, the weighted count of households experiencing high-frequency poverty is then plotted separately for the low-frequency poor and non-poor subgroups. It is worth noting that the monthly totals shown here do not sum to the full population, since only households experiencing high-frequency poverty in that month are included. In other words, monthly non-poor households within both subgroups are excluded.

Figure 4 reveals that in every month of the study period, the absolute number of low-frequency non-poor households experiencing high-frequency poverty exceeds the number of low-frequency poor households. With the exception of the COVID-19 spike in mid-2020, the weighted count of households are generally stable, with the low-frequency non-poor households consistently tracking between 8 and 10 million households in monthly poverty and the low-frequency poor households remaining between 3 and 5 million.

Figure 5: (Weighted) Proportion of Monthly Poverty due to yearly non-poor households

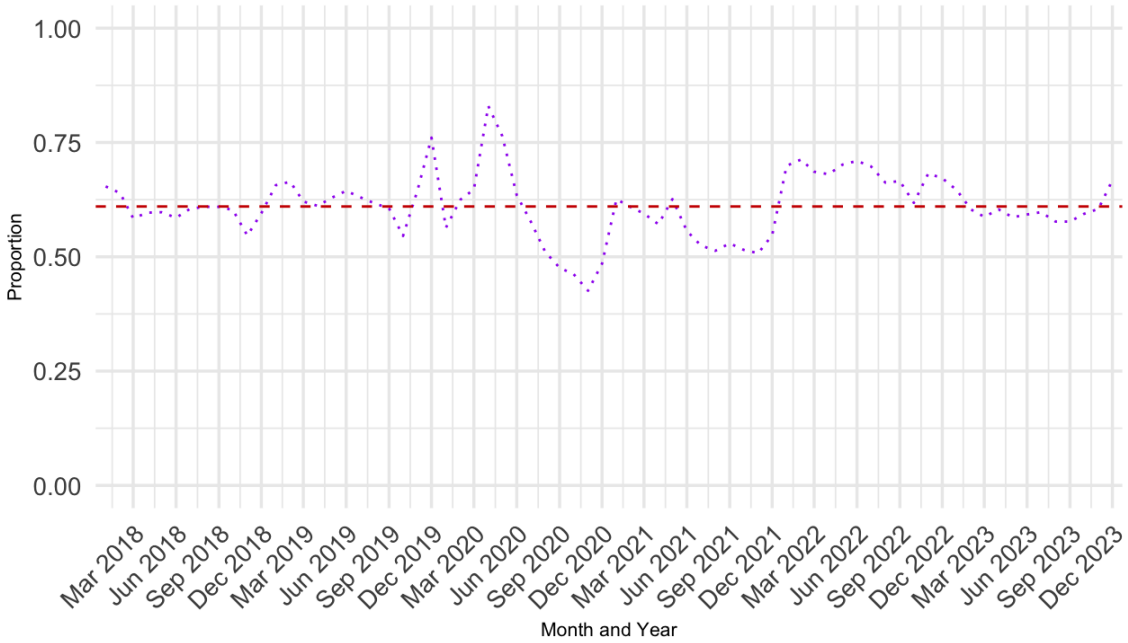


Figure 5 translates this into proportional terms. In each month, it plots the share of all high-frequency poor households that belong to the low frequency non-poor group. The dashed red line, representing the period average, sits at approximately 0.60, implying that across the study period, roughly 60 percent of all household-months spent in poverty are contributed by households that are not classified as poor on an annual basis.

This finding confirms the first result of the high-frequency framework: poverty is more widespread than annual measures suggest. A majority of the monthly poverty observed in the sample is invisible to conventional measurement, not because it does not exist, but because it is averaged away when expenditure is aggregated to the annual level.

It should be noted, however, that this result is partly a mechanical consequence of the sample composition. Low-frequency poor households represent only 2 percent of the weighted sample, as shown in Table 3, while non-poor households account for the remaining 98 percent. Even if non-poor households experience monthly poverty at a much lower rate than poor households, their sheer numerical dominance means they will contribute a disproportionately large share of the total count of household-months in poverty. Figure 5 should therefore be interpreted with this in mind: it does not imply that non-poor households are more vulnerable than poor ones, but rather that the large base of non-poor households amplifies their aggregate contribution to monthly high-frequency poverty counts.

Another thing to note is that the COVID-19 period represents a clear structural break in both figures. Figure 4 shows a sharp spike in monthly poverty in mid-2020, driven by the non-poor households, which momentarily exceeds 60 million. Figure 5 shows the corresponding volatility during this period, with the proportion of monthly poverty attributable to non-poor households fluctuating sharply before stabilising by late 2020. As noted in Section 2.1, this period reflects an

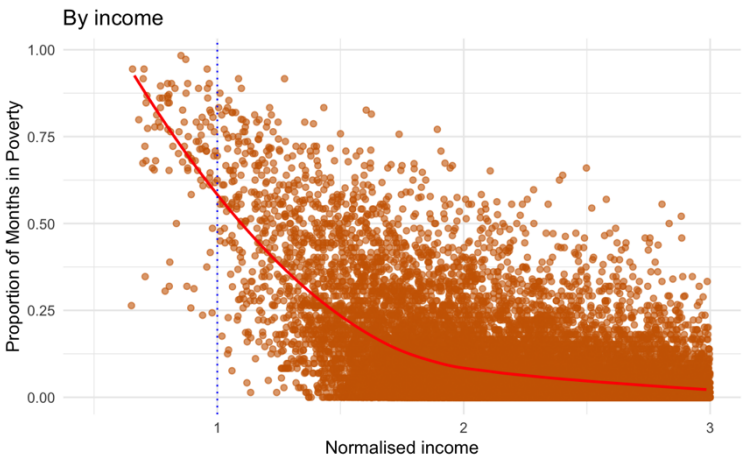
unusual disruption rather than the routine within-year uncertainty that the high-frequency framework is designed to capture, and results from this period should be interpreted accordingly.

4.2.2. Poverty Exposure

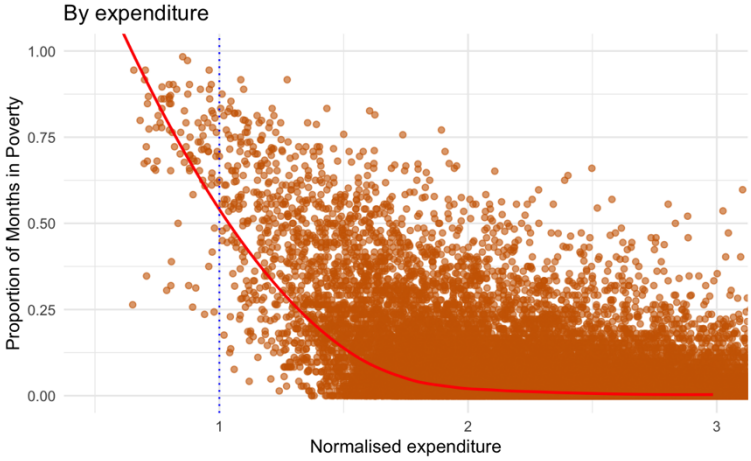
The monthly poverty analysis in Section 4.2.1 established that a substantial share of households classified as non-poor on an annual basis nonetheless experience at least some months of poverty within the year. However, it does not speak to the relationship between a household's average living standard and its vulnerability to monthly poverty. A household just above the poverty line and one at twice the poverty line are both classified as non-poor, yet their exposure to deprivation is likely very different. Figure 6 addresses this by examining how a household's proportion of months spent in poverty varies with its average income and expenditure relative to the poverty line.

Figure 6: (Unweighted) Poverty Exposure

Panel A



Panel B



For each household in the sub-sample, the monthly average normalised income and expenditure are computed by taking the mean across all six years of the study period and dividing by the average high-frequency poverty line across the same period. The proportion of months in poverty is similarly averaged across all six years. Each point in the scatter plot, therefore, represents a single household's average position over the entire study period, and the smooth red curve is estimated using local polynomial regression.

Two observations can be made from Figure 6. First, the poverty exposure curve is downward sloping, indicating that even households with normalised income or expenditure marginally above 1 are exposed to a considerable proportion of months in poverty. Second, the curve does not approach zero until household expenditure reaches approximately twice the poverty line (thrice in the case of household income). This implies that a household should spend (or earn) at roughly double (triple in the case of household income) the poverty threshold before its exposure to monthly poverty becomes negligible. Being classified as non-poor, in other words, is not the same as being completely insulated from exposure to poverty.

These two observations have direct implications for the high-frequency framework. A binary classification of households as poor or non-poor, whether based on monthly or annual measures, obscures the degree of deprivation among households both above and below the poverty line. The

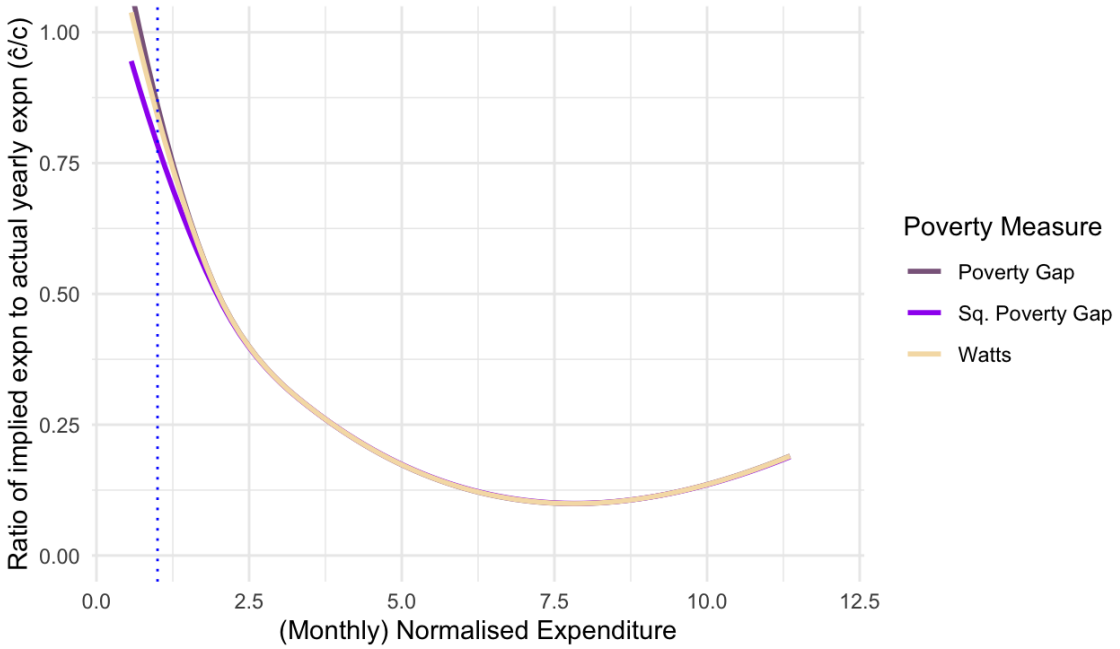
poverty exposure curve, in turn, makes this explicit: it shows that the risk of experiencing poverty in any given month diminishes gradually and non-linearly with increasing distance from the poverty line, rather than dropping to zero at the threshold as a headcount measure would suggest otherwise.

Even though the two panels tell a consistent story, there is a notable difference in the dispersion of points around the curve. The income-based panel exhibits a wider spread than the expenditure-based panel. This is consistent with the higher coefficient of variation of income relative to expenditure documented in Section 3.4, which indicates that income is more volatile than expenditure, and that a household's average normalised income is therefore a noisier signal of its underlying living standard than its average normalised expenditure. The tighter clustering of points around the red curve in the expenditure panel reflects the fact that expenditure is a more reliable predictor of a household's actual exposure to poverty, as previously discussed in Section 3.4. This reinforces the earlier methodological choice to use expenditure as the primary measure of poverty in this analysis.

4.2.3. Monthly vs. Yearly Poverty for Distributionally Sensitive Indicators

The analyses in sections 4.2.1 and 4.2.2 relied primarily on the poverty headcount to compare low- and high-frequency poverty. While useful for establishing the broader result that poverty is more widespread than annual measures suggest, the headcount is distributionally insensitive, i.e., it treats all households below the poverty line identically, regardless of how far below it they fall. As established in Section 3.5, distributionally sensitive measures such as the squared poverty gap and the Watts index are designed to capture precisely this dimension, penalising larger shortfalls disproportionately. Based on this, Figure 7 plots the divergence between low and high-frequency poverty measures for comparison between the Poverty gap, which is distributionally insensitive to the Squared Poverty gap and the Watts Index, which are distributionally sensitive indicators.

Figure 7: (Unweighted) Implied yearly expenditures, non-headcount measures



The analysis is restricted to the sub-sample of households with at least four full years of monthly data, as described in Section 3.3. For each household, three quantities are computed: the average

of the annual average monthly expenditure across the study period, the average value of the household high frequency poverty line across the study period, and the average value of each high-frequency poverty indicator (poverty gap, squared poverty gap, and Watts index) across the study period.

Let $\bar{c}_i$ denote the average of the annual average monthly expenditure of household i across the study period, and let $\bar{PL}_i$ denote the average household poverty line across the same period. Further, let $\bar{PG}_i^{hf}$, $\bar{SPG}_i^{hf}$, and $\bar{W}_i^{hf}$ denote the average high-frequency poverty gap, squared poverty gap, and Watts index for household i respectively, each averaged across all months and years in the study period. The primary operation is to translate each household's average high-frequency poverty measure into the level of expenditure that would generate an equivalent poverty value under the low-frequency measurement framework. This is referred to as implied expenditure, denoted by $\tilde{c}_i$. Solving for $\tilde{c}_i$in each case:

For the poverty gap:

$$\bar{PG}_i^{hf} = \left(\frac{\bar{PL}_i - \tilde{c}_i}{\bar{PL}_i}\right) \quad \ldots (24)$$
$$\Longrightarrow \tilde{c}_i = \bar{PL}_i\left(1 - \bar{PG}_i^{hf}\right)$$

For the squared poverty gap:

$$\bar{SPG}_i^{hf} = \left(\frac{\bar{PL}_i - \tilde{c}_i}{\bar{PL}_i}\right)^2 \quad \ldots (25)$$
$$\Longrightarrow \tilde{c}_i = \bar{PL}_i\left(1 - \sqrt{\bar{SPG}_i^{hf}}\right)$$

For the Watts Index:

$$\bar{W}_i^{hf} = \ln\left(\frac{\bar{PL}_i}{\tilde{c}_i}\right) \quad \ldots (26)$$
$$\Longrightarrow \tilde{c}_i = \bar{PL}_i \,.\, e^{-\bar{W}_i^{hf}}$$

Finally, the ratio of implied expenditure to actual expenditure is computed for each of the three indicators as:

$$\frac{\tilde{c}_i}{\bar{c}_i} \quad \ldots (27)$$

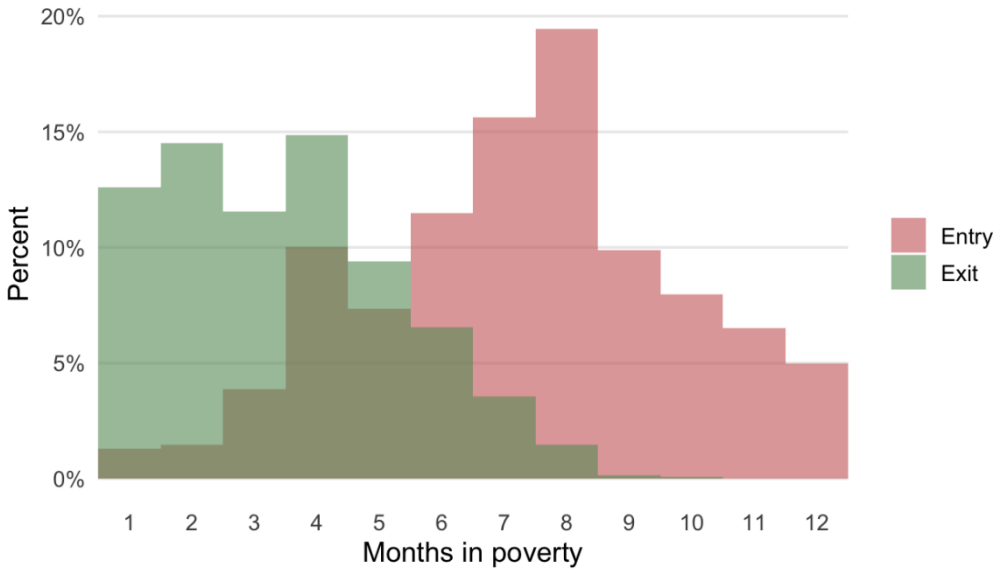
Based on this, Figure 7 plots the ratio on the vertical axis against normalised expenditure on the horizontal axis. Here, a value of less than 1 indicates that the expenditure level implied by the high-frequency poverty measure is lower than the household's actual average expenditure. This is true for all three poverty indicators across the entire range of normalised expenditure. This implies that when poverty is measured at the monthly level, it registers as deeper than when the same household's expenditure is averaged over the year and compared against an annual threshold.

However, it should be noted that the ratio is not constant across the distribution. The curve reaches its minimum at approximately 7.5 times the poverty line, meaning that the divergence between implied and actual expenditure is largest for households at this point. These are households whose average expenditure is well above the poverty line and who would not typically be considered vulnerable under any conventional measure of poverty. Finally, all three curves appear to overlap closely in the figure, and this is partly due to scaling constraints. In actuality, the precise values of the three indicators differ slightly before converging as the distance from the poverty line increases. Regardless, their behaviour across the distribution is consistent, and the above interpretation applies to all three.

4.2.4. Poverty Entry and Exit

So far, the analysis has established that poverty is more widespread than annual measures suggest, and that households classified as non-poor on an annual basis are nonetheless exposed to months of poverty within the year. Naturally, the next step is to examine what happens on the knife-edge of annual poverty classification. To this effect, figure 8 plots the distribution of months spent in high-frequency poverty separately for households that transition into and out of poverty between consecutive years.

Figure 8: (Unweighted) Poverty Entry and Exit



For each household, low-frequency poverty status is determined in each year and compared to the preceding year. A household is classified as 'entry' in a given year if it was not poor in the previous year but is poor in the current year, and as an 'exit' if it was poor in the previous year but not in the current year. Since the lagged poverty status cannot be computed for 2018, the analysis covers the period 2019 to 2023. The analysis is restricted to the sub-sample of households with at least four full years of complete monthly data, as described in Section 3.3. For each household-year observation classified as either entry or exit, the total number of months spent in high-frequency poverty in that year is computed, and the distribution of these counts is plotted separately for the two groups[12].

[12] The figure plots year-on-year transitions in annual poverty status, not households. A household is observed for at least four years and can transition more than once, so it may appear several times, as both an entry and an exit. Of the 46,382 households in the sub-sample, only 2,717 households ever cross the annual poverty line. This implies that a household that is poor in every year records no entry. Likewise, a household that is non-poor in every year registers

As evident in figure 8, among households classified as having exited poverty, the distribution of months in poverty is concentrated at the lower end, with the mode at 2 and 4 months. The majority of exiting households spend fewer than 5 months in poverty in the year of exit. This is as expected, since low frequency non-poor tend to spend relatively few months below the monthly poverty threshold. However, the distribution does not reach zero since a small share of exiting households still experience 3 to 6 months in poverty, indicating that exit from poverty as measured annually does not imply a complete absence of monthly deprivation.

Among households classified as having entered poverty, the distribution is concentrated at the higher end, with the mode at 8 months and a leftward tail. This implies that households entering poverty in a given year tend to spend more than half of that year below the poverty threshold. Therefore, it becomes clear that transitions into and out of poverty as measured annually are not discrete events. These transitions occur across a spectrum of uncertain experiences that the low-frequency framework does not capture.

Figure 8 helps reinforce the third result of the high-frequency framework: transitions in and out of poverty are not as clear-cut as implied by annual measures. A household that has exited poverty may still spend several months of the year below the poverty line. Similarly, a household that has entered poverty may not spend the entire twelve months of the year staying in poverty. The high-frequency framework makes these degrees of deprivation visible in a way that the annual measures cannot.

4.3. Testing consumption smoothing

The analysis in Section 4.2 established that poverty is more widespread than what annual measures suggest, and that households' within-year experience of deprivation varies considerably with their average living standard. A cross-cutting idea through this analysis is the role of consumption smoothing, which is understood as the degree to which households are able to stabilise their consumption expenditure in the face of uncertainties. While Section 3.4 provided preliminary evidence that households do smooth consumption, given that expenditure registers lower variability than income, this section tests this hypothesis more closely using two complementary approaches. The first is a fixed effects regression model that estimates the relationship between income and consumption across five model specifications, using education level as a proxy for socioeconomic status. The second is a paired t-test that directly compares three poverty headcounts, namely, income-based high-frequency, expenditure-based high-frequency, and expenditure-based low-frequency, to determine whether households smooth consumption perfectly, not at all, or imperfectly.

4.3.1. Fixed Effects Model

4.3.1.1. Hypothesis

The underlying hypothesis is that if households were perfectly smoothing their consumption, changes in monthly income would have no bearing on monthly expenditure after controlling for fixed effects and other covariates. Formally:

no exit. The figure counts only transitions with at least one month of high-frequency poverty in the transition year. This leaves 4,320 such events (not individual households), comprising of 2,393 entries and 1,927 exits.

H_0: Household consumption is perfectly smoothed with respect to income shocks, i.e., changes in household income have no effect on household expenditure after controlling for fixed effects and other covariates.

H_1: Household consumption is not perfectly smoothed, i.e., changes in household income are positively associated with changes in household expenditure, even after accounting for fixed effects and other covariates.

4.3.1.2. Model Specifications

Table 4 provides a summary of the five specifications considered for regression.

Table 4: Model Specifications

Specification	Equation	Description
M1	$c_{it} = \alpha + \beta y_{it} + \gamma(y_{it} \times Edu_{it}) + \varepsilon_{it}$	Baseline: measures income-consumption relationship, by education group.
M2	$c_{it} = \alpha + \beta y_{it} + \gamma(y_{it} \times Edu_{it}) + \lambda_t + \varepsilon_{it}$	Year-month fixed effects: Controls for time shocks common to all households.
M3	$c_{it} = \alpha + \beta y_{it} + \gamma(y_{it} \times Edu_{it}) + \mu_i + \lambda_t + \varepsilon_{it}$	Household, year-month fixed effects: Controls for time-invariant household traits and time shocks.
M4	$c_{it} = \alpha + \beta y_{it} + \gamma(y_{it} \times Edu_{it}) + \delta lagc_{it} + \mu_i + \lambda_t + \varepsilon_{it}$	M3 + lagged mean expenditure: M3 + controls for household's average past-year expenditure
M5	$c_{it} = \alpha + \beta y_{it} + \gamma(y_{it} \times Edu_{it}) + \mu_i + \upsilon_{vt} + \varepsilon_{it}$	Household, village-year-month fixed effects: Controls for village or community level factors.

The dependent variable in all five specifications is adjusted monthly household expenditure[13] given by c_{it}. The primary regressor of interest is monthly household income represented by y_{it}. Under perfect consumption smoothing, the coefficient of y_{it} should be statistically indistinguishable from zero. A positive and significant coefficient would instead imply that households are adjusting their expenditure in response to income fluctuations, thereby constituting evidence of imperfect smoothing.

Education level is included as a proxy for socioeconomic status, since it is associated with access to higher-paying and more stable employment and with greater financial literacy (David Card, 1999; Lusardi & Mitchell, 2014). Additionally, since formal education is typically completed in early adulthood, it is relatively fixed over time and straightforward to collect in a survey setting, making

[13] For some variables, such as monthly expenditure, CMIE collects information from the household for the month as well as for the week, assuming that weekly expenditures may be more reliably reported than monthly expenditures. Then, it applies an adjustment factor to account for discrepancies between the two and reports it as monthly 'adjusted' expenditure.

it a practical proxy for the household's capacity to manage uncertainty(Galobardes, et al., 2006). The education classification used here draws on CMIE's household-level education grouping, which categorises households based on the distribution of education levels among members aged 25 and above, rather than the education of the household head alone. CMIE's rationale for this approach is that the household head is often less educated than younger members, and relying solely on the head's education level risks misclassifying households. The 11 education groups provided by CMIE are collapsed into three: low, comprising households where all or some members are illiterate; medium, comprising households where at least some members have passed matriculation; and high, comprising households where at least some members are graduates. The higher education group serves as the baseline category in all specifications. In other words, the coefficients of education level are interpreted relative to the higher education baseline.

An interaction term between income and education level ($\gamma(y_{it} \times Edu_{it})$), is included to test the joint effect of income and education on expenditure. This captures how households with different levels of education smooth consumption in response to monthly changes in income, over and above the baseline effect of income alone.

The five model specifications are as follows:

(i) M1 is the baseline, estimating the income-consumption relationship by education group without any fixed effects.

(ii) M2 adds year-month fixed effects $\lambda_t$to control for aggregate time shocks common to all households, such as seasonal price movements or macroeconomic fluctuations.

(iii) M3 adds household fixed effects $\mu_i$to M2, controlling for time-invariant household characteristics that may confound the income-expenditure relationship

(iv) M4 builds on M3 by incorporating the lagged mean expenditure variable $lagc_{it}$, defined as the average household expenditure over the preceding 12 months, to test whether past expenditure is a reliable predictor of current expenditure.

(v) M5 replaces the year-month fixed effects in M3 with village-year-month fixed effects υ_{vt}, which absorb local community-level factors that vary over time.

It is worth noting that M4 and M5 are best understood as two distinct extensions of M3 rather than a sequential progression. M4 isolates the role of past expenditure within the household-time framework of M3, while M5 replaces the time fixed effect with a more granular community-level equivalent. Including all three types of fixed effects simultaneously (household, time, and village-year-month) would risk introducing multicollinearity among the fixed effects terms, which could inflate standard errors, destabilise coefficient estimates, and distort the true magnitude and direction of the effects of interest.

4.3.1.3. Regression Results

A summary of the regression results is presented in Table 5.

Table 5: Regression Results

	M1	M2	M3	M4	M5
Constant	9,742.4***				
	(119.3)				
Monthly Income	0.1919***	0.1872***	0.0880***	0.0546***	0.0860***
	(0.0047)	(0.0047)	(0.0037)	(0.0025)	(0.0041)

Education level: low	-1,586.2***	-627.1***	257.9***	187.2***	22.80
	(156.7)	(146.9)	(78.05)	(55.26)	(75.41)
Education level: medium	2,337.3***	2,305.3***	696.7***	391.6**	455.2*
	(176.3)	(174.1)	(154.0)	(139.1)	(187.2)
Income x Education level: low	-0.0987***	-0.1044***	-0.0432***	-0.0282***	-0.0247***
	(0.0084)	(0.0077)	(0.0044)	(0.0030)	(0.0043)
Income x Education level: medium	-0.0980***	-0.1004***	-0.0399***	-0.0180**	-0.0271**
	(0.0076)	(0.0075)	(0.0072)	(0.0066)	(0.0089)
Mean lagged Expenditure (12 months)				0.9983***	
				(0.0127)	
Fixed-Effects:	---	---	---	---	---
Month x year	No	Yes	Yes	Yes	No
Household	No	No	Yes	Yes	Yes
Village x year x month	No	No	No	No	Yes
S.E.: Clustered	by: Household	by: Household	by: Household	by: Household	by: Household
Observations	30,30,862	30,30,862	30,30,862	25,21,103[14]	30,30,862
R2	0.02259	0.02639	0.06414	0.12784	0.50171
Within R2	--	0.02091	0.00298	0.07101	0.00433

Four observations can be summarised from the results in Table 5. First, the coefficient of monthly income is positive and statistically significant across all five specifications, confirming that households do not perfectly smooth their consumption in response to income shocks. Second, the magnitude of this coefficient diminishes progressively as household and time fixed effects are introduced, suggesting that a portion of the raw income-expenditure correlation is attributable to time-invariant household characteristics and aggregate time shocks rather than only within-household consumption sensitivity. Third, the interaction coefficients on income and education level are negative for both the low- and medium-education groups across all specifications, implying that more-educated households are better able to absorb income shocks without adjusting their expenditure. Fourth, the coefficient on lagged mean expenditure in M4 is close to one and highly significant, confirming that past expenditure is a strong predictor of current spending. Taken together, these observations imply that consumption smoothing is not a uniform capacity distributed equally across households. This is because it is contingent on socioeconomic status, with more educated households being better insulated from the effect of uncertainty on expenditure. Details are presented below.

First, the coefficient of monthly income is positive and statistically significant at the 1 percent level across all five specifications, providing strong and consistent evidence against the null hypothesis of perfect consumption smoothing. In M1, a one rupee increase in monthly income is associated with a 0.191 rupee increase in monthly expenditure. This effect diminishes progressively as additional controls are introduced, falling to 0.187 in M2, 0.088 in M3, 0.054 in M4, and 0.086 in M5. This suggests that a portion of the raw income-expenditure correlation is attributable to time-

[14] The reduction in observations in M4 from 30,30,862 to 25,21,103 is a mechanical consequence of the lagged variable construction: since the lag is computed over the preceding 12 months, the first year of data for each household cannot be used, reducing the effective sample by one full year per household.

invariant household characteristics and aggregate time shocks rather than genuine within-household consumption sensitivity to income. Nevertheless, the coefficient remains positive and significant throughout, confirming that households do not perfectly smooth their consumption.

Second, lower-income and lower-education households engage in more active and deliberate consumption smoothing precisely because their income volatility is higher. Simultaneously, for such households, the consequences of failing to smooth are more severe. First looking at M1 and M2, it can be seen that households with low education levels exhibit negative coefficients, indicating worse consumption smoothing relative to highly educated households. However, from M3 onwards, after accounting for household fixed effects, the sign reverses to positive for low education households, while medium education households display positive coefficients throughout. While this may appear counterintuitive, it is consistent with the earlier finding in Section 3.4 that households in the 25th income percentile, arguably a group that overlaps with the low education category, exhibit a lower coefficient of variation in expenditure relative to income than their counterparts in the 75th percentile.

Third, the interaction coefficients of income and education level are negative for both low and medium education groups across all five specifications, and statistically significant in most cases. Since these are measured relative to the high education baseline, negative coefficients imply that less and moderately educated households smooth a smaller fraction of any given income change into expenditure compared to highly educated households. In other words, more educated households are better able to absorb income shocks without adjusting their expenditure, while less educated households pass a larger share of income fluctuations through to consumption.

Lastly, the coefficient of lagged mean expenditure in M4 is 0.998 and statistically significant at the 1 percent level, implying that past expenditure is an almost one-for-one predictor of current expenditure. This is consistent with the hypothesis that household spending is highly persistent and that consumption habits formed over the preceding year are a strong determinant of current period expenditure.

4.3.2. Test of Equality

4.3.2.1. Hypothesis

The regression framework in the previous section confirmed that households do not perfectly smooth their consumption. However, it does not directly situate the degree of smoothing relative to the two theoretical benchmarks, segregated into perfect smoothing and no smoothing. The test of equality addresses this by comparing three poverty headcounts for each education group: income-based high-frequency poverty $P(m,y)^{hf}$, expenditure-based high-frequency poverty $P(m,c)^{hf}$, and expenditure-based low-frequency poverty $P(a,c)^{lf}$. The underlying logic here is that if households smooth consumption perfectly, monthly expenditure will be insulated from income fluctuations and high-frequency expenditure-based poverty will converge to its low-frequency counterpart. Formally, it can be hypothesised as:

$$P(m,c)^{hf} = P(a,c)^{lf} \quad \dots (28)$$

On the other hand, If households do not smooth consumption at all, expenditure will mirror income exactly, and income-based and expenditure-based high-frequency poverty will be identical. This can be represented as:

$$P(m,c)^{hf} = P(m,y)^{hf} \quad \ldots (29)$$

Finally, if households smooth consumption imperfectly, expenditure-based high-frequency poverty will lie strictly between these two benchmarks:

$$P(a,c)^{lf} \leq P(m,c)^{hf} \leq P(m,y)^{hf} \quad \ldots (30)$$

4.3.2.2. Testing the hypothesis

A paired t-test is used to evaluate each of the two equality conditions. The paired t-test is appropriate here because the three headcounts are computed from the same pool of households. In other words, the comparison is not between two distinct groups but between two different measurements of the same underlying concept applied to the same sample. Table 6 summarises the results of the paired t-test.

Table 6: Paired t-test

	Low education	Medium education	High education
P (month, income)	0.171	0.073	0.068
P (month, expenditure)	0.188	0.027	0.045
P (annual, expenditures)	0.109	0.008	0.021
Test of equality (p):			
$P(m,c)^{hf} = P(m,y)^{hf}$	0	~0	0
$P(m,c)^{hf} = P(a,c)^{lf}$	0	~0	0

Both paired t-tests result in rejections at conventional significance levels across all education groups. The p-value for the test $P(m,c)^{hf} = P(m,y)^{hf}$ is approximately zero for all three groups, allowing rejection of the hypothesis that households do not smooth consumption at all. The p-value for the test $P(m,c)^{hf} = P(a,c)^{lf}$ is similarly approximately zero, allowing rejection of the hypothesis that households smooth consumption perfectly. Together, these results confirm that households largely align with the intermediate case: they do smooth consumption in the face of income volatility, but they do so imperfectly, and the degree of imperfection varies with their education level. For the medium- and high-education groups, the point estimates satisfy the ordering in equation (30), confirming the intermediate case: these households smooth consumption in the face of income volatility, but do so imperfectly. The low-education group is the exception. Here the upper bound of equation (30) is violated i.e., the value of expenditure-based high-frequency poverty $P(m,c)^{hf}$ exceeds not only the expenditure-based low-frequency poverty $P(a,c)^{lf}$, as expected, but also income-based high-frequency poverty $P(m,y)^{hf}$. In other

words, it is the opposite of what consumption smoothing would produce, i.e., monthly poverty measured on expenditure is higher than monthly poverty measured on income. This suggests that the least-educated households are unable to effectively draw on their coping strategies to stabilise their consumption. As a result, a drop in income translates to a drop in consumption. Nonetheless, this is consistent with the fixed-effects results in Section 4.3.1, where these households transmitthe largest share of income fluctuations to expenditure. This constitutes evidence for justifying the second result of the high-frequency framework: households' ability to smooth consumption varies with socioeconomic status, and this variation directly shapes their exposure to poverty throughout the year.

4.4. Policy Experiment

The analysis in Sections 4.2 established that poverty is more widespread and dynamic than annual measures suggest, and that households experience deprivation in an uneven manner across months within a year. This section tests the implications of this observation for poverty-alleviation programs using a hypothetical cash-transfer experiment.

The experiment considers a fixed annual transfer of INR 12,000 per household, equivalent to INR 1,000 per month. While the amount is chosen arbitrarily, it is broadly indicative of the magnitude of existing government cash transfer schemes in India[15]. . It should be noted that the experiment is restricted to low-frequency poor households, i.e., those whose annual average monthly expenditure falls below the annual household poverty line. This is an obvious choice given that any cash transfer program would target only poor households. The experiment consists of four allocation strategies, which are evaluated against a no-transfer baseline:

(i) Baseline: No transfer
(ii) Strategy 1: INR 1,000 per month distributed evenly across all 12 months
(iii) Strategy 2: INR 2,000 per month distributed across the 6 poorest months.
(iv) Strategy 3: INR 4,000 per month distributed across the 3 poorest months.
(v) Strategy 4: The full INR 12,000 is distributed exclusively across months in which the household is already in high-frequency poverty, with the per-month amount varying inversely with the number of poor months.

For strategies 2 and 3, the poorest months are identified at the household level by ranking each household's monthly expenditure from lowest to highest within each year and applying the transfer to the bottom 6 and 3 months, respectively. For strategy 4, the transfer is applied only to months in which the household's expenditure falls below the monthly poverty line, with the total annual amount divided equally across those months. The impact of each strategy is evaluated using two high-frequency poverty measures, as shown in table 7.

Table 7: Impact on Poverty Indicators

Strategy	Headcount	Watts
Baseline	0.040	0.009

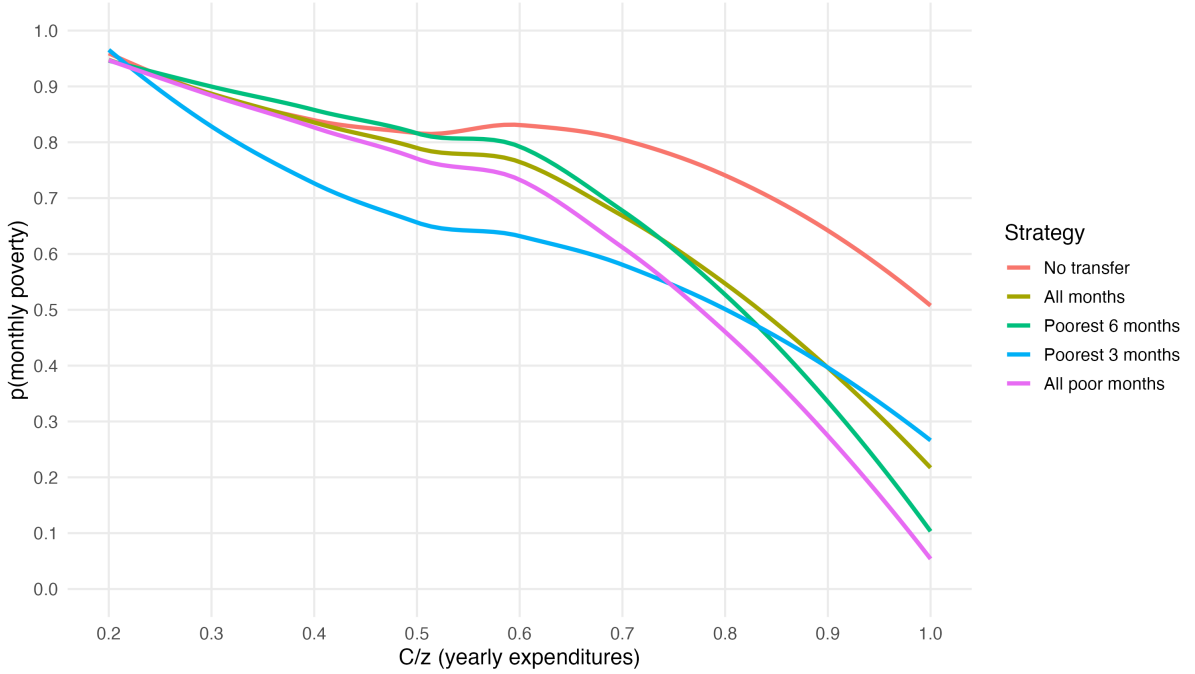
[15] In Chhattisgarh, Mahtari Vandan Yojana provides INR 1,000 as DBT to eligible women in the household. Similarly, in Tamil Nadu, the Kalaignar Magalir Urimai Thodai scheme provdes INR 1000 to women heads of eligible households. Other states, such as Madhya Pradesh, provide higher benefit amounts of INR 1500 to eligible women through the Chief Miniter's Ladli behna Yojana. In Jharkhand, the benefit amount under Mukhyamantri Maiya Samman Yojana now sits at INR 2500. States regularly undertake periodic revisions of their cash transfer schemes.

S1	0.018	0.003
S2	0.011	0.002
S3	0.014	0.003
S4	0.008	0.001

This confirms that all four strategies reduce both the headcount and the Watts index relative to the baseline. The baseline high-frequency headcount of 0.040 falls to 0.018 under S1, 0.011 under S2, 0.014 under S3 and 0.008 under S4. A similar pattern holds for the Watts index, which falls from 0.009 at baseline to 0.003, 0.002, 0.003, and 0.001 from S1 through S4, respectively. Strategy 4 performs best on both measures, confirming that concentrating the transfer precisely on months of observed poverty is the most effective allocation design. Among the remaining three strategies, S2 outperforms both S1 and S3 on both indicators.

Figure 9 illustrates this pattern for poverty headcount across the distribution of normalised annual expenditure.

Figure 9: (Unweighted) Policy Experiment- Headcount



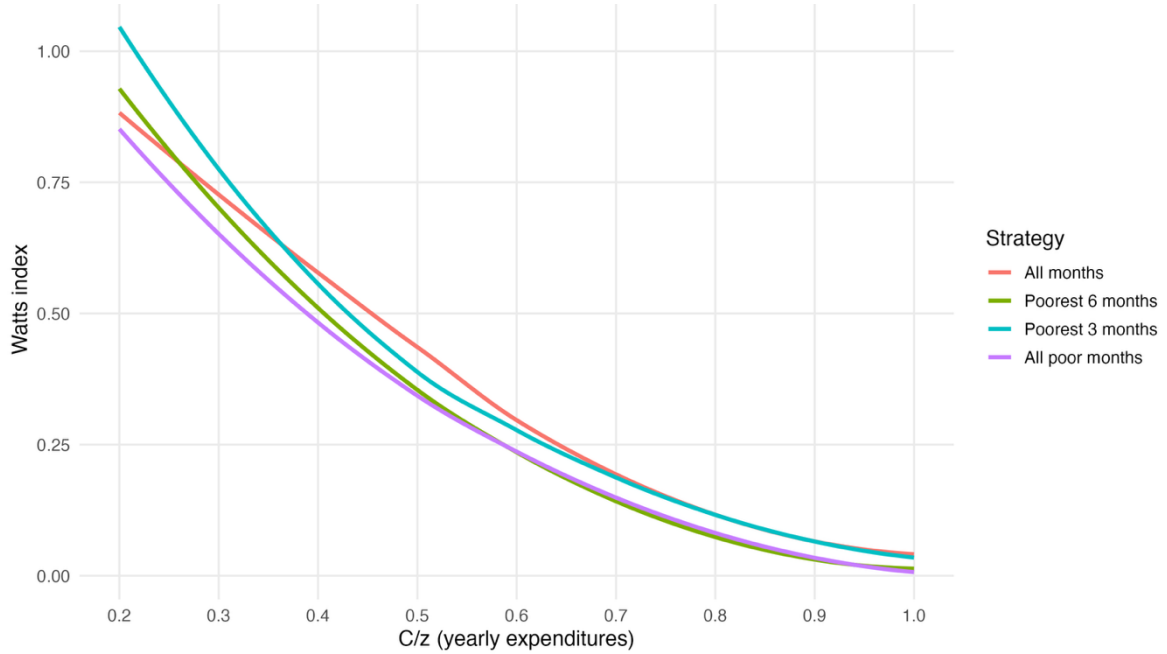
All four transfer strategies as well as the baseline exhibit a downward sloping trajectory. This implies that as normalised yearly expenditures move closer to the poverty line (given by 1 on the x axis), households' exposure to poverty, expressed as a proportion of months in poverty, reduces. The figure does not plot the distributions beyond 1 on the x-axis as that would simply be the normalised expenditure of households who have crossed the poverty threshold and therefore, are no longer considered poor. Notably, all four transfer strategies lie below the no-transfer baseline throughout, confirming that a direct cash transfer reduces monthly poverty for all poor households. The no-transfer curve is notably elevated across the middle of the distribution, particularly between 0.5 and 0.7 times the poverty line, and declines only gradually toward 1.0, suggesting that without targeted support, households in this range spend a disproportionately large share of their months in poverty despite being relatively close to the annual poverty threshold.

As seen in the figure, there is no one size (strategy) that fits all (best). The most effective way to provide financial assistance for poverty alleviation depends on how far below the poverty line a household already sits. For households farthest from the poverty line, exposure to poverty is highest, i.e., these households are poor in almost every month of the year. Therefore, targeting the transfer to the poorest 3 months (S3) delivers the highest relief from deprivation. Notably, up until 0.5 times the poverty line, strategies S1, S2, and S4 do not show any significant deviation from the baseline of no transfers. However, as the households' normalised expenditure moves closer to the poverty line, strategies S1, S2, and S4 become more effective in reducing their exposure to poverty, as compared to S3, which now falls short.

The implication here is that effective policy design must be sensitive to knowledge about not just who is poor, but how poor they are. A single nationwide program, irrespective of whether it spreads assistance evenly throughout the year or targets specific months within the year, will yield differential outcomes for different households, whose experience of poverty varies significantly.

Figure 10 demonstrates this pattern for the Watts Index.

Figure 10: (Unweighted) Policy Experiment- Watts



S4 and S2 register the lowest values across the distribution. It is worth noting that S3, despite concentrating the transfer into fewer months at a higher per-month value, performs worse than S2. This is particularly significant because the Watts index is distributionally sensitive and therefore penalises deeper shortfalls more heavily. This implies that concentrating the transfer into just 3 months does not adequately address the depth of deprivation in the remaining months. As a result, this inadequacy is captured by the Watts index in a way that the headcount cannot. S1, the uniform distribution, consistently registers the highest values among the transfer strategies, confirming that evenly spreading the transfer across all months, including months in which the household is not poor, is the least efficient use of a direct cash transfer.

These results illustrate the practical value of the high-frequency framework. A conventional annual poverty measure, by averaging out within-year fluctuations, provides no basis for distinguishing between these four strategies since all four involve the same total annual transfer and would appear identical under a low-frequency lens. It is only by measuring poverty at the monthly level that the temporal effect becomes visible, and it is precisely this visibility that can enable more effective targeting of poverty-alleviation programs.

5. Discussion

This section compares the results of this paper with those of the original study, discusses where and why they diverge, and reflects on what the modifications in this paper contribute to the high-frequency framework.

5.1. Comparing with M&M's prior work on high-frequency poverty

The first point of comparison is the magnitude of the gap between high and low frequency poverty. M&M find that the high-frequency headcount is 26 percent larger than the conventional headcount (37 percent at high frequency poverty against 29 percent at low frequency poverty) in their sample of 945 rural Indian households between 2010 and 2015. In this paper, the high-frequency headcount of 4 percent is twice the low-frequency headcount of 2 percent.

Several factors likely contribute to this difference. First, the CPHS covers a fundamentally different population from the VDSA. M&M use the VDSA survey, which covers a balanced panel of 945 low-income households from rural India, all of whom are, by construction, close to or below the poverty line. The CPHS, by contrast, is a nationally representative survey covering both rural and urban India across all income levels. The 2 percent low-frequency poverty rate in this paper reflects the fact that the majority of the sample sits comfortably above the poverty line. Notably, this figure matches estimates provided by the World Bank (2025), according to which, extreme poverty in India (those living on less than $2.15 per day) stood at 2.3 percent as of 2022-23. Therefore, the two figures are not directly comparable and the differences in results between this reproduction and M&M's original should be understood in this light. Additionally, there are minor differences in the composition of variables between the two studies that warrant acknowledgement. The VDSA survey collects detailed monthly data on household income and expenditure across a narrow set of agricultural communities in rural India. The CPHS, by contrast, collects a broader set of socioeconomic indicators across a nationally representative sample. The specific variables used to construct income and expenditure measures differ in their scope and definition. In particular, the expenditure variable used in this analysis, i.e., the total adjusted household expenditure from CPHS, includes a wider range of consumption categories than the expenditure data in the VDSA, which focuses more on food and basic consumption items. Similarly, the income variable in CPHS captures total household income from all sources, which may be measured differently from the income data in the VDSA. These differences in variable construction mean that the poverty indicators computed here are not perfectly comparable to those in M&M's original study, and some of the divergence in results may reflect differences in what the two datasets are measuring rather than genuine differences in the poverty experience of the two samples.

The second factor that is likely contributing to this difference is that the poverty line used in this paper is dynamic, since it varies by state, region, and month. Meanwhile, M&M use a static rural poverty line. The dynamic poverty line introduces additional month-to-month variation in the threshold against which household expenditure is compared, potentially classifying some households as poor in certain months that a static line would not. This is the primary

methodological contribution of this paper, and the larger gap between high and low frequency poverty measures in this exercise partly reflects the additional sensitivity introduced by the variation in the poverty line across states and regions. Third, the study periods differ in ways that are consequential. M&M cover 2010 to 2015, while this analysis covers 2018 to 2023, which includes the COVID-19 pandemic. As noted in Section 2.1, the lockdowns of 2020 and 2021 represent a structural break in household expenditure patterns that is unlikely to reflect routine within-year volatility. The spike in monthly poverty observed in mid-2020 in Figures 4 and 5 suggests that the COVID-19 period is inflating the high-frequency headcount.

The second point of comparison concerns the contribution of non-poor households to monthly poverty. M&M find that 47 percent of households labelled as non-poor annually experience at least one month of poverty, and that 35 percent of all months in poverty are experienced by non-poor households. In this reproduction, Figure 5 shows that approximately 60 percent of all household-months in poverty are contributed by low-frequency non-poor households, which is a larger share than what M&M find. As noted in Section 4.2.1, however, this is partly a mechanical consequence of sample composition. With 98 percent of the weighted sample classified as non-poor, even a very small monthly poverty rate among non-poor households produces a large aggregate contribution to the monthly poverty count. Nonetheless, the finding that non-poor households contribute meaningfully to the aggregate experience of monthly poverty is consistent with M&M.

The third point of comparison is the consumption smoothing results. M&M find evidence consistent with substantial but imperfect consumption smoothing, reporting that the ratio of the median coefficient of variation of consumption (0.25) to the median coefficient of variation of income (0.87) is approximately 27 percent for poor households. This indicates that households smooth out about a significant proportion of income volatility through expenditure smoothing[16]. The fixed effects regression in Section 4.3.1 is consistent with this finding. This paper reports that the coefficient on monthly income is positive and significant across all five model specifications, but diminishes as fixed effects are introduced, confirming that households do smooth consumption but imperfectly. The test of equality in Section 4.3.2 further confirms that, for households with medium and high levels of education, expenditure-based high-frequency poverty lies strictly between income-based high-frequency poverty and expenditure-based low-frequency poverty across all three education groups, satisfying the inequality $P(a,c)^{lf} \leq P(m,c)^{hf} \leq P(m,y)^{hf}$. For the least-educated households the inequality does not hold, indicating that they pass income shocks to consumption rather than buffering them through coping mechanism. Nonetheless, the finding that the ability to absorb income shocks without adjusting expenditure rises with education is consistent with M&M's result that exposure to poverty is greatest for households with the least education.

The fourth point of comparison is the cash transfer experiment. M&M demonstrate that targeting transfers to the most challenging months is more cost-effective than distributing them evenly across the year. The policy experiment in Section 4.4 reproduces the essence of this result. Strategy 4, which concentrates the full annual transfer exclusively on months of observed high-frequency poverty, achieves the largest reduction in both the headcount and the Watts index. Strategy 2, which targets the six poorest months, outperforms both Strategy 1 (uniform distribution) and Strategy 3 (three poorest months) on the Watts index, consistent with M&M's finding that targeting the 6 poorest months is better than concentrating support in only the three poorest months on on

[16] If there were no smoothing, the variability of monthly consumption would be identical to variability of monthly income, and the ratio of their coefficients of variation would be 100 percent. Conversely, if households could smooth consumption perfectly, the ratio would be zero.

distributionally sensitive measures. The result that uniform distribution is the least effective strategy is also consistent with M&M and reinforces the practical value of the high-frequency framework for designing anti-poverty interventions. However, as the discussion in Section 4.3. highlights, the effect of a cash transfer is heterogenous and contingent on how far below the poverty line the household sits.

5.2. Modifications to the methodology

The methodological modifications applied to create a dynamic poverty line in this paper are significant, since they expand M&M's framework in meaningful ways. M&M use a static monthly poverty line that is identical across all households and all periods. This simplification, while analytically convenient, does not account for the fact that the cost of meeting basic needs varies substantially across Indian states and between rural and urban areas, and that these costs change over time with inflation. By constructing a poverty line that varies by state, region, and month using the Tendulkar Committee's baseline and MoSPI's CPI series, this paper allows the poverty threshold to reflect the true purchasing power of households at a given point in space and time.

The second modification concerns the data. M&M draw on the Village Dynamics in South Asia (VDSA) survey, while this paper instead uses the Consumer Pyramids Household Survey (CPHS). This change offers two advantages. First, the CPHS covers a much wider range of households across the income distribution. This makes it possible to examine within-year poverty among households with widely different socioeconomic status. Second, CPHS observes households over a longer and more recent period, allowing their movements in and out of poverty to be tracked across six years rather than five. This advantage is partly offset by substantial attrition, which imposes a restriction on the analyses to a sub-sample of households observed for at least four years. As mentioned earlier, the longer period also includes the COVID-19 pandemic. This complicates direct comparison with M&M's work, since the nationwide lockdowns of 2020 and 2021 represent an aggregate shock rather than the routine within-year volatility the framework is designed to capture but it also shows how the gap between monthly and annual poverty behaves under such a shock.

5.3. Limitations

Several limitations of this analysis are worth acknowledging.

First, the COVID-19 pandemic represents a structural break in the study period, complicating the interpretation of the results. The spike in monthly poverty in mid-2020 and the associated volatility in the proportion of monthly poverty attributable to non-poor households are unlikely to reflect the routine within-year dynamics that the high-frequency framework is designed to capture. Excluding 2020 and 2021 from the analysis would likely produce results closer to M&M's findings, and future work should consider additional robustness checks.

Second, the sub-sample used for longitudinal analyses, which includes households observed for at least 48 consecutive months, may introduce selection bias. As noted in Section 3.3, households that persist in the survey for longer periods are likely to be less mobile and more economically stable than those who attrit, which may cause the sub-sample to underrepresent precisely the kind of highly volatile, transient households whose poverty experience the high-frequency framework is designed to capture.

Third, the Tendulkar Committee's poverty lines, which serve as the baseline for the dynamic poverty line construction, which itself is based on prices in 2011-12. While the CPI adjustment

carries these baselines forward to the study period, the basket of goods underlying the Tendulkar lines may no longer accurately reflect the consumption needs of Indian households in 2018-2023, particularly given structural changes in consumption patterns over the last two decades. Accordingly, following the release of the new Purchasing Power Parities (PPPs) by the International Comparison Program in May 2024, the World Bank has revised the International Poverty Line from $ 2.15 (2017 PPP) to $ 3.0 (2021 PPP).

Notwithstanding these limitations, the results of this reproduction are broadly consistent with M&M's findings across all three results: (i) Poverty is more widespread than annual measures suggest, (ii) consumption smoothing is imperfect and varies with socioeconomic status, (iii) and transitions into and out of poverty are not as clear-cut as implied by annual measures. The modifications introduced in this paper add analytical value to M&M's framework without undermining its core insights, and they open avenues for further work and extension using other nationally representative datasets from India and elsewhere.

6. Conclusion

Poverty, as experienced by low-income households, is not a static condition that can be fully captured by a single annual number. It is dynamic, spatially uneven, and shaped by households' unequal capacity to smooth consumption in the face of uncertainty. The high-frequency framework introduced by Merfeld and Morduch and extended in this paper offers a way to measure poverty that is closer to how it is actually experienced.

This reproduction confirms that the gap between what annual measures tell us and what monthly measures reveal is measurable and meaningful. Moreover, it shows that poverty varies across regions in ways that a purely temporal framework cannot identify. The spatial dimension of poverty introduced here adds a geographic layer to M&M's framework. The policy experiment further demonstrates that there's practical value in knowing when poverty is most acute, since it enables more effective targeting of anti-poverty interventions than annual measures can support.

This paper is offered as an invitation to further work. The framework can be extended to understand poverty for specific geographical regions, applied to a wider range of interventions, and deepened through finer spatial units and advanced spatial econometric techniques. Notably, one result from M&M's original study that this paper does not attempt to reproduce is the predictive power of high-frequency poverty indicators for anthropometric outcomes. Specifically, predictors of height-for-age of children would require linking the CPHS to external health datasets, such as the National Family Health Survey, which represents a meaningful avenue for future research.

Ultimately, the case for measuring poverty at higher frequency rests not on methodological elegance but on the recognition that the people whose lives poverty statistics are meant to describe do not experience deprivation as statistical averages. Getting poverty measurement right in its timing, its geography, and its sensitivity to the depth of deprivation is a necessary step toward designing policies that respond to poverty as it is actually lived, rather than as it is conventionally counted.

Bibliography

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

Mas, I., & Murthy, G. (2017, August). Money, Decisions, and Control. *Working Paper*. Consultative Group to Assist the Poor.

Booth, C. (1893). Life and labour of the people of London: First results of an inquiry based on the 1891 census. *Journal of the Royal Statistical Society*, 557-93.

Rowntree, B. S. (1901). *Poverty : a study of town life*. London: Macmillian.

Himmelfarb, G. (1985). *The Idea of Poverty: England in the Early Industrial Age*. New York.

Gillie, A. (1996). The Origin of the Poverty Line. *The Economic History Review, 49*(4), 715-730.

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

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

United Nations. (2025). *The Sustainable Development Goals Report 2025*. United Nations.

Chambers, R. (1983). *Rural Development: Putting the Last First*. Pearson Education Ltd.

Dercon, S., & Krishnan, P. (2000). *Vulnerability, seasonality and poverty in Ethiopia* (Vol. 36(6)). Journal of Development Studies.

Azevedo, J. P., & Seitz, W. (2016). *How persistent is poverty in the short run?* World Bank Blogs.

Deaton, A. (1991). Savings and Liquidity Constraints. *Econometrica, 59*(6), 1221-1248.

Fink, G., Jack, B. K., & Masiye, F. (2020). Seasonal Liquidity, Rural Labor Markets, and Agricultural Production. *American Economic Review, 110*(11), 3351-92.

Augenblick, N., Jack, B. K., Kaur, S., Masiye, F., & Swanson, N. (2023). Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Poverty. *Working Paper*. University of California, Berkeley.

Breza, E., Kaur, S., & Shamdasani, Y. (2021). Labour Rationing. *American Economic Review, 111*(10), 3184-3224.

Bryan, G., Chowdhury, S., & Mobarak, A. M. (2014). Underinvestment in Profitable Technology: The Case of Seasonal Migration in Bangladesh. *Econometrica, 82*(5), 1671-1748.

Stephen Devereux, R. S.-W. (2012). *Seasonality, Rural Livelihoods and Development*. Routledge.

Conceptual and Measurement Issues. (2001). In B. Bradbury, S. P. Jenkins, J. Micklewright, B. Bradbury, S. P. Jenkins, & J. Micklewright (Eds.), *The Dyamics of Child Poverty in Industrialised Countries*. Cambridge University Press.

Foster, J., Greer, J., & Thorbecke, E. (1984). A Class of Decomposable Poverty Measures. *Econometrica, 52*(3), 761-766.

Sen, A. (1976). Poverty: An Ordinal Approach to Measurement. *Econometrica, 44*(2), 219-231.

Mancini, G., & Vecchi, G. (2022). On the Construction of a Consumption Aggregate for Inequality and Poverty Analysis. Washington, D.C.: The World Bank Group.

Deaton, A. (1997). *The Analysis of Household SPolicyurveys: A Microeconomic Approach to Development* . The World Bank Group & The Johns Hopkins University Press.

Deaton, A., & Zaidi, S. (2002). Guidelines for Constructing Consumption Aggregates for Welfare Analysis. *LSMS Working Paper Number 135*. The World Bank Group.

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

World Bank. (2025). India: Trends in Poverty from 2011-12 to 2022-23. World Bank.

Card, D. (1999). The Causal Effects of Education on Earnings. In O. Ashenfelter, *Handbook of Labor Economics* (pp. 1801-1963).

Lusardi, A., & Mitchell, O. S. (2014). The Economic Importance of Financial Literacy: Theory and Evidence. *Journal of Economic Literature, 52*(1), 5-44.

Galobardes, B., Shaw, M., Lawlor, D. A., Lynch, J. W., & Smith, G. D. (2006). Indicators of socioeconomic position (part 1). *Journal of Epidemiology & Community Health, 60*(1), 7-12.

Merfeld, J., & Morduch, J. (2022). Poverty at Higher Frequency. NYU.
Merfeld, J., & Morduch, J. (2025). Poverty at Higher Frequency. NYU.
Planning Commission. (2013). *Press Note on Poverty Estimates, 2011-12*. New Delhi: Government of India.
Ministry of Statistics and Programme Implementation (MoSPI). (2024). Consumer Price Index (Rural, Urban), base 2012=100: State-wise monthly indices. Government of India.