Policy Brief

Measuring Poverty at Higher Frequency

Overview

Similar to most low and middle income countries, the measurement of poverty in India is understood as an annual affair. Traditional policy discourse documents it as an annual headcount metric that is typically described as the proportion of all households facing an income shortfall within a year. Consequently, all national poverty rates and anti-poverty policies are constructed and implemented, respectively, based on this understanding of poverty. However, such poverty statistics do not adequately describe the deprivations experienced by households. This is because low-income households face uncertainties of insufficiency, instability, and illiquidity in a manner that cannot be quantified into statistical averages. Moreover, households have developed coping mechanisms to address these uncertainties by smoothing their consumption, but they do so imperfectly and unevenly across the income distribution. Therefore, a more nuanced and grounded understanding of poverty is a necessary step toward designing policies that respond to poverty as it is actually lived, rather than how it is conventionally measured.

In the study based on Merfeld and Morduch's Poverty at Higher Frequency framework, we use the Consumer Pyramids Household Survey (CPHS) for January 2018 to December 2023, and find that (i) high-frequency poverty[1] is roughly twice low-frequency poverty[2], and that households classified as poor in a given year experience nine out of twelve months in poverty; (ii) the ability to smooth consumption varies with socioeconomic status, as more educated households better absorb income shocks without adjusting expenditure, while less educated households pass a larger share of income fluctuations through to consumption; and (iii) households above the poverty line remain exposed to deprivation, with exposure diminishing gradually and non-linearly with distance from the line and not becoming negligible until a household spends at roughly twice the poverty line. Ultimately, the contribution of the high-frequency framework is not to produce a different aggregate statistic but to make the temporal structure of poverty explicit, showing when, where, and for whom poverty accrues within the year.

Each of these findings points to a dimension of the poverty experience that the annual measure cannot capture: poverty is an episodic manifestation of uncertainty rather than a continuous status, households absorb shocks unevenly and smooth their consumption imperfectly, and exposure to deprivation extends well above the poverty line. Consequently, policy instruments constructed on the basis of annual measurement inherit the same limitations, whether they are statistical estimates, cash transfer schedules, or assessments of a household's capacity to repay loans. To this effect, we submit five recommendations, addressed to MoSPI, central and state governments, the RBI, financial service providers, and researchers, and set them out in the next section.

[1] High-frequency poverty compares a household's expenditure in a given month against the poverty line for that month, region and state

[2] Low-frequency poverty aggregates those monthly thresholds into an annual benchmark and compares it against average monthly expenditure over the year, yielding a measure comparable to conventional poverty statistics.

Recommendations

Summary of Recommendations

⇒ R1- MoSPI: Publish HCES microdata and report poverty estimates by period of survey, so that within-year variation already embedded in the data is not averaged away. Apply the same to the National Household Income Survey, where studying income volatility is particularly important.

⇒ R2- Central & state governments: Calibrate the timing and frequency of cash transfers to the lean periods households actually experience, rather than disbursing in evenly spaced instalments. No additional fiscal commitment, only a revision of scheme guidelines.

⇒ R3- RBI: Revisit income assessment methodology to incorporate an explicit measure of month-to-month income variability, rather than a single average figure. The Microfinance Directions (2022) are one place to give this effect but the principle applies wherever compliance rests on assessed income.

⇒ R4- FSPs: Collect and monitor high-frequency cash flow data to design products and repayment structures matched to household cash flow, drawing on Account Aggregator infrastructure and on loan officers' existing knowledge of borrower cash flows.

⇒ R5- Researchers: Reproduce the high-frequency framework beyond CPHS, with priority to the seasonality of deprivation across regions, and treat the temporal specification of the poverty line as a design choice in work on updating India's poverty benchmarks.

R1: For the Ministry of Statistics and Programme Implementation (MoSPI)

Household surveys that follow certain data-collection guidelines may already produce aggregate estimates that approximate high frequency poverty estimates or timecounts. This is particularly true in the case of the Household Consumption Expenditure Survey's (HCES) survey design, which distributes its sample across the survey year, as a result of which the annual headcount it produces already embeds within-year variation in deprivation. However, during the pre-processing and compilation of the data, the details regarding the within-year variability in household consumption are averaged out.

Two recommendations follow:

- **On existing surveys:** We recommend that the microdata already collected through HCES be published in a form that permits consumption to be studied through a high-frequency lens, and that poverty estimates be additionally reported by period of survey. This would reveal how deprivation varies across the calendar year rather than only in aggregate. Notably, this only requires cross tabulation and publication of the data that has already been collected.

- **On forthcoming surveys:** The National Household Income Survey (2026) is the first pan-India survey focused on measuring household income. We recommend that its estimates likewise be reported by period of survey, and that the underlying microdata be published in a form permitting the same high-frequency lens to be applied to income as to consumption. This is of particular importance because household income is considerably more volatile than household expenditure. An income survey reported only in annual aggregate would obscure precisely the variability that matters most.

R2: For central and state governments designing cash transfer schemes

Existing unconditional cash transfer schemes such as the Mahtari Vandan Yojana in Chhattisgarh, the Chief Minister's Ladli Behna Yojana in Madhya Pradesh, as well as other similar state-level schemes targeting women disburse financial assistance across evenly spaced instalments. Here, the implicit assumption is that household demand is uniform across the year. However, evidence from the paper suggests that concentrating the transfer on months of observed poverty roughly halves headcount poverty relative to spreading the same amount evenly. In other words, an evenly spaced distribution is an inefficient strategy because it delivers assistance in months when the household is not poor. Therefore, we recommend that the timing and frequency of transfers be calibrated to the lean periods experienced by households in a given region, and that the design account for how far below the line the target households sit. Notably, this requires no additional fiscal commitment but only a revision of scheme guidelines.

R3: For the Reserve Bank of India (RBI)

Regulatory frameworks governing lending to low-income households rest on income as the basis for credit assessment. For instance, The Reserve Bank of India (Regulatory Framework for Microfinance Loans) Directions (2022) state that a household's monthly loan repayment obligations across all outstanding loans may not exceed 50 per cent of its monthly household income, referred to as the Fixed Obligations to Income Ratio (FOIR). Consequently, regulated entities assess eligibility using an income assessment policy, which in practice yields a single figure of household income. However, household income is considerably more volatile than household expenditure. A household assessed as being within the 50 per cent limit on its average income will therefore breach that limit in months when income falls below the average. This in turn weakens the household's capacity to repay its loan obligations. Therefore, we recommend that the income assessment methodology be revisited to incorporate an explicit measure of month-to-month income variability, rather than a single average figure. The Microfinance Directions are one place where this could be given effect, but the principle can be applied wherever regulatory compliance requires assessment on the basis of income.

R4: For Financial Service Providers (FSPs)

As our previous work on this subject has noted, the condition under which a low-income household borrows is a cash flow deficit arising at high frequency, while the condition under which an FSP lends is income sufficiency assessed at low frequency. The findings presented in the paper extend empirical support to this misalignment. Such misalignment can invite delinquency and push households back toward informal sources to service formal debt. Therefore, we recommend that high-frequency cash flow data be collected at the point of credit assessment and continuously monitored thereafter, which can then be used to design suitable products and repayment structures that match the household's actual cash flow position.

That said, we recognise that this is not costless. Collecting cash flow data at high frequency such as monthly or even weekly imposes additional operational demands on the field staff and is difficult to justify given the thin margins on small-ticket lending. Two considerations can mitigate these challenges. First, the Account Aggregator framework already permits machine-readable access to bank statements, deposit and credit details, and cash-flow metrics. This substantially lowers the cost of building a separate cash flow viewer for households with bank accounts. However, this may present to be a significant challenge for households whose income flows remain in liquid cash. Second, loan officers already hold considerable knowledge of borrowers' cash flow patterns, accumulated through recurring interactions at the last mile. Such knowledge would be valuable to regulated entities in designing products suited to their borrowers, and to the regulator in framing the income assessment methodology recommended in R3.

R5: For researchers studying the financial lives of low-income households

Evidence on intra-year poverty dynamics in India rests on a small number of datasets. We recommend that the high-frequency framework be applied to other existing panel data with priority given to the seasonality of deprivation by different geographical regions. Three limitations of the present analysis mark out further priorities.

First, the study period includes the COVID-19 pandemic and the associated nationwide lockdowns of 2020 and 2021, a structural break unlikely to reflect the routine uncertainty the framework is designed to capture. Reproduction over a period without such a break would establish how much of the observed within-year variation is expected. Second, households surveyed continuously for at least four years are likely to be less residentially mobile and more economically stable than those dropped from the survey. The sub-sample used in the paper may therefore under-represent precisely the households whose poverty experience is most volatile. Third, the Tendulkar Committee's poverty lines are based on 2011-12 prices, and the Consumer Price Index (CPI) adjustment carries them forward. Given the larger economic and infrastructural developments in India, the underlying basket of consumption goods may no longer reflect the needs of Indian households. Successive committees such as the Rangarajan Committee (2014) have revised the poverty line by revisiting the methodology for calculation of poverty as well as the composition and cost of the consumption basket, but the frequency at which the resulting benchmark is applied has received relatively little attention. A high-frequency framework emphasises that any new benchmark should be specified not only in terms of what a household needs to consume, but over what period that need is assessed. Further work on updating India's poverty benchmarks would therefore benefit from treating the temporal specification of the poverty line as a critical design choice rather than as a matter of convention.