Identifying Hand-to-Mouth Households: Evidence from India

Vivek Gupta
Shiv Nadar Institution of Eminence

Fiorella Pizzolon
Hamilton College

Aditi Singh*
CAFRAL, Reserve Bank of India

Click here for latest version

Abstract

Like most emerging economies, India has no dataset that jointly observes household income, consumption, and wealth. We build the joint distribution by harmonizing India's balance-sheet survey with the world's largest household panel. Roughly half of Indian households are wealthy hand-to-mouth, holding substantial wealth in land and gold but too little liquidity to cover even half a month of income. The treatment of gold matters substantially: counting it as liquid rather than illiquid reduces hand-to-mouth prevalence by half. The usability of this wealth, not only its amount, shapes monetary transmission. Four months after a contractionary surprise, poor hand-to-mouth households cut consumption 0.17 log points more than wealthy ones; by month seven, wealthy households without potential gold collateral cut consumption 0.79 log points more than those with it, a collateral margin with no counterpart in the mortgage channels of advanced economies. The same margin governs marginal propensities to consume: out of a transitory income change, wealthy hand-to-mouth households without pledgeable gold consume 0.12 over a survey wave, against 0.08 for the unconstrained and 0.05 for those whose gold suffices to pledge, and among predicted gold non-owners the classic ranking of hand-to-mouth over unconstrained MPCs reappears. Counting only the poor hand-to-mouth as constrained would therefore miss most of the households that respond as if constrained: those whose wealth sits in land and in gold too small to pledge.

Keywords: Hand-to-mouth households; liquidity constraints; monetary policy; household consumption; machine learning

JEL codes: E21, E52, D14, D91, G51

*We gratefully acknowledge the excellent research support provided by Ken Lam. The views expressed here are solely the responsibility of the author and should not be interpreted as reflecting the views of the Center for Advanced Financial Research and Learning (CAFRAL) or the Reserve Bank of India.

1 Introduction

The share of households living hand-to-mouth helps determine how changes in income translate into aggregate consumption. These households hold little liquid wealth relative to income and tend to adjust spending when income changes. Their prevalence, together with their exposure to changes in earnings, taxes, and transfers, therefore shapes the transmission of monetary and fiscal policy. It also provides an empirical target for heterogeneous-agent models used to evaluate policy counterfactuals (Kaplan et al., 2018). Measuring this population requires looking beyond poverty and net worth. Kaplan et al. (2014), henceforth KVW, distinguish poor hand-to-mouth households (P-HtM), which have no positive net illiquid wealth, from wealthy hand-to-mouth households (W-HtM), which own positive net illiquid wealth despite holding little liquidity. Quantifying both groups helps identify households whose consumption may be sensitive to short-run income movements, including those whose asset holdings would otherwise suggest financial security.

Recovering comparable HtM shares in developing economies faces a demanding data requirement. The KVW classification requires household income and detailed holdings of liquid and illiquid assets and debts, but this information is often spread across surveys. Income and expenditure surveys may lack asset values, while balance-sheet surveys may omit income or consumption. Existing international evidence uses measures of financial fragility to study HtM prevalence and its policy implications in emerging economies (Bracco et al., 2026), but these measures do not by themselves recover the distinction between poor and wealthy HtM households.

To our knowledge, we provide the first national estimates distinguishing poor and wealthy HtM households in an emerging economy, namely India, under the KVW classification, and connect these classifications to household consumption dynamics. We harmonize the All-India Debt and Investment Survey (AIDIS), which records detailed balance sheets but not income, with the Consumer Pyramids Household Survey (CPHS), which records monthly income and expenditure but not asset amounts. Using common covariates, we predict income in AIDIS. We then reverse the prediction direction, estimating liquid and illquid assets and debt in CPHS. This then allows us to classify households into the KVW hand-to-mouth categories.

Our benchmark estimates show that 59.6 percent of Indian households are hand-to-mouth: 8.2 percent are poor HtM and 51.4 percent are wealthy HtM. These shares are sensitive to the treatment of gold, which households may sell or pledge to finance current spending. Counting gold as illiquid raises the wealthy HtM share to 56.5 percent while leaving total HtM prevalence at 59.6 percent. Counting gold as liquid reduces total prevalence to 24.6 percent. We then consider whether gold held as an illiquid asset could provide sufficient collateral to move a household above the liquidity threshold. Among households classified as W-HtM when gold is illiquid, we define potential collateral capacity as net liquid wealth plus 75 percent of gold value exceeding one-half of monthly income. Under this rule, 31.9 percent of all households are W-HtM with potential collateral capacity, while 24.7 percent are W-HtM without it. For the latter group, liquid wealth and the assumed pledgeable value of gold are together insufficient to exceed the threshold.

We first examine responses to transitory income changes. Using the covariance restrictions of Blundell et al. (2008), we estimate marginal propensities to consume (MPCs) in CPHS. Over a survey wave, W-HtM households without gold collateral capacity have an estimated MPC of 0.12, significantly above the non-hand-to-mouth (N-HtM) estimate of 0.08. Collateral-capable W-HtM households have an MPC of 0.05, significantly below it. The difference within the W-HtM group is 0.07, with a standard error of 0.01. Among the 17.5 percent of households predicted to own no gold, HtM households have higher MPCs than N-HtM households at every horizon up to a year. The aggregate MPC ordering also reverses when gold moves between liquid and illiquid wealth in the classification. These results suggest that the composition of the HtM population matters alongside its size: households with little measured liquidity can differ in their consumption responses according to the potential usability of their illiquid assets.

We then study responses to monetary-policy surprises using local projections (Jordà, 2005). We combine the monthly CPHS panel over 2014–2025 with the target and path surprises of Lakdawala and Sengupta (2025). Four months after a one-standard-deviation contractionary target surprise, the consumption response of P-HtM households is 0.17 log points lower than that of W-HtM households, with a standard error of 0.08. A second contrast appears within the W-HtM group. At month seven, the consumption response of collateral-capable households is 0.79 log points higher than that of households without potential collateral capacity, with a standard error of 0.45. At month ten, the corresponding income-response difference is 0.75 log points, with a standard error of 0.33. Together with the MPC estimates, these patterns are consistent with potential access to gold-backed funds cushioning household cash flow.

Our paper relates first to the measurement of household liquidity constraints. Zeldes (1989) uses net worth to identify households likely to face borrowing constraints, while Kaplan et al. (2014) establishes the importance of separating liquid and illiquid wealth. Applications document wealthy hand-to-mouth households in Japan, South Korea, and Belgium (Hara et al., 2016; Song, 2020; Cherchye et al., 2023), and **?** connects hand-to-mouth shares to social-transfer multipliers across developed and emerging economies. We apply the liquidity distinction to detailed Indian balance sheets and show how measured prevalence depends on the treatment of gold. We also identify a measurement problem that arises when balance sheets must be predicted: recovering conditional means does not ensure that a model recovers the ownership patterns and wealth thresholds that define hand-to-mouth status.

The paper also speaks to consumption insurance and heterogeneous responses to monetary policy. Blundell et al. (2008) provide a framework for estimating how consumption responds to permanent and transitory income changes. Kaplan et al. (2018) and Auclert (2019) explain how differences in households' consumption responses and exposure to policy-induced changes in income and wealth shape monetary transmission. Empirically, Cloyne et al. (2020) show that consumption responses to monetary policy differ across households with different housing and mortgage positions. We examine a related distinction within the wealthy HtM population: whether illiquid assets provide potential collateral capacity. Both our MPC estimates and our local projections associate this capacity with smaller consumption responses. Over a survey wave, the estimated transitory MPC is 0.12 for W-HtM households without gold collateral capacity and 0.05 for those with it. These findings suggest that potential asset usability helps distinguish consumption behavior among households with little measured liquidity. They also provide empirical targets for models of monetary and fiscal policy in economies where physical assets account for much of household wealth.

The remainder of the paper develops the measurement framework before turning to its applications. Sections 2 and 3 describe and harmonize the surveys. Section 4 maps Indian balance sheets into liquid and illiquid wealth, and section 5 develops the prediction procedure. Sections 6 and 7 define hand-to-mouth status and estimate prevalence under the KVW benchmark and alternative asset classifications. Section 8 estimates MPCs by hand-to-mouth status. Section 9 studies potential gold collateral capacity and monetary transmission; and Section 10 concludes. The appendix documents data construction, the full balance-sheet map, prediction diagnostics, classification robustness, and additional monetary-policy results. We begin with the information each survey provides and the gaps that the prediction procedure must fill.

2 Data

In order to classify households into different hand-to-mouth categories, we need to know three things about them: their income, their net liquid wealth, and their net illiquid wealth[1]. No survey in India jointly measures all three. The All-India Debt and Investment Survey (AIDIS) records detailed household balance sheets but does not ask household income. The Consumer Pyramids Household Survey (CPHS) records income and expenditure at a monthly frequency but does not quantify assets or liabilities, making wealth calculations impossible. We use AIDIS to estimate the level and composition of household wealth, and CPHS to evaluate income, consumption, and their dynamics. We combine these sources with high-frequency monetary-policy surprises measured around Reserve Bank of India announcements to quantify the effects of monetary policy on Indian households.

2.1 The All-India Debt and Investment Survey

We use Visit 1 of the 2019 AIDIS, conducted by the National Statistical Office as part of the seventy-seventh round of the National Sample Survey (National Statistical Office, 2021). The survey records household assets item by item, including cash, deposits, securities, gold, land, buildings, livestock, transport equipment, and business machinery. It also records each outstanding loan, its source and purpose, and repayments made over the reference period. The original Visit-1 file contains 116,461 households. We drop 24 households whose head is recorded in the survey's transgender category, which has no counterpart in the CPHS head-gender variable, leaving 116,437 households. Non-head transgender household members remain in household size and age-composition measures.

The timing of AIDIS variables is not uniform. Asset stocks refer to June 30, 2018, whereas

[1]See Section 6 for details.

outstanding debt refers to the household's survey date in 2019. Loan repayments are flows since July 1, 2018. We convert these flows to monthly rates using the exact number of days between July 1 and the interview date. All AIDIS statistics use the National Sample Survey multiplier divided by 100.

We construct liquid and illiquid positions from the underlying asset and loan records. Benchmark liquid assets include cash; current and savings accounts; fixed-income, post-office, cooperative-bank, and nonbank-financial institution deposits; mutual funds; listed shares; and debentures and bonds. We subtract short-term liabilities for consumption, education, medical treatment, litigation, and Kisan Credit Card borrowing. Benchmark illiquid assets include land and buildings, provident and pension funds, cooperative-society shares, selected self-help-group claims, and receivables. Life-insurance sum assured is excluded because it is a contingent benefit, not wealth owned at the survey date. We subtract housing debt from the illiquid position. These definitions follow the economic distinction in Kaplan et al. (2014): a wealthy hand-to-mouth household has positive net worth but too little liquid wealth to finance ordinary expenditure between income receipts.

Gold, transport equipment, livestock, and machinery do not enter either benchmark position. Their liquidity is the question rather than an assumption we can settle from the survey. Later sections therefore move these assets through a placement ladder (outside both aggregates, illiquid, collateralizable, and liquid) and show how each placement changes the classification. Keeping them outside the benchmark makes each alternative an explicit statement about economic usability rather than an implicit feature of a broad balance-sheet aggregate.

2.2 The Consumer Pyramids Household Survey

CPHS is a longitudinal household survey produced by the Centre for Monitoring Indian Economy (Centre for Monitoring Indian Economy, 2025). It surveys households three times per year and records monthly reference-period data on expenditure, income, and household members. We use four modules. Consumption Pyramids provides expenditures by detailed category; Income Pyramids provides total income and its components; People of India provides member demographics, employment, health, and education; and Aspirational India provides asset-ownership indicators, saving behavior, and borrowing indicators. These asset variables measure ownership, not rupee amounts. CPHS therefore cannot reveal where a household lies in the liquid-wealth distribution without the prediction step developed below.

The cross-sectional bridge to AIDIS uses asset ownership from CPHS wave 14, fielded from May through August 2018, and income, expenditure, and demographics from the 2019 reference months. A household must appear in the selected asset wave and in at least eight of the twelve 2019 months in each of the consumption, income, and people modules. We average continuous variables over the observed months, define binary variables as ever reported over the window, and retain the first observed value of categorical variables. These rules yield 131,789 households. We use the CMIE household weights as released.

The released CPHS weights do not eliminate all differences between CPHS and national household surveys. In our common-variable sample, CPHS households are more urban and more educated than their AIDIS counterparts before weighting, and differences remain in several financial-ownership margins after weighting. We therefore use AIDIS, which directly observes balance sheets, for national level statements. We use CPHS for within-sample dynamics and report a post-stratification exercise that reweights CPHS to AIDIS household-composition margins.

For the monetary-policy application, we use the CPHS expenditure panel as it enters our estimation. The three outcomes are nominal total expenditure, nominal food expenditure, and nominal expenditure on appliances. Total expenditure in these regressions is the released CPHS total; unlike the cross-survey consumption measure in Section 3, it is not stripped of equated monthly installments, augmented by self-production, or deflated by the consumer price index. We take logarithms of positive expenditure and multiply by 100, so coefficients are in log points.

Each available hand-to-mouth classification is attached to households by identifier and held fixed over the monthly panel. The principal direct classification is measured in the 2019 cross-section using observed CPHS income and predicted balance sheets. Thus, for outcomes before 2019, status is measured after the monetary-policy surprise; for outcomes after 2019, status predates the surprise. The design treats the fixed label as a proxy for a persistent portfolio type, not as a contemporaneous balance sheet. This timing is a limitation of the current implementation, and the monetary-policy section evaluates results across several independently constructed fixed classifications rather than interpreting the 2019 label as predetermined in every month.

2.3 Monetary-policy surprises

We use the high-frequency monetary-policy surprises constructed by Lakdawala and Sengupta (2025). The underlying series measures changes in Overnight Index Swap rates around Reserve Bank of India announcements and decomposes them into a target factor and a path factor. The target factor captures news about the current policy rate; the path factor captures news about its expected future path.

The shock series runs from November 2003 through February 2023. We sum announcement-level surprises within calendar months and assign zero to months with no announcement. The household application uses the overlap beginning in 2014 and ending in February 2023; CPHS observations after that date do not contribute to regressions requiring the shock. The headline specification includes the target and path factors jointly using the one-day event window. A first principal component, a two-day event window, and a sample excluding March 2020 through June 2021 serve as robustness checks. Section 3 next explains how we put AIDIS and CPHS on a common measurement footing before predicting the quantities that each survey omits.

3 Harmonization and Measurement

The two surveys are not linked household by household. We first construct a common schema and then use the overlap in observed variables to predict the quantities missing from each survey. The credibility of that exercise depends on whether similarly named variables measure the same economic object. We therefore harmonize items, reference periods, and units before estimating any prediction model.

3.1 Reference periods and common covariates

We align stocks with stocks and flows with flows. The CPHS asset indicators come from May–August 2018, bracketing AIDIS's June 30, 2018 asset date. CPHS income, expenditure, and demographics are household averages over the observed 2019 reference months, while their AIDIS counterparts refer to the 2019 interview or the survey's stated recall window. CPHS binary indicators equal one if the household reports the item in any observed month. This rule is the closest counterpart to the AIDIS positive-value indicator, but the two survey instruments remain different: CPHS asks whether the household holds an asset, whereas AIDIS records a positive rupee value.

Both surveys identify states and districts, but AIDIS uses National Sample Survey district codes and CPHS uses CMIE district names. We match 498 of the 683 AIDIS districts exactly or through a verified alias. We pool remaining districts within state when the prediction model requires geographic fixed effects. Nine small states and union territories, representing 0.77 percent of AIDIS households, do not appear in CPHS; for prediction we map them to a neighboring covered state while retaining the original state for descriptive statistics. A common-states robustness exercise drops them.

We harmonize household size and age composition directly. Head age, gender, education, religion, and caste are mapped to common categories. Employment type follows the National Sample Survey distinction among self-employment, regular wage work, casual labor, and other work, with rural employment further divided between agricultural and nonagricultural activity. These mappings use member-level employment arrangements and industry rather than the older, coarser occupation-group crosswalk.

3.2 Consumption and income

The baseline consumption measure is designed to capture resources consumed, not cash leaving the household in a particular month. In AIDIS it equals purchased consumption, self-produced consumption, in-kind receipts, and one-twelfth of annual durable expenditure. In CPHS it equals total expenditure minus all equated monthly installments plus income from self-production. The EMI subtraction matters because the released CPHS total includes debt service, while AIDIS consumer expenditure does not. The self-production addition matters because AIDIS values home-grown consumption directly, whereas CPHS records the same resource on the income side.

Relative to a purchases-only measure, this construction raises mean monthly consumption to Rs. 10,048 in AIDIS and Rs. 12,092 in CPHS. The remaining 20 percent gap is concentrated in rural India; in urban India the weighted means are approximately equal. Dropping durables does not close the gap. This pattern points to the surveys' different recall designs (a single AIDIS question versus itemized monthly CPHS reports) rather than a single mismatched expenditure category. We retain five alternative consumption definitions, including versions that exclude self-production and durables, as robustness exercises.

We construct three CPHS income measures. Total income includes every source recorded by CMIE and is the baseline. Frequent income retains regularly recurring sources (wages, pensions, self-production, private and government transfers, and business profits) and equals 99 percent of total income at the weighted mean. Narrow income contains wages and government transfers. Total income pairs a complete resource measure with the complete consumption measure and is also the best-predicted target. Because self-produced resources do not pass through a liquid account, we also construct a cash-income threshold that subtracts predicted self-production. That change moves the wealthy hand-to-mouth share by 0.3 percentage points, so the main classification is not driven by treating in-kind resources as cash.

3.3 Balance sheets and descriptive comparison

Table 1 reports the moments that discipline the cross-survey comparison. Household size and head age are close across sources. Resource measures are not identical: CPHS consumption is 20 percent higher at the mean, and only CPHS observes income. Balance sheets are observed only in AIDIS and are extremely skewed. Median benchmark liquid assets are Rs. 5,200, compared with median illiquid assets of Rs. 665,000. This contrast, more than a factor of 100, is the central measurement fact behind the wealthy hand-to-mouth classification.

3.4 Validation and limits of the comparison

We use three sets of checks before transporting information across surveys. First, internal checks test whether separate AIDIS blocks describe the same household balance sheet. All 94,023 households reporting positive owned-and-possessed land in the household block have a corresponding member-level ownership report, and 99.6 percent also have plot-level land records. The small residual consists primarily of homestead plots with a median area of 0.02 acres.

Second, we compare weighted distributions rather than relying only on means. AIDIS consumption is strongly rounded: 48.9 percent of positive values are multiples of Rs. 100, compared with 0.2 percent in the household-mean CPHS cross-section. The baseline leaves reported values unchanged. A multiple-imputation jitter exercise bounds the resulting attenuation in the income prediction at approximately 0.2 percent, and the classification is nearly unchanged. Third, we compare AIDIS aggregates for assets, debt, and investment with Reserve Bank of India statistics and CPHS consumption with the National Sample Sur-

Table 1: Summary Statistics for the Harmonized Samples

	AIDIS		CPHS	
Variable	Mean	Median	Mean	Median
Total consumption (Rs./month)	10,048	8,250	12,092	10,875
Total income (Rs./month)	—	—	20,867	15,809
Benchmark liquid assets (Rs.)	35,822	5,200	—	—
Benchmark illiquid assets (Rs.)	1,784,434	665,000	—	—
Benchmark liquid debt (Rs.)	25,918	0	—	—
Benchmark illiquid debt (Rs.)	35,305	0	—	—
Age of household head	47.7	47.0	50.6	50.0
Household size	4.3	4.0	4.0	4.0

Notes: The AIDIS sample contains 116,437 households and the CPHS sample contains 131,789 households. Statistics use AIDIS multipliers divided by 100 and CMIE household weights, respectively. AIDIS asset stocks refer to June 30, 2018; debts are outstanding at the 2019 interview date. CPHS income and consumption are means over observed 2019 reference months. Gold, transport equipment, livestock, and machinery are excluded from both benchmark asset aggregates.
Source: Authors' calculations from AIDIS and CPHS.

vey benchmark. These external comparisons are order-of-magnitude diagnostics and do not rescale either survey. Appendix A records the implemented sample rules, reference periods, concordance, and validation limits.

These checks do not make the surveys interchangeable. The data support comparisons of ranks, ownership margins, and hand-to-mouth shares within a stated definition. They support response comparisons across household groups when the group definition is held fixed. They are less informative about exact cross-survey level differences, particularly for rural consumption; about the tails of balance sheets predicted by conditional-mean models; and about very small groups. The remainder of the paper therefore treats AIDIS as the anchor for level estimates and uses multiple prediction routes to show which conclusions survive when the missing quantity is changed.

4 Mapping Indian Balance Sheets into Liquid and Illiquid Wealth

Every result in the paper rests on a division of household assets and liabilities into liquid and illiquid positions. This section fixes that division before any quantity is predicted and before any household is classified.

Our benchmark net liquid position is liquid assets minus liquid debt, and our benchmark net illiquid position is illiquid assets minus housing debt. We call this mapping the KVW benchmark, after Kaplan et al. (2014): it is the definition we use for headline prevalence because it preserves international comparability, not a claim that it is the preferred economic classification for India; which assets can finance current spending is precisely the question the placement ladder and the collateral-capacity split address. The mapping is consequential because AIDIS records forms of wealth that are important in India but do not fit neatly into a financial-assets-versus- housing division. We impose four corrections before constructing any classification. Life-insurance sum assured is excluded because it is a contingent benefit rather than an asset owned at the survey date. Cooperative-society shares are treated as illiquid. Kisan Credit Card balances enter liquid debt because they are the revolving facility observed in AIDIS. Post-office saving certificates enter only in the broader definition ladder. Appendix B.1 gives the complete item map and shows how we prevent double counting.

Underlying the whole exercise is the position that liquidity is not an attribute of an asset in isolation. It is an attribute of the asset together with the market in which the household operates: whether a secondary market exists, how deep it is, what transaction costs it imposes, and whether institutions will lend against the asset. The same asset can therefore be classified reasonably differently in India and in the United States, and a classification built for one economy should not be transplanted to the other unexamined. We retain the KVW benchmark precisely to make this point measurable: it fixes the international-comparison boundary, and every departure from it in the placement ladder and the collateral-capacity split quantifies how much the Indian market context moves the answer.[2]

Gold, art, transport equipment, livestock, and machinery sit outside both benchmark asset aggregates. This choice preserves the financial-assets- plus-housing boundary used for international comparison. It is not an assertion that these assets are economically irrelevant. We instead vary their placement explicitly. Treating gold as illiquid wealth changes the P/W split without changing total HtM. Treating it as liquid wealth allows gold to move households out of HtM. A broader specification also admits physical assets and additional debt categories. This ordered ladder makes clear whether a result follows from the amount of wealth a household owns or from the ease with which it can use that wealth for current spending.

We also implement an India-specific real-estate screen. It removes asset items that are not plausibly usable as housing or collateral and applies a Rs. 1 lakh floor to agricultural land and the dwelling bundle. The adjustment affects net illiquid wealth, so it can move households from W-HtM to P-HtM but cannot change total HtM. We retain the unadjusted KVW benchmark as the headline definition and report the real-estate screen as a main-text robustness exercise.[3] Appendix B.2 treats this floor and the credit-limit alternatives as implemented sensitivity parameters, not as household-level observations of policy eligibility

[2]The same logic applies to the small financial categories. Listed shares and mutual funds are conventionally liquid, though their usability in an aggregate downturn is weaker than in normal times because prices fall exactly when liquidity is needed; cooperative-society shares are locally held and hard to transact, which is why we classify them as illiquid. These categories are small shares of Indian household portfolios, and moving them across the boundary changes the classification shares negligibly.

[3]Like gold, land and housing are classified by effective liquidity rather than ownership alone: a household may own land without a functioning market through which to sell it quickly or borrow against it. Ongoing digitisation of land and property records could raise effective pledgeability and, through the same logic as the gold-capacity split, reduce measured wealthy hand-to-mouth status; but digitisation by itself creates neither a secondary market nor lenders willing to accept the collateral, so we treat such improvements as counterfactual movements along the placement ladder rather than as features of the current data.

or loan approval.

Debt enters the classification only through net positions, never as a marker of constraint by itself.[4]

With the aggregates fixed, the two surveys share a common balance-sheet schema and a common covariate set. The next step is to fill each survey's missing side: income in AIDIS and wealth amounts in CPHS.

5 The Prediction Bridge

The two surveys leave opposite holes. AIDIS observes wealth but not income; CPHS observes income but not wealth amounts. We fill these holes using only variables that have been harmonized across the two surveys, and we evaluate the predictions where they will matter for the classification. In brief, a household is hand-to-mouth when its net liquid wealth is either small, no more than one-half of monthly income, or close enough to its borrowing limit that current income determines current spending; hand-to-mouth households with positive net illiquid wealth are wealthy hand-to-mouth, and those without are poor hand-to-mouth. Section 6 states the rule formally. The classification thus turns on a rupee threshold in net liquid wealth and a kink at zero in net illiquid wealth, and we judge predictions at those points in addition to average fit. The exercise requires a transportability assumption: conditional on the common covariates, the mapping from consumption and household characteristics to income can be carried from CPHS to AIDIS, and the mapping from the same covariates to balance sheets can be carried from AIDIS to CPHS. No model-fit statistic can test this assumption directly. We instead combine held-out within-survey prediction, out-of-survey checks on common objects, covariate-shift diagnostics, and alternative prediction rules that preserve different features of the target distribution.

5.1 Common evaluation protocol

All headline prediction models use the same fixed 20 percent household holdout and common cross-validation folds. We estimate and evaluate the models with survey weights. Predictors include harmonized consumption, demographics, geography, employment, durable ownership, borrowing, saving, and financial-access measures. Monetary predictors enter in logarithms or inverse-hyperbolic-sine form. The holdout is never used to choose the reported fit, the wealth-ownership threshold, the Duan smearing factor, or the quantile map.

The common holdout makes the model comparisons meaningful, but it addresses prediction

[4] A household with outstanding debt is not automatically hand-to-mouth under any definition in this paper. Unsecured debt reduces net liquid wealth and housing debt reduces net illiquid wealth, so a leveraged household with ample liquid assets and good credit access is classified by its net position, exactly as a debt-free household would be. The borrower branch of the classification rule additionally requires the liquid position to sit near the assumed credit limit, and Appendix B.2 varies that limit across institutional alternatives. Distinguishing debt by liquidity class, rather than treating all debt symmetrically, is what the two-sided netting accomplishes.

error only within the training survey. The transfer from one survey to the other also depends on common support and on whether the same reported characteristic has the same informational content in both survey instruments. The harmonization checks in Section 3 address the observable part of this problem. Appendix B.3 describes the complete model ladder and the shift and reweighting exercises.

5.2 Predicting income in AIDIS

We estimate monthly total income in CPHS and deploy the fitted mapping in AIDIS. The target is log total income. Table 2 compares the active headline models on 26,344 held-out CPHS households. Histogram gradient boosting performs best: its holdout R^2 is 0.804 and its rank correlation is 0.886. The gain over consumption bins is 0.230 in R^2, and the gain over elastic net is 0.071. Reweighting the training sample toward AIDIS covariates changes the gradient-boosting R^2 from 0.804 to 0.801, which indicates that the model's within-CPHS performance is not driven by a small region of obvious covariate mismatch.

Table 2: Prediction of Monthly Total Income in CPHS Holdout

Model	RMSE	MAE	R^2	Rank correlation
Consumption bins	0.388	0.303	0.574	0.749
Elastic net	0.307	0.233	0.733	0.848
Random forest	0.290	0.217	0.762	0.858
Histogram gradient boosting (selected)	0.263	0.197	0.804	0.886
Histogram gradient boosting, reweighted	0.265	0.200	0.801	0.882

Notes: The target is log monthly total income. Statistics are survey-weighted and evaluated on the common 20 percent household holdout (N = 26,344). Rank correlation is the weighted correlation of ranks. The reweighted model gives greater weight to CPHS observations that resemble the AIDIS covariate distribution. The selected model, histogram gradient boosting, supplies the predicted income used in the remainder of the paper.
Source: Authors' calculations from CPHS.

We use histogram gradient boosting for the headline AIDIS classification. We also estimate frequent income, narrow income, a version that excludes consumption from the predictor set, and a quantile-gradient- boosting draw. These variants address different concerns rather than serve as interchangeable searches for the highest fit. Frequent and narrow income test whether the liquidity threshold should exclude irregular resources. The no-consumption model tests whether the predictor set depends too heavily on household expenditure. The quantile draw preserves residual dispersion. The benchmark P/W/N shares change little between the headline gradient-boosting income and the quantile draw.

CPHS also records self-production income, and AIDIS independently records the value of self-produced consumption. We use that overlap as an out-of-survey check: the model is trained on the CPHS concept and compared with the separately elicited AIDIS item. This is more demanding than a random holdout because it tests the mapping across survey instruments. It does not validate total income, which remains unobserved in AIDIS, but it tests an economically important component for rural households.

5.3 Predicting balance sheets in CPHS

We predict liquid assets, liquid debt, illiquid assets, and illiquid debt separately. Each aggregate uses a two-part model. The first part predicts whether the household has a positive amount; we refer to this ownership margin as *participation* throughout, the extensive margin of the balance sheet, as distinct from the amount held conditional on holding. The second predicts the amount conditional on a positive value. The combined prediction equals the conditional amount when the predicted ownership probability exceeds a calibrated threshold and zero otherwise. We calibrate this threshold in AIDIS so that predicted and observed ownership rates agree, then freeze it before applying the model to CPHS. Preserving genuine zeros is essential: the classification has a kink at zero, and a conventional conditional-mean prediction would mechanically empty parts of the P-HtM group.

Table 3 summarizes performance on the common AIDIS holdout, in the format of Table 2. Two facts stand out. First, the two parts of the model behave very differently: ownership is predicted almost perfectly at the calibrated thresholds, while amounts among holders carry moderate fit at best. Second, the aggregates differ sharply. Debt positions and illiquid assets are well predicted; liquid assets are not. Gradient boosting, the preferred conditional-mean model, reaches a combined IHS-R^2 of only 0.289 for liquid assets, against 0.920 for liquid debt, 0.782 for illiquid assets, and 0.987 for illiquid debt. The asymmetry is economic rather than computational. Liquid assets are nearly universally held and their amounts are highly dispersed and volatile, leaving the model with little extensive-margin variation and substantial idiosyncratic intensive-margin variation that consumption and demographic covariates cannot capture. Alternative predictor sets and tuning variants move the liquid-assets figure only to between 0.32 and 0.35, so the weak level fit reflects the information set, not the choice of estimator. Yet liquid assets are precisely the object compared with one-half of monthly income. The weakest level prediction therefore enters the most consequential classification margin.

We address this problem with distribution-preserving legs on both margins of the two-part model. The quantile-mapped model retains the gradient-boosting rank of predicted positive amounts and replaces each amount with the observed AIDIS amount at the same training-sample quantile; this alone raises the combined liquid-assets IHS-R^2 on the holdout from 0.289 to 0.557. Amounts, however, are only half of the deployment problem. The P-HtM group requires nonpositive net illiquid wealth, and the ownership models deployed on CPHS covariates predict nearly universal illiquid-asset holding: the deployed probabilities lack the low-participation mass observed in AIDIS. Mapping amounts among predicted holders cannot repair that margin. Our headline deployment route therefore also rank-maps the CPHS distribution of predicted ownership probabilities onto its weighted AIDIS counterpart for each predicted aggregate, while retaining the quantile-mapped amounts. This participation map is applied only at deployment and is the identity on AIDIS itself, so it changes no within-AIDIS validation statistic. It adds one explicit assumption: AIDIS participation rates in assets and liabilities apply to the CPHS population. The stochastic alternative

Table 3: Prediction of Balance-Sheet Aggregates in AIDIS Holdout

	Liquid assets	Liquid debt	Illiquid assets	Illiquid debt
Combined prediction, holdout IHS-R^2				
Wealth bins	−0.554	−1.022	−0.849	−1.549
Consumption bins	−0.289	−2.704	−0.215	−9.510
Two-part elastic net	0.276	0.896	0.765	0.979
Two-part random forest	0.345	0.886	0.773	0.987
Two-part gradient boosting (selected)	0.289	0.920	0.782	0.987
Quantile-GB draw	0.323	0.938	0.827	0.983
Memo: two-part gradient boosting, by margin				
Ownership accuracy at calibrated threshold	0.971	0.989	0.970	1.000
Amount R^2 among positive holders (log)	0.436	0.409	0.582	0.584
Rank correlation, quantile-mapped amounts	0.741	0.966	0.854	0.999

Notes: Entries are survey-weighted statistics on the common 20 percent AIDIS household holdout. The combined prediction follows equation (7): the conditional amount when the predicted ownership probability exceeds the calibrated threshold and zero otherwise. IHS-R^2 is computed on the inverse-hyperbolic-sine scale; negative entries indicate fit worse than a constant. Ownership accuracy is the share of households whose predicted holding status at the calibrated threshold matches observed status. Amount R^2 is on log amounts among observed positive holders. Rank correlation is the weighted correlation between quantile-mapped amounts and observed amounts. Two-part gradient boosting is the selected model; the remainder of the paper deploys it in CPHS with its amounts quantile-mapped and its participation margin mapped onto the AIDIS distributions, as described below.
Source: Authors' calculations from AIDIS and CPHS.

combines quantile-gradient-boosting amounts with draws on the extensive and residual margins. These procedures answer different questions: quantile mapping seeks household-level agreement while restoring dispersion; the participation map additionally seeks the correct ownership margins; the draw seeks the correct marginal distribution and accepts additional household-level noise.

The classification test reveals the cost of relying on conditional means. On the AIDIS holdout, where the true balance sheet is observed, substituting gradient-boosting conditional means yields 67.6 percent agreement with the observed three-way classification. Quantile mapping raises agreement to 73.3 percent; the participation-mapped route coincides with it on the holdout by construction. The draw reproduces the aggregate HtM share more closely but lowers household-level agreement to 61.5 percent. These results explain why the CPHS direct-route level is method-dependent: good rank prediction does not by itself locate a household correctly relative to a rupee threshold. Appendix B.4 reports the corresponding class shares, ownership thresholds, and aggregate-level fit.

The participation margin matters only in deployment, where it is decisive for the P-HtM group. Under every amount-only route, the deployed CPHS P-HtM share is roughly 0.1 percent, against 8.2 percent in AIDIS (Appendix B.5). The chain from the observed AIDIS holdout share (8.0 percent) through the holdout prediction under quantile mapping (9.5 percent) to deployment locates the collapse: it is not model error, which the holdout rules out, but the smoothness of the deployed participation probabilities, compounded by CPHS's documented undercoverage of the poorest, asset-less households. Mapping the participation margin restores a deployed P-HtM share of 3.7 percent, stable between 3.6 and 4.3 percent across the calibrated ownership thresholds and consistent with a composition-adjusted share below the AIDIS level. The resulting classification, quantile map plus participation map, is the headline CPHS route in the remainder of the paper.

Gold receives the same two-part treatment, with one repair that the validation statistics themselves exposed. In the harmonized AIDIS data the indicator for gold savings is constructed from the recorded gold value, so the ownership model for gold was effectively shown its own target: its holdout classification was exact and its calibrated threshold collapsed to zero. Deployed on CPHS, where the same indicator is a narrow survey question answered affirmatively by about one percent of households, that threshold classified every household as a gold owner, leaving the 17.5 percent AIDIS non-ownership margin with no CPHS counterpart and overstating collateral capacity: on the holdout, the capacity condition under that predicted gold labels 81 percent of W-HtM households collateral-capable against 58 percent under observed gold. The headline gold prediction therefore refits the ownership model with the derived indicator excluded, which yields an interior calibrated threshold and a holdout ownership agreement of 84.3 percent against an always-owner benchmark of 83.2 percent: the honest measure of how weakly covariates separate owners in a population where gold ownership is nearly universal. Deployment then follows the participation-map logic exactly as for the other assets: the clean ownership scores are rank-mapped onto their AIDIS distribution and cut at the AIDIS-calibrated threshold, so 17.5 percent of CPHS households hold zero predicted gold, with the covariates determining which households they are, and positive amounts are rank-mapped onto the observed positive AIDIS gold distribution fitted on training rows. Under this piecewise gold prediction the deployed capacity split of the W-HtM group is 15.2 percent collateral-capable and 10.0 percent not, and the capable share of the wealthy group, 60.3 percent, sits next to the observed-gold holdout benchmark of 57.6 percent. Because gold enters the classification only through the capacity condition, the P-HtM, W-HtM, and N-HtM margins are unchanged by construction; Section 9 and Appendix D use this corrected capacity split throughout.

6 Classifying Hand-to-Mouth Households

The distinction between poor and wealthy hand-to-mouth households is a distinction between two balance sheets, not two income levels. Both groups hold little liquid wealth relative to their income. Wealthy hand-to-mouth households nevertheless own positive net illiquid wealth; poor hand-to-mouth households do not. With the balance-sheet aggregates of Section 4 and the predictions of Section 5 in hand, this section sets out the classification rule, distinguishes the four empirical routes used in the paper, and explains how the 2019 classification is carried to the monthly panel. Section 7 then reports the resulting prevalence estimates.

6.1 The classification rule

We follow the timing convention in Kaplan et al. (2014). The economic idea is that a household is hand-to-mouth when its liquid position is sufficiently close to either zero or its borrowing limit that current income determines current spending. Let m_i denote household i's net liquid wealth, a_i its net illiquid wealth, and y_i its monthly income. The baseline assumes a monthly pay period and sets $\tau = 2$, so the zero-liquid-wealth neighborhood extends to one-half of monthly income. A household satisfies the liquidity condition if

$$H_i^L = \mathbf{1}\left\{0 \le m_i \le \frac{y_i}{2}\right\} \;\vee\; \mathbf{1}\left\{m_i \le 0 \text{ and } m_i \le \frac{y_i}{2} - \bar{m}_i\right\}, \tag{1}$$

where $\bar{m}_i$ is the unsecured credit limit. We set $\bar{m}_i = y_i$ in the baseline.[5] The first branch of equation (1) captures households with small nonnegative liquid balances. The second captures borrowers whose negative liquid position places them sufficiently close to or beyond the assumed credit limit. We classify a household as hand-to-mouth when either branch is satisfied.

[5] The $y_i/2$ threshold is the Kaplan et al. (2014) value: if income arrives monthly and is consumed uniformly over the month, average cash on hand over the pay period is half of monthly income, so a household holding less than that cannot smooth within the period. We retain it for comparability with the international literature rather than re-derive it for India. Two qualifications apply. First, Indian expenditure is not uniform within the month: rent, utilities, school fees, and similar committed payments are typically due near the start, which front-loads cash needs and makes $y_i/2$ conservative as a measure of the buffer actually required. Second, the classification's other free parameters are varied explicitly rather than defended a priori: the credit limit $\bar{m}_i$ across institutional alternatives in Appendix B.2, and the calibrated ownership thresholds of the prediction models, across which the deployed P-HtM share moves only between 3.6 and 4.3 percent.

We then use net illiquid wealth to divide these households into two groups:

$$\text{P-HtM:} \quad H_i^L = 1 \quad \text{and} \quad a_i \le 0, \tag{2}$$
$$\text{W-HtM:} \quad H_i^L = 1 \quad \text{and} \quad a_i > 0. \tag{3}$$

Households that satisfy neither condition are N-HtM. This formulation differs from a net-worth test. A household with little liquidity and a home can be W-HtM even when its net worth is large. The net-worth classification therefore answers a different question and appears only as a robustness exercise in Appendix B.2.

The AIDIS income prediction is positive by construction because the model predicts log income and exponentiates the result. In the CPHS direct route, we restrict the classification to households with positive observed monthly income. All reported shares use survey weights.

6.2 What each classification route establishes

Neither survey observes both sides of the classification. We therefore use four related procedures, but they do not have equal evidentiary status. Table 4 records what is observed, what is predicted, and how each output enters the paper.

Table 4: Classification Routes and Their Empirical Roles

Route	Observed object	Predicted object	Role
AIDIS anchor	Liquid and illiquid balance sheet	Monthly income	Headline level shares
CPHS direct	Monthly income	Four balance-sheet aggregates	Cross-survey diagnostic; 2019 labels fixed in the monthly panel
AIDIS-label transport	CPHS harmonized covariates	AIDIS class probabilities	Cross-survey robustness
Within-CPHS extrapolation	2019 direct-route labels and CPHS-native covariates	Household HtM class	Robustness labels for the monthly panel

Notes: The panel classification is fixed at the household level; the local projections do not reclassify households after each shock. Alternative direct-route labels are retained for robustness.

We treat AIDIS as the anchor because it observes the balance sheet, the object that determines whether a household lies near the liquidity kink. The CPHS direct route is a diagnostic rather than an independent validation sample: CPHS does not reveal the balance sheet against which its predicted values could be checked. Its level estimates also depend on how the prediction restores the dispersion of wealth. Table 5 reports the deployed classification under each implemented wealth-prediction method, holding the KVW benchmark boundary fixed. Conditional-mean models compress the predicted balance sheet and generate the lowest HtM shares. Quantile mapping and the stochastic draw restore dispersion and raise W-HtM, from 19.0 percent to 28.9 and 36.4 percent; wealth bins yield the largest share, 49.6 percent. The P-HtM share is essentially zero under every route that predicts only amounts and 3.7 percent under the headline route, which also maps the participation margin onto its AIDIS distribution.

Table 5: Direct CPHS Classification by Wealth-Prediction Method

Prediction method	P-HtM	W-HtM	N-HtM
Wealth bins	0.0	49.6	50.4
Logit plus log-OLS	0.1	21.8	78.1
Elastic net	0.1	19.5	80.5
Random forest	0.1	25.0	74.9
Histogram gradient boosting	0.1	19.0	80.9
Gradient boosting, quantile-mapped	0.1	28.9	71.0
Quantile map plus participation map (headline)	3.7	25.2	71.1
Quantile-GB draw	0.2	36.4	63.4

Notes: Entries are weighted percentages of CPHS households. All rows use observed positive CPHS income, the KVW benchmark balance-sheet boundary, the combined rule, and each model's ownership threshold calibrated in AIDIS. Rounding can cause rows not to sum to exactly 100.
Source: Authors' calculations from AIDIS and CPHS.

These rows should not be read as eight estimates of one parameter. They make different choices about the feature of the wealth distribution to preserve. These differences combine prediction error with genuine cross-survey composition differences, and they are too large to interpret as independent confirmation of a single HtM level.

Even under the headline route, the CPHS classification is far less hand-to-mouth than AIDIS: 28.9 percent of households against 59.6, with 71.1 percent N-HtM. Both margins of the comparison contribute. CPHS undercovers the poorest, asset-less households and is more urban and more educated than AIDIS (Section 2), so part of the gap is genuine composition; the rest is the method dependence documented in Table 5. This is why AIDIS, which observes the balance sheet, anchors all national level statements. Within CPHS, however, the asset-placement experiments of Section 4 can be repeated on the predicted balance sheets. Table 6 reports that ladder under the headline route, with gold entering through the piecewise prediction of Section 5. The pattern mirrors the AIDIS anchor reported in Section 7: moving gold into illiquid wealth reallocates households from P-HtM to W-HtM while leaving total HtM and N-HtM unchanged, and moving gold or the broader physical bundle into liquid wealth collapses total prevalence, here from 28.9 to 9.4 and 5.9 percent. The gold-illiquid P-HtM cell remains far smaller in CPHS (0.8 percent) than in AIDIS (3.1 percent): even with the corrected 17.5 percent non-ownership margin, predicted gold holding is more widespread among the deployed P-HtM candidates than observed holding is among their AIDIS counterparts. Table 6 therefore doubles as a caution for CPHS-only measurement. Without the AIDIS anchor, and without mapping both prediction margins onto their AIDIS distributions, the level of hand-to-mouth prevalence and the entire poor hand-to-mouth margin would be artifacts of the prediction method rather than features of the population.

Table 6: Hand-to-Mouth Shares in CPHS under Alternative Balance-Sheet Boundaries

Definition	P-HtM	W-HtM	N-HtM	Total HtM
KVW benchmark	3.69	25.22	71.09	28.91
Benchmark with real-estate screen	3.70	25.21	71.09	28.91
Gold included in illiquid wealth	0.75	28.16	71.09	28.91
Gold included in liquid wealth	1.23	8.14	90.62	9.38
Broader physical assets liquid	0.81	5.11	94.08	5.92
Broad illiquid wealth and debt	0.41	28.70	70.89	29.11

Notes: Entries are weighted percentages of CPHS households under the headline direct route: observed positive CPHS income and participation-mapped predicted wealth, with each ownership threshold calibrated in AIDIS and gold from the piecewise prediction. Rows follow the definition ladder of Appendix B.1 and match the AIDIS ladder reported in Section 7. All rows use monthly income, $\tau = 2$, and the combined zero-kink and borrower rule.
Source: Authors' calculations from AIDIS and CPHS.

This sensitivity is the principal measurement threat in the paper. Conditional-mean predictions can rank households well while compressing the wealth distribution away from the threshold that defines HtM. It is why Section 5 evaluates prediction methods by the classification they imply, in addition to conventional loss functions, and why AIDIS remains the level anchor throughout.

6.3 From the 2019 classification to the monthly panel

The local projections require a household classification throughout the CPHS panel, not a new predicted balance sheet in every month. The headline panel classification attaches the 2019 direct-route labels formed from observed income and the participation-mapped predicted wealth to the monthly panel and holds them fixed. Under these labels the panel contains a usable P-HtM cell (4,035 of the 131,789 bridge households, 3.7 percent weighted), so the three-way contrasts in Section 9 can be estimated for the first time alongside the gold-capacity split.

Holding the label fixed prevents the monetary-policy shock from changing the group definition mechanically. It also means that the local projections inherit the direct route's sensitivity to the wealth model used to create the labels. We therefore retain the quantile-mapped, conditional-mean, and distribution-preserving label constructions as robustness classifications, together with a within-CPHS extrapolation route that predicts the 2019 labels from the survey's full native covariate set. The selected extrapolation model is gradient boosting; on the common 20 percent CPHS holdout, its full-native specification has overall accuracy of 0.948 and balanced accuracy of 0.737, with recall of 0.99 for N-HtM, 0.45 for P-HtM, and 0.77 for W-HtM. The low P-HtM recall reflects the very small P-HtM cell produced by its conditional-mean training labels (Appendix B.6), not high precision about that group. The appendix records the level differences across label constructions; the monetary-policy section shows whether the response contrasts survive those differences.

7 The Prevalence and Balance Sheets of Hand-to-Mouth Households

India's hand-to-mouth population is large and overwhelmingly wealthy. Under the KVW benchmark definition, 59.6 percent of households are hand-to-mouth. Fully 51.4 percent are W-HtM, while 8.2 percent are P-HtM. Thus, more than six W-HtM households appear for every P-HtM household, and W-HtM households account for 86 percent of the hand-to-mouth population. The central fact is not simply that many Indian households have little liquidity. It is that most of those households simultaneously hold substantial illiquid wealth.

7.1 Headline prevalence

Table 7 reports how the result changes as assets are moved across the liquid–illiquid boundary. The first row applies the KVW benchmark. The second applies the India-specific real-estate screen. Because that screen changes only illiquid wealth, it reallocates households between P-HtM and W-HtM without changing total HtM. The remaining rows progressively expand the assets that households may use for current spending or count against a broader set of debts.

Table 7: Hand-to-Mouth Shares under Alternative Balance-Sheet Boundaries

Definition	P-HtM	W-HtM	N-HtM	Total HtM
KVW benchmark	8.16	51.42	40.41	59.59
Benchmark with real-estate screen	10.66	48.92	40.41	59.59
Gold included in illiquid wealth	3.06	56.53	40.41	59.59
Gold included in liquid wealth	3.96	20.65	75.39	24.61
Broader physical assets liquid	3.07	11.12	85.80	14.20
Broad illiquid wealth and debt	2.23	59.24	38.53	61.47

Notes: Entries are weighted percentages of AIDIS households. The KVW benchmark follows the balance-sheet boundary of Kaplan et al. (2014), with gold, transport equipment, livestock, and machinery outside both asset aggregates. Income is predicted by the headline gradient-boosting model. All rows use monthly income, $\tau = 2$, and the combined zero-kink and borrower rule. The first five rows set the credit limit equal to one month of income. The broad row uses the broad asset and debt aggregates. See Appendix B.1 for the item map.
Source: Authors' calculations from AIDIS and CPHS.

The placement of gold is the consequential margin. Counting gold as illiquid wealth raises the W-HtM share to 56.5 percent but leaves total HtM unchanged. Counting the same asset as liquid reduces total HtM to 24.6 percent. Moving a broader set of physical assets into liquid wealth reduces the share further, to 14.2 percent. These rows are not alternative estimates of a single parameter. They answer different economic questions about which assets can finance current expenditure. Table 6 in Section 6 reports the same ladder for the deployed CPHS classification; the gold reallocation and the liquidity collapse replicate there.

7.2 Balance sheets and household characteristics

Table 8 makes clear why income alone does not identify the groups. Median predicted monthly income rises across the three groups, but the differences are modest relative to the contrast in their portfolios. P-HtM and W-HtM households both hold roughly Rs. 900 in net liquid wealth. Yet the median W-HtM household holds Rs. 675,000 in net illiquid wealth, compared with zero for the median P-HtM household. The median N-HtM household differs primarily on the liquid margin, holding Rs. 19,764 in net liquid wealth.

Table 8: Balance Sheets and Characteristics by Hand-to-Mouth Status

	P-HtM	W-HtM	N-HtM
Share of households (percent)	8.2	51.4	40.4
Median net liquid wealth (Rs.)	844	900	19,764
Median net illiquid wealth (Rs.)	0	675,000	910,000
Owns a house (percent)	2	95	84
Has outstanding borrowing (percent)	25	52	29
Rural household (percent)	32	76	62
Median predicted monthly income (Rs.)	10,964	12,094	13,133

Notes: Shares and ownership rates use AIDIS survey weights; medians are weighted medians. Net liquid and illiquid wealth use the benchmark balance-sheet mapping. Income is predicted by the headline gradient-boosting model. The entries are descriptive and are not adjusted for differences in household composition.
Source: Authors' calculations from AIDIS and CPHS.

The other characteristics reinforce the balance-sheet distinction. Nearly all W-HtM households own a home, compared with 2 percent of P-HtM households. W-HtM households are also more likely to have outstanding borrowing and to live in rural areas. These patterns are consistent with portfolios centered on housing and other illiquid assets. They should not be read as causal effects of geography, homeownership, or borrowing on hand-to-mouth status.

7.3 Gold, liquidity, and collateral capacity

Gold can matter without being immediately liquid. Rather than forcing gold into either side of a binary liquid-illiquid boundary, the capacity rule asks the question that matters for consumption smoothing: how much cash or credit can the household actually obtain from the gold it owns? To separate ownership from potential collateral capacity, we reclassify the 56.5 percent of households who are W-HtM when gold is illiquid. A household is labeled collateral-capable when its liquid position plus 75 percent of its gold value exceeds one-half of monthly income, $m_i + 0.75g_i > y_i/2$. The 75 percent fraction is a haircut in the range of regulated gold-loan lending, and the value-based condition means that owning a small piece of jewellery does not by itself confer capacity: holdings too small to span the liquidity gap leave the household NC-W-HtM regardless of ownership.[6] No minimum gold-value floor is

[6]Both parameters are varied. In the CPHS mirror of this split, computed with the piecewise gold prediction of Section 6.3, the headline haircut and zero floor give a collateral-capable share of 17.3 percent of

imposed in the headline calculation.

Table 9: Hand-to-Mouth Status with Gold as Illiquid Wealth

Group	Share of households
P-HtM	3.06
W-HtM, collateral-capable	31.87
W-HtM, not collateral-capable	24.66
N-HtM	40.41

Notes: Entries are weighted percentages of AIDIS households. Gold is included in illiquid wealth. Collateral capacity is defined using a 75 percent gold-value fraction and no minimum-value floor. It measures potential capacity under the implemented rule, not observed gold-loan use or approval.
Source: Authors' calculations from AIDIS and CPHS.

Potential collateral capacity therefore divides the gold-illiquid W-HtM group almost evenly: 31.9 percent of all households are collateral-capable, and 24.7 percent are not. This distinction is useful precisely because ownership is not equivalent to access. The calculation does not observe whether a household seeks a gold loan, is approved, or is willing to pledge the asset.

The capacity rule is the paper's answer to a question that a binary liquid-or-illiquid label cannot pose: how much cash or credit can the household actually obtain from its gold? Rather than declaring gold liquid because it is pledgeable or illiquid because it is jewellery, the rule values each household's holding at the credit it could realistically generate, a haircut fraction θ of value above a pledge floor, and asks whether that credit would move the household off the liquidity kink. Ownership alone therefore confers nothing: a token jewellery holding fails the condition because the credit it can generate is too small to matter, which is exactly the economics of small gold holdings in practice.[7] The split is also not knife-edge in the rule's two parameters. Varying the haircut between 0.5 and 0.85 and the pledge floor between zero and Rs. 10,000 moves the capable share by a few percentage points and cannot move P-HtM, total HtM, or N-HtM at all, because the condition only partitions the existing W-HtM group.[8]

households and a non-capable share of 10.8 percent; moving the haircut between 0.50 and 0.85 and imposing minimum pledge floors up to Rs. 10,000 moves the capable share between 14.7 and 18.0 percent, with the non-capable share absorbing the difference (10.1 to 13.5 percent), because the capacity condition partitions W-HtM and leaves the W-HtM total, the P-HtM share, and total HtM unchanged by construction. A fully continuous treatment, in which every asset carries its own pledgeability and the household solves a portfolio problem over that continuum, would require a different and substantially harder model; the capacity split is the tractable middle ground between a binary boundary and that continuum.

[7] The natural generalization, a continuous liquidity weight on every asset inside a portfolio-choice model, is substantially harder: it requires taking a stand on the joint pricing of pledgeability across assets and turns the classification into a model solution rather than a measurement. The capacity condition is the tractable middle ground: gold stays outside the liquid aggregate, and its realizable credit enters only where it can change the household's constraint status.

[8] On the CPHS side, where the grid is tabulated for the headline route with gold in illiquid wealth under

7.4 Robustness and geographic heterogeneity

Routine changes to the predicted-income measure leave the main result largely intact. The quantile-income draw yields P/W/N shares of 8.1/51.5/40.4, and frequent income yields 8.1/51.3/40.6. Narrow income lowers W-HtM to 46.6 percent, in part because that concept omits resources that are particularly important to rural households. Thus, the W-HtM majority is not a consequence of one fitted-income specification.

Credit access matters, but no implemented alternative eliminates the large HtM population. With a one-month total-income credit limit, total HtM is 59.59 percent. The India-specific institutional limit yields 59.00 percent; the pledge-augmented limit yields 51.98 percent; a two-month limit yields 55.76 percent; and a zero-credit limit yields 61.83 percent. These are economic counterfactuals about available credit, not sampling-error bands. Appendix B.2 describes the implemented limits and the associated classification rules.

The prevalence of W-HtM also varies geographically. In the standardized rural–urban battery, W-HtM is 56.7 percent in rural areas and 39.3 percent in urban areas. The standardized national comparison uses a slightly different common sample and produces P/W/N shares of 7.3/50.6/42.1, so its levels should not be mixed mechanically with the headline estimates. Across states, the W-HtM share ranges from approximately 15 to 72 percent. We treat the state estimates as descriptive rather than ranking states: model-based income, incomplete geographic overlap, and small samples in some states can all contribute to the dispersion.

7.5 What prevalence does and does not establish

The prevalence estimates establish that limited liquidity in India is usually paired with substantial illiquid wealth. They do not establish that the three groups respond differently to an income or monetary-policy shock. That question requires panel evidence with household status held fixed. Section 9 provides it, attaching the fixed 2019 labels of Section 6.3 to the monthly CPHS panel.

8 The Marginal Propensity to Consume by Hand-to-Mouth Status

The local projections of Section 9 measure responses to a common monetary-policy shock, so the identifying variation is aggregate and the informative objects are between-group contrasts. This section asks the complementary micro question: how much of an *idiosyncratic* transitory income change does each group consume? The two designs share the same fixed 2019 classification and the same monthly panel, so together they discipline the same mecha-

the piecewise gold prediction, the collateral-capable share ranges from 14.7 to 18.0 percent and the non-capable share from 10.1 to 13.5 percent across this grid, holding W-HtM fixed at 28.2 percent; on the benchmark boundary used in estimation the headline split is 15.2 and 10.0 percent. Appendix B.2 treats these as implemented sensitivity parameters.

nism from two directions. If the wealthy hand-to-mouth population were hand-to-mouth in name only, its marginal propensity to consume (MPC) should look like that of the unconstrained; if gold-collateral capacity is genuine smoothing capacity, the capacity split should separate MPCs exactly as it separates the wage and consumption responses of the previous section.

8.1 Empirical design

We estimate transitory MPCs with the covariance restrictions of Blundell et al. (2008). Let Δc_{it} and Δy_{it} be residualized log changes of real per-capita consumption and income over a window of h months. Under the standard permanent–transitory income process, the transitory MPC is identified as

$$\psi(h) \;=\; \frac{\mathrm{cov}(\Delta c_{it},\, \Delta y_{i,t+1})}{\mathrm{cov}(\Delta y_{it},\, \Delta y_{i,t+1})}, \tag{4}$$

where t indexes consecutive windows. The denominator is $-(\sigma^2_{\varepsilon} + \sigma^2_{me})$, the transitory-plus-measurement-error variance, so classical income measurement error attenuates ψ toward zero rather than inflating it; short-window estimates are therefore best read as lower bounds on the impact response.

The panel is the full CPHS: 182,816 households survive the estimation filters (accepted consumption responses, positive weights, positive real per-capita income and consumption, at least twenty-four usable months), of which 131,731 carry a fixed 2019 label under the headline participation-mapped classification with the piecewise gold prediction (Section 6.3) — the identical labels used in Section 9. We estimate Equation (4) at four window lengths $h \in \{1, 4, 12, 24\}$ months: monthly, the four-month CPHS survey wave, annual, and biennial. Because $\psi(h)$ cumulates the consumption response over the same window in which the income change is measured, the four estimates trace an intertemporal MPC profile rather than four estimates of one number; comparing levels across windows conflates response accumulation with the mechanical decline of attenuation as income noise averages out. Standard errors come from a stratified household-cluster bootstrap with 500 replications, resampling households within group cells, and group contrasts are computed replication-by-replication so their standard errors account for the covariance between jointly estimated group MPCs. The headline sample excludes March 2020 through June 2021, when lockdown-driven income and expenditure swings dominate the growth moments; full-sample estimates are discussed with the robustness results below.

8.2 MPCs by group

Table 10 reports $\psi(h)$ by group. Three features stand out. First, the pooled MPC rises monotonically with the measurement window, from 0.016 at one month to 0.157 at twenty-four months — the time-aggregation staircase familiar from Crawley (2020): most of the two-year cumulative response is in place within the first waves, and the small monthly

Table 10: Transitory MPCs by Hand-to-Mouth Status

	Monthly	Wave (4mo)	Annual	Biennial	Households
Panel A: cumulative MPC $\psi(h)$					
P-HtM	0.018 (0.003)	0.064 (0.012)	0.114 (0.022)	0.210 (0.033)	4,034
C-W-HtM	0.017 (0.001)	0.047 (0.006)	0.101 (0.011)	0.124 (0.014)	17,143
NC-W-HtM	0.019 (0.002)	0.118 (0.008)	0.167 (0.016)	0.146 (0.018)	9,945
N-HtM	0.015 (0.001)	0.083 (0.003)	0.140 (0.005)	0.165 (0.006)	100,366
All households	0.016 (0.000)	0.079 (0.002)	0.135 (0.005)	0.157 (0.006)	131,488
Panel B: contrasts against N-HtM					
NC-W-HtM − C-W-HtM	+0.003 (0.002)	+0.071 (0.010)	+0.067 (0.019)	+0.022 (0.023)	
NC-W-HtM − N-HtM	+0.004 (0.002)	+0.035 (0.009)	+0.027 (0.017)	−0.019 (0.019)	
C-W-HtM − N-HtM	+0.001 (0.001)	−0.036 (0.007)	−0.040 (0.012)	−0.041 (0.015)	
(P+NC)-HtM − N-HtM	+0.004 (0.001)	+0.019 (0.007)	+0.012 (0.014)	−0.003 (0.017)	

Notes: Entries are transitory MPCs from Equation (4) at four measurement windows, with stratified household-cluster bootstrap standard errors (500 replications) in parentheses; contrast standard errors use the replication-level covariance of the group estimates. The classification is the participation-mapped headline route with the piecewise gold prediction, measured in 2019 and held fixed; the C/NC split uses $\theta = 0.75$ and no gold-value floor. Household counts are for the wave frequency; counts decline mechanically at longer windows. The sample excludes March 2020–June 2021. Columns are not four estimates of one parameter: $\psi(h)$ cumulates the consumption response over an h-month window, and shorter windows are more attenuated by income measurement error. Source: Authors' calculations from AIDIS and CPHS.

estimate reflects the attenuation floor, not the absence of a response. Figure 1 plots the profiles.

Second, at the wave frequency — the survey's natural reporting unit — the collateral-capacity split of Section 7.3 separates the wealthy hand-to-mouth sharply and in opposite directions. W-HtM households without potential gold collateral have an MPC of 0.118, significantly *above* the non-hand-to-mouth value of 0.083 (difference +0.035, standard error 0.009), while collateral-capable W-HtM households sit significantly *below* it at 0.047 (difference −0.036, standard error 0.007). The within-W-HtM contrast, +0.071 with a standard error of 0.010, is the sharpest group difference in the paper. The population that responds as if constrained — the poor plus the wealthy without usable collateral, the same pooled group whose wage income and labor supply carry the local-projection results — has a wave MPC of 0.102 against 0.083 for the unconstrained (difference +0.019, standard error 0.007).

Third, the poor hand-to-mouth cell is too small for precision. Its point estimates sit above the non-hand-to-mouth at one month and twenty-four months and below in between, but with 4,034 households no individual contrast is significant; the informative poor-cell evidence comes from the gold-free subsample below. By the biennial window the groups largely converge, as they should once even constrained households have consumed through a transitory change, and the biennial column inherits the weakest denominator, so we do not lean on it.

Taken at face value, the wave-frequency ordering is not the textbook one: the wealthy hand-to-mouth as a whole sit slightly *below* the unconstrained (difference −0.011, standard error 0.006), driven entirely by the collateral-capable majority of the group. On the logic of the

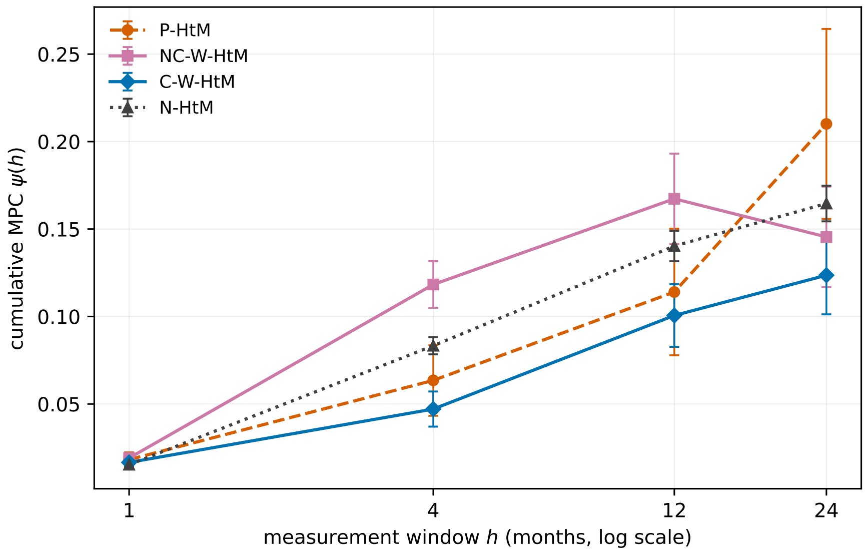

Figure 1: Cumulative MPC Profiles by Hand-to-Mouth Status

Notes: Points are the estimates of Table 10, Panel A, with 90 percent bootstrap intervals, plotted against the measurement window on a log scale. The rising profiles mix response accumulation with declining measurement-error attenuation, so the informative comparisons are between groups at a given window, not across windows for a given group. Source: Authors' calculations from AIDIS and CPHS.

previous section, that is what gold-backed liquidity should produce: a household that can convert gold into cash at short notice can smooth a transitory income change even though its measured liquid wealth is negligible. The classification exercises below test this reading directly.

8.3 The gold margin tested twice

8.3.1 Households without gold

The piecewise gold prediction implies that 20,842 classified households (17.5 percent, matching the AIDIS non-owner share by construction) own no gold at all. For these households the gold margin cannot operate, so the textbook ordering should reappear. Table 11 shows that it does, strongly. Among predicted non-owners every hand-to-mouth MPC exceeds the non-hand-to-mouth MPC at the monthly, wave, and annual windows: the pooled HtM−N difference is +0.030 (standard error 0.009) at one month, +0.090 (0.020) at the wave, and +0.064 (0.028) at the annual window, converging by the biennial window. MPC levels are roughly double the full sample's at short windows, consistent with a poorer and less-buffered subpopulation. The poor cell is thin here (686 households) and its contrasts, though positive at every window, are individually imprecise; the wealthy cell carries the precision.

Table 11: Transitory MPCs among Predicted Gold Non-Owners

	Monthly	Wave (4mo)	Annual	Biennial	Households
P-HtM	0.030 (0.022)	0.165 (0.048)	0.264 (0.063)	0.361 (0.097)	685
W-HtM	0.074 (0.009)	0.198 (0.018)	0.292 (0.027)	0.268 (0.035)	4,280
N-HtM	0.038 (0.005)	0.104 (0.011)	0.225 (0.014)	0.304 (0.016)	15,787
All	0.046 (0.004)	0.132 (0.010)	0.242 (0.012)	0.298 (0.015)	20,752
HtM − N-HtM	+0.030 (0.009)	+0.090 (0.020)	+0.064 (0.028)	−0.023 (0.036)	
W-HtM − N-HtM	+0.036 (0.010)	+0.094 (0.020)	+0.067 (0.030)	−0.036 (0.039)	

Notes: Same estimator, sample rules, and fixed 2019 labels as Table 10, restricted to households the piecewise gold prediction classifies as owning no gold (17.5 percent of classified households, matching the AIDIS weighted non-owner share by construction). Because no household in this subsample has gold, the collateral-capacity split is degenerate and the three-way classification is reported. Household counts are for the wave frequency. The sample excludes March 2020–June 2021. Source: Authors' calculations from AIDIS and CPHS.

8.3.2 Moving gold across the balance-sheet boundary

The second test applies each alternative balance-sheet boundary of Table 6, re-classifies every household, and re-estimates the MPCs, holding the estimator, the panel, and the gold prediction fixed. Table 12 reports the wave-frequency contrasts. When gold (alone, or together with other physical assets) is counted as liquid, the households still classified as hand-to-mouth are those without a gold buffer, and the textbook ordering appears with force: the hand-to-mouth MPC exceeds the non-hand-to-mouth MPC by +0.058 (standard error 0.012) when gold and property are liquid and by +0.042 (0.009) when gold alone is. When gold is instead illiquid, the wealthy hand-to-mouth pool absorbs the gold-rich, the aggregate contrast turns slightly negative — and the collateral-capacity split resolves it: the within-W-HtM contrast lies between +0.065 and +0.071 across definitions. Whether the textbook MPC ordering appears is governed entirely by which side of the liquidity boundary gold occupies. That is the MPC counterpart of the prevalence result that the asset boundary is an economic object rather than a bookkeeping choice, and it is the same conclusion the gold-free subsample delivers without any reclassification.

8.4 Robustness and interpretation

Including the March 2020–June 2021 window shrinks but does not overturn the capacity results: the within-W-HtM contrast at the wave frequency is +0.043 (standard error 0.008) with the COVID months included, against +0.071 without, and the gold-free subsample's pooled wave contrast is +0.039 (0.016) against +0.090. The NC-W-HtM minus N-HtM contrast loses significance when the COVID months are included, which is why we report the exclusion as the headline: the lockdown moves income and expenditure together at enormous scale, and the resulting growth moments say more about the lockdown than about transitory-income pass-through.

Three limitations parallel those of the previous section. The standard errors condition on the

Table 12: Wave-Frequency MPC Contrasts under Alternative Balance-Sheet Boundaries

Definition	HtM − N-HtM	W-HtM − N-HtM	NC − C (within W)
KVW benchmark (headline)	−0.012 (0.005)	−0.011 (0.006)	+0.071 (0.010)
Gold included in illiquid wealth	−0.012 (0.005)	−0.013 (0.006)	+0.065 (0.009)
Gold included in liquid wealth	+0.042 (0.009)	+0.045 (0.010)	—
Broader physical assets liquid	+0.058 (0.012)	+0.063 (0.014)	—
Broad illiquid wealth and debt	−0.014 (0.006)	−0.015 (0.005)	+0.052 (0.009)

Notes: Entries are wave-frequency MPC differences with bootstrap standard errors in parentheses. Rows follow the definition ladder of Table 6, using the piecewise gold prediction throughout; each row re-classifies households and re-estimates Table 10's Panel A. The real-estate screen row is omitted: in CPHS it re-classifies almost no households (Table 6). The collateral-capacity column applies the $\theta = 0.75$, no-floor capacity rule within W-HtM where gold is outside liquid wealth. Under the gold-liquid definitions the hand-to-mouth cells are small (roughly 6,000 and 9,000 households); their poor cells are too thin to report separately. The sample excludes March 2020–June 2021. Source: Authors' calculations from AIDIS and CPHS.

predicted classification and do not propagate uncertainty from the wealth and gold models. The gold-free subsample is model-selected — predicted rather than observed non-owners — although the identical household list drives the labor-market results of the long-horizon local projections. And short-window estimates are attenuated by income measurement error, which is most severe for irregular agricultural incomes; this is the natural reading of the near-zero monthly differences and of the poor cell's flat monthly estimate.

Subject to those limitations, the MPC evidence closes the loop opened by the classification. Hand-to-mouth status predicts high transitory MPCs exactly where the balance sheet says wealth is unusable: among households without gold, among households whose gold is too small to pledge, and under any classification that does not let gold masquerade as a buffer it cannot provide — while potential gold collateral marks the one hand-to-mouth subgroup that smooths like the unconstrained. The same capacity margin that separates wage and consumption responses to a common monetary shock separates marginal propensities to consume out of idiosyncratic income, at magnitudes (0.05–0.12 at the wave frequency, 0.10–0.17 at the annual) squarely within the range the quasi-experimental literature reports for constrained households.

9 Gold Collateral Capacity and Monetary Transmission

Two transmission differences emerge under the participation-mapped classification. In the first half-year after a contractionary target surprise, P-HtM households cut consumption by more than W-HtM households: the difference reaches −1.90 log points per unit surprise at month four (standard error 0.89), or −0.17 log points for a one-standard-deviation monthly surprise (standard error 0.08). This three-way contrast is estimable at all only because the participation map restores the P-HtM cell that amount-only predictions empty. The second difference lies within the wealthy hand-to-mouth group and shows first in labor income:

from month six onward the wage income of W-HtM households with potential gold-collateral capacity holds up relative to W-HtM households without such capacity, by 4.6 log points per unit surprise at month six (standard error 2.2) and 7.1 (3.0) at month eleven. The corresponding consumption difference is positive in the second half-year but individually imprecise within twelve months, 5.4 log points per unit surprise at month seven (standard error 5.5); at the longer horizons of Appendix D it peaks at month thirteen at 9.5 (5.1). Beyond these two contrasts, the group comparisons do not produce a stable full ordering, so we organize this section around the comparisons that the data support.

9.1 Empirical design

We combine the fixed 2019 CPHS classifications from Section 6.3 with the target and path surprises of Lakdawala and Sengupta (2025). The estimation sample begins in 2014 and ends in February 2023, the final month in the shock series. Months without an RBI announcement receive a zero surprise. The headline specification enters the target and path factors jointly using the one-day event window.

Following Jordà (2005) and the panel practice in Jordà and Taylor (2025), for each horizon $h \in \{0, \ldots, 12\}$ we estimate a household fixed-effects local projection of the form

$$\begin{aligned} y_{i,t+h} - y_{i,t-1} = \sum_{g} \sum_{k \in \{T,P\}} \beta_h^{g,k} s_t^k \mathbf{1}\{G_i = g\} + \sum_{\ell=1}^{3} \rho_{h,\ell} \Delta y_{i,t-\ell} \\ + \sum_{k \in \{T,P\}} \sum_{\ell=1}^{3} \pi_{h,\ell}^k s_{t-\ell}^k + \delta_{m(t),h} + \alpha_{i,h} + \varepsilon_{i,t+h}, \end{aligned} \tag{5}$$

where s_t^T and s_t^P are the target and path surprises, G_i is the fixed household group, and $\delta_{m(t),h}$ denotes month-of-year effects. For consumption and income, y is 100 times the logarithm of a positive monthly amount, so the dependent variable is a cumulative log change. For employment, y is the percentage of household members age 15 or older who are employed. Estimation is unweighted. Standard errors are clustered by calendar month, the level at which the shock varies. The millions of household-month observations improve measurement of group outcomes, but the identifying time-series variation comes from roughly one hundred months.

The shock is common to every household in a month, so calendar-month fixed effects would absorb it. This makes the between-group contrast more informative than any group response in isolation. Our headline estimands are therefore $\beta_h^{P,T} - \beta_h^{W,T}$ and $\beta_h^{CW,T} - \beta_h^{NCW,T}$. We estimate the group coefficients jointly and uses their covariance to form the standard error of each difference. The headline classification is the participation-mapped direct route of Section 6.3; the quantile-mapped route is retained as a robustness classification.

Group status is held fixed at its 2019 value. This prevents the policy shock from changing the classification mechanically, but it does not make status predetermined throughout the sample. For shocks before 2019, the label is a subsequent measure of portfolio type. For shocks after 2019, it is a prior measure. We interpret the fixed label as a proxy for a persistent portfolio position and do not claim that the observed 2019 balance sheet predates every shock.

9.2 The poor hand-to-mouth and consumption

Figure 2 plots the P-HtM minus W-HtM cumulative total-consumption response to a one-standard-deviation monthly target surprise under the participation-mapped classification. The difference is close to zero on impact, turns negative immediately, and remains between -0.12 and -0.17 log points over months one through five, with 90 percent pointwise confidence intervals that exclude zero at each of those horizons. The largest gap is -0.17 log points at month four (standard error 0.08), or -1.90 log points per unit target surprise (standard error 0.89). The difference fades in the second half-year and is indistinguishable from zero by month seven.

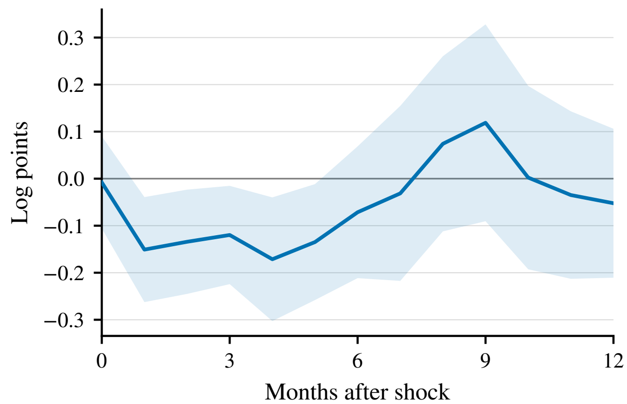

Figure 2: P-HtM Minus W-HtM Consumption Response to a Target Surprise

Notes: The line is the P-HtM minus W-HtM cumulative total-consumption response to a one-standard-deviation monthly target surprise. Negative values indicate a larger consumption cut among P-HtM households. The shaded area is the 90 percent pointwise confidence interval based on standard errors clustered by calendar month. The classification uses observed CPHS income and participation-mapped predicted wealth; status is measured in 2019 and held fixed. The target and path factors enter jointly using the one-day event window.

The pattern matches the economics of the classification. P-HtM households hold neither liquid nor illiquid buffers, so a contractionary surprise should pass through to their consumption faster and harder than for W-HtM households, whose illiquid wealth provides at least some scope for adjustment. The individual group paths are noisy, since the P-HtM cell contains only 4,035 households, but the contrast is estimated tightly because both coefficients come from the same regression and share the same monthly shocks. The corresponding P-HtM minus N-HtM difference has the same sign and a larger magnitude at month four but is much less precise (Appendix C.3).

9.3 Gold capacity and consumption

The participation-mapped classification divides W-HtM households according to the gold-capacity rule used in Section 7.3. A W-HtM household is C-W-HtM when its net liquid position plus 75 percent of piecewise-predicted gold exceeds one-half of monthly income; otherwise it is NC-W-HtM. The rule uses no minimum gold-value floor. It measures potential collateral capacity, not observed borrowing or approval for a gold loan.

Figure 3 plots the C-W-HtM minus NC-W-HtM response directly. A positive value means that the collateral-capable group has a less negative response. All three panels scale the target surprise by its monthly standard deviation, 0.090, over the 2014–February 2023 overlap. The consumption difference is close to zero on impact, briefly negative in months two through four, and positive in the second half-year, reaching 0.49 log points at month seven (5.38 log points per unit target surprise, standard error 5.48). Within twelve months the consumption difference is individually imprecise at every horizon; the long-horizon estimates of Appendix D show it peaking just past the first year, at 9.48 log points per unit surprise at month thirteen (standard error 5.08), before decaying. The sharper twelve-month evidence for the capacity margin is in wage income, below.

Figure 3: Gold Collateral Capacity and Responses to a Target Surprise

Notes: Each panel reports the C-W-HtM minus NC-W-HtM cumulative response to a one-standard-deviation monthly target surprise. Positive values indicate a less negative response among W-HtM households with potential gold-collateral capacity. Shaded areas are 90 percent pointwise confidence intervals based on standard errors clustered by calendar month. The classification uses observed CPHS income, participation-mapped predicted wealth, predicted gold, $\theta = 0.75$, and no gold-value floor; status is measured in 2019 and held fixed. The target and path factors enter jointly using the one-day event window. The three outcomes have different vertical scales. Employment is measured in percentage points.

The individual C-W-HtM and NC-W-HtM response paths are imprecise. The contrast is estimated more tightly because both coefficients are identified by the same monthly shocks and their estimation errors are positively correlated. We therefore read Figure 3 as evidence about a response difference, not as precise evidence about the level response of either group.

9.4 Income and employment

The income response moves in the same direction as consumption, with the wage component carrying the precision. The C-W-HtM minus NC-W-HtM difference in wage income is 4.57 log points per unit target surprise at month six (standard error 2.15) and 4.75 (1.92) at month seven, and it widens to 7.09 (3.01) at month eleven and 7.08 (3.20) at month twelve. The total-income difference has the same sign but is imprecise within the year, 4.22 log points per unit surprise at month ten (standard error 3.11); at the longer horizons of Appendix D it peaks at 7.90 (3.64) at month seventeen. These estimates are consistent with a cash-flow difference between the two groups.

Employment does not display the same separation within the first year. At month twelve, the C-W-HtM minus NC-W-HtM difference is 0.95 percentage points per unit target surprise (standard error 1.21), equivalent to 0.09 percentage points for a one-standard-deviation surprise (standard error 0.11). We construct employment as the share of household members age 15 or older who are employed. Because CPHS observes that outcome approximately once per survey wave, we carry it forward for at most three months. The employment result is therefore a coarse mechanism check, not a precisely measured monthly response.

The corresponding P-HtM minus W-HtM differences in income and employment are mostly imprecise and alternate in sign over the horizon. The one exception is a briefly negative wage difference around month three, so a differential income decline cannot be ruled out as a contributor to the early P-HtM consumption gap; the employment difference is never statistically distinguishable from zero.

The joint pattern of a consumption difference, a similarly signed income difference, and little employment separation is consistent with potential gold liquidity cushioning household cash flow. It does not establish that households borrow against gold. The data do not observe applications, approval, loan proceeds, or the willingness to pledge the asset. Gold capacity may also proxy for other persistent differences within the W-HtM population.

9.5 Robustness and interpretation

Table 13 shows which parts of the two consumption contrasts survive the implemented alternatives. The gold-capacity consumption difference at month seven is 0.49 log points under the headline classification with the corrected piecewise gold, 0.72 under the quantile-mapped classification, and 0.48 excluding March 2020 through June 2021; none of the month-seven consumption entries under the corrected gold measure is individually significant, and the two-day event window reduces the estimate to 0.21 with a standard error more than twice its size. The first-principal-component shock produces a positive 0.50 log-point difference that is marginally significant at the 10 percent level. The month-seven consumption contrast is therefore positive across every implemented alternative but not individually significant under the corrected gold measure; the capacity margin's statistically sharp results are the wage contrast above and the mid-horizon estimates in Appendix D.

The P-HtM minus W-HtM difference at month four keeps its negative sign in every implemented alternative but retains 5 percent significance only in the headline specification. Excluding the COVID window attenuates it to −0.12 log points (standard error 0.07); the two-day window and the principal-component shock leave estimates of similar or larger relative magnitude that are less precise. The quantile-mapped entry is shown for completeness but is not informative: that route assigns roughly 0.1 percent of households to P-HtM, so the contrast is estimated from a nearly empty cell.

Table 13: Robustness of the Consumption Contrasts

Specification	P-HtM minus W-HtM (month 4)	C-W-HtM minus NC-W-HtM (month 7)
Participation-mapped headline	−0.171 (0.080)	0.485 (0.494)
Quantile-mapped classification	−0.587 (0.347)	0.716 (0.369)
Exclude March 2020–June 2021	−0.123 (0.072)	0.477 (0.514)
Two-day event window	−0.185 (0.132)	0.208 (0.479)
First principal component	−0.118 (0.074)	0.496 (0.295)

Notes: Entries are differences in the cumulative total-consumption response, in log points, with standard errors in parentheses. Each shock is scaled by its monthly standard deviation over the 2014–February 2023 overlap. Standard errors use the covariance between the jointly estimated group coefficients and are clustered by calendar month. The quantile-mapped P-HtM cell contains roughly 0.1 percent of households, so its P-HtM minus W-HtM entry is reported only for completeness. These are pointwise, not simultaneous, comparisons across horizons.
Source: Authors' calculations from AIDIS and CPHS.

The alternative synthetic-series estimator is uninformative on this margin at month seven under the corrected split: its unscaled C-W-HtM response is −18.4 log points against −17.7 for NC-W-HtM, with standard errors several times those magnitudes. The synthetic-series implementation does not report a covariance-correct standard error for their difference, so we treat it as an appendix diagnostic rather than independent confirmation of the headline result. The path factor and the remaining three-way P/W/N comparisons are reported as diagnostics in Appendix C. Food and appliance outcomes are estimated but do not support the headline interpretation.

Finally, months four and seven were not selected before examining the response paths, and the displayed bands are pointwise across thirteen horizons. The figures should therefore be read as dynamic patterns rather than as single pre-specified hypothesis tests. The reported standard errors also condition on the predicted classification and do not propagate uncertainty from the wealth and gold models. Subject to those limitations, the estimates show that both margins of the classification matter: the absence of any wealth buffer separates the early consumption response of poor from wealthy hand-to-mouth households, and within the wealthy group, potential access to gold-backed funds separates consumption and income responses even when employment responses remain similar.

The group paths, though individually imprecise, line up with this reading, with the timing now explicit: within the first year the separation runs through the poor and through wages. Over months two through eight, mean consumption of P-HtM households falls relative to N-HtM by 0.09 log points per one-standard-deviation surprise, while the two wealthy subgroups both track N-HtM in consumption even as their wage paths separate; the consumption separation between the wealthy subgroups opens just past the first year and peaks at month thirteen (Appendix D). The population that responds as if constrained is the poor plus the wealthy without usable collateral: 27.7 percent of households under the gold-illiquid classification, of which the poor in the balance-sheet sense are only 3.1 points. A framework that stopped at poor hand-to-mouth, or that treated all wealth as equally usable, would miss most of that population.

Appendix D tests this reading directly and extends it in two directions. Pooling P-HtM and NC-W-HtM into a single constrained group and re-estimating the projections sharpens the labor results: the pooled group's wage income falls significantly relative to both the capable wealthy and the non-hand-to-mouth, its total-income and consumption gaps against the capable wealthy are widest at mid horizons (months thirteen to seventeen) and largely closed by month forty-eight, and the pooled group's relative employment rises two to three years out, consistent with self-insurance through labor supply where collateral is unavailable. A placebo that sorts households by income at the same group shares reproduces the shape of these responses but roughly doubles their standard errors, so the balance-sheet classification is not a relabeled income distribution.

10 Conclusion

Most hand-to-mouth households in India are not poor in the balance-sheet sense. Under the KVW benchmark definition, 59.6 percent of households hold too little liquid wealth to span one-half of monthly income. Yet only 8.2 percent of households are P-HtM; 51.4 percent are W-HtM. Wealthy hand-to-mouth households therefore account for 86 percent of India's hand-to-mouth population. These households own homes, land, or other illiquid assets but cannot readily use those assets for current spending. A classification based only on income or net worth misses this portfolio distinction. The central measurement result is thus not simply that liquidity is scarce. It is that limited liquidity commonly coexists with substantial wealth.

We reach this result by assigning different empirical roles to two household surveys. AIDIS anchors national prevalence because it observes detailed balance sheets. CPHS supplies monthly income, expenditure, and employment needed to study household dynamics. Neither survey jointly observes all of these quantities, so we predict the variable that each survey omits rather than describe a synthetic panel in which every variable is observed. Prediction at the classification threshold requires more than fitting conditional means. On held-out AIDIS households, conditional-mean gradient boosting classifies 37.5 percent as hand-to-mouth, 21.3 percentage points below the observed share of 58.8 percent. Distribution-preserving methods come much closer to the observed group shares, and only a prediction that also maps the participation margins of assets and liabilities onto their AIDIS distributions preserves the poor hand-to-mouth group in deployment. This sensitivity is not an incidental modeling detail: when an economic group is defined by a threshold, the predicted distribution can matter as much as average predictive fit.

Gold makes the asset boundary economically consequential. Treating gold as illiquid wealth raises the W-HtM share to 56.5 percent while leaving total HtM unchanged; treating it as liquid wealth reduces total HtM from 59.6 to 24.6 percent. Under the implemented collateral-capacity rule, 31.9 percent of all households are W-HtM with potential gold capacity and 24.7 percent are W-HtM without it. Both classification margins appear in the monetary-policy application. Four months after a contractionary target surprise, poor hand-to-mouth households, restored by the participation-mapped prediction, cut consumption 0.17 log points more than wealthy hand-to-mouth households for a one-standard-deviation monthly shock. Within the wealthy group, the wage income of the collateral-capable W-HtM group holds up relative to the group without capacity, by 0.41 log points at month six and 0.64 at month eleven; the consumption difference is positive but imprecise within the first year and peaks just past it. The corresponding employment difference at month twelve is 0.09 percentage points and imprecise. The joint pattern is consistent with the absence of any buffer accelerating pass-through for the poor hand-to-mouth and with potential access to gold-backed liquidity cushioning cash flow within the wealthy group. It does not show that households apply for, receive, or use gold loans. Counting only the poor hand-to-mouth as constrained would nevertheless miss most of this responsive population: the wealthy without capacity outnumber the poor by roughly eight to one.

The classification exercises also admit a counterfactual reading that connects them to financial inclusion. Each rung of the placement ladder and each credit-limit alternative answers a question of the form: how many wealthy hand-to-mouth households would there be if households could actually use a given asset for current spending? If gold became fully usable, total HtM would fall from 59.6 to 24.6 percent; a credit limit augmented by pledgeable gold lowers total HtM from 59.6 to 52.0 percent; and improvements in property collateralization would operate on the same margin for land and housing. Read this way, the gap between the KVW benchmark and the most liquid rung is a measure of the consumption insurance that deeper collateral markets and broader financial inclusion could unlock.[9]

The classification ladder also admits a counterfactual reading that speaks to policy. Each placement case answers a question of the form: what would hand-to-mouth prevalence be if households could actually use this asset for current spending? If gold became fully spendable, total HtM would fall from 59.6 to 24.6 percent; if pledgeable gold merely augmented the credit limit, total HtM would fall to 52.0 percent. The distance between these numbers and the benchmark measures, in shares of households, what better collateralization infrastructure could accomplish. Several current policy margins map directly into the exercise. Digitisation of land and property records could raise the effective pledgeability of real estate, though records alone do not create the secondary markets and willing lenders that effective liquidity requires. Financial inclusion operates on the same margin from the credit side. And for gold

[9]India's gold-loan market makes this margin unusually concrete: disbursal against ornaments is nearly instantaneous at regulated non-bank lenders, so the effective fixed cost of accessing the collateral is close to zero, while interest spreads are large. The Reserve Bank of India's 2025 directions on lending against gold and silver collateral replaced the uniform 75 percent loan-to-value ceiling with a tiered schedule of 85, 80, and 75 percent by loan size. That is a dated, legislated change in exactly the haircut parameter of our capacity rule, differential by loan size, and studying its consumption effects through the lens of the classification developed here is a natural next step that we leave to future work.

the relevant institutional parameter is legislated: gold lending in India combines near-zero origination costs with wide spreads, and the loan-to-value ceiling that regulated lenders face is precisely the haircut θ of our capacity rule.[10]

The evidence has four limits. Wealth and gold are predicted in CPHS rather than observed. Household status is measured in 2019 and held fixed, so it postdates the monetary surprises used before 2019. The month-four and month-seven consumption comparisons were not selected in advance, and the reported confidence intervals are pointwise across horizons. Finally, the local- projection standard errors condition on the predicted classification rather than propagate uncertainty from the wealth and gold models. These limits rule out a welfare claim or a definitive causal verdict on gold access. They do not alter the paper's measurement lesson. In an economy where households hold much of their wealth in housing, land, and gold, the usability of wealth matters as much as the amount households own, both for identifying who is hand-to-mouth and for interpreting how monetary policy reaches them.

References

Almuzara, Martín and Victor Sancibrián, "Micro Responses to Macro Shocks," Staff Report 1090, Federal Reserve Bank of New York 2024.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, *109* (6), 2333–2367.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

Bracco, Jessica, Luciana Galeano, Pedro Juarros, Daniel Riera-Crichton, and Guillermo Vuletin, "Social transfer multipliers in developed and emerging countries: The role of hand-to-mouth consumers," *The World Bank Economic Review*, 2026, p. lhag024.

Centre for Monitoring Indian Economy, "Consumer Pyramids Household Survey," Consumer Pyramids Dx 2025. Household panel data through December 2025, accessed under license.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen, "Identifying Wealthy Hand-to-Mouth Households in Belgium," Working Paper 419, National Bank of Belgium 2023.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico, "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism," *Review of Economic Studies*, 2020, *87* (1), 102–129.

[10] The Reserve Bank of India's 2025 directions on lending against gold and silver collateral replaced the uniform 75 percent loan-to-value ceiling with a tiered schedule, 85 percent for the smallest loans, stepping down to 75 percent for the largest. A dated, legislated, loan-size-differentiated change in θ is a natural experiment for the capacity mechanism, and we regard it as a promising design for future work rather than part of this paper's evidence.

Crawley, Edmund, "In Search of Lost Time Aggregation," *Economics Letters*, 2020, *189*, 108998.

Hara, Kengo, Takashi Unayama, and Justin Weidner, "The Wealthy Hand-to-Mouth in Japan," RIETI Discussion Paper Series 16-E-073, Research Institute of Economy, Trade and Industry 2016.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

_ **and Alan M. Taylor**, "Local Projections," *Journal of Economic Literature*, 2025, *63* (1), 59–110.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

_ **, Giovanni L. Violante, and Justin Weidner**, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 77–153.

Lakdawala, Aeimit and Rajeswari Sengupta, "Measuring Monetary Policy Shocks in Emerging Economies: Evidence from India," *Journal of Money, Credit and Banking*, 2025, *57* (2-3), 407–437.

National Statistical Office, "All India Debt & Investment Survey – 2019," Technical Report NSS Report No. 588 (77/18.2), Ministry of Statistics and Programme Implementation, Government of India 2021.

Song, Seong-Yoon, "Leverage, Hand-to-Mouth Households, and Heterogeneity of the Marginal Propensity to Consume: Evidence from South Korea," *Review of Economics of the Household*, 2020, *18* (4), 1213–1244.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

A Data Appendix

This appendix records the sample rules, reference periods, units, and cross-survey concordance used in estimation. The purpose is to make clear which quantities are observed, which are constructed, and which are predicted. It does not treat similarly named variables in AIDIS and CPHS as interchangeable.

A.1 Sample construction and empirical roles

Table 14 follows each source from its released unit to its role in the paper. The distinction between weighted cross-sectional estimates and unweighted panel regressions is deliberate. AIDIS supplies national prevalence because it observes balance sheets and provides National Sample Survey multipliers. CPHS supplies household dynamics; the monetary local projections compare households within that estimation sample.

Table 14: Sample Construction and Empirical Roles

Source	Implemented sample rule	Final unit and size	Weight and role
AIDIS Visit 1	Start with 116,461 households; drop 24 whose household head is in the unmatched transgender category	2019 household cross-section; 116,437 households	NSS multiplier divided by 100; prevalence anchor and wealth-model training sample
CPHS bridge cross-section	Require asset ownership in wave 14 and at least eight of twelve 2019 months in each of the consumption, income, and people modules	2019 household cross-section; 131,789 households	Released CMIE household weight; cross-survey prediction and 2019 labels
CPHS monthly application	Retain accepted responses and one observation per household and month; attach one fixed household classification	Household-month panel over the available outcome and shock overlap	Unweighted household fixed-effects local projections; within-sample dynamics
Monetary surprises	Sum announcement-level surprises by month and assign zero in nonannouncement months	Monthly target and path factors; application window 2014–February 2023	Common monthly regressors; target and path enter jointly

Notes: CPHS wave 14 was fielded from May through August 2018. The shock file begins in November 2003, but the household regressions use only the CPHS overlap beginning in 2014. The monthly panel size varies with the outcome, horizon, and availability of positive values after the log transformation.
Source: Authors' documentation of the estimation samples.

A.2 Reference periods and variable concordance

The surveys align most closely on their 2018 asset information and 2019 household characteristics. AIDIS asset stocks refer to June 30, 2018. CPHS asset-ownership indicators come from May–August 2018. AIDIS outstanding debt is measured at the 2019 interview, whereas CPHS income, expenditure, and demographics are aggregated over observed 2019 reference months. Table 15 states the operational concordance.

A.3 Units, weights, and missing values

AIDIS asset and debt amounts remain nominal rupee stocks. Repayments since July 1, 2018 are converted to monthly flows using the exact number of days between that date and the interview. CPHS continuous bridge variables are household means over observed 2019 months. Binary variables record any report within the window, and categorical variables use the first observed value. Cross-sectional AIDIS and CPHS statistics use their respective survey weights. The monetary-policy regressions are unweighted.

The local-projection outcomes follow the released CPHS measures rather than the cross-survey resource concept. Positive nominal expenditure and income are transformed to 100 times their natural logarithm. Zero and negative values become missing under this transformation. Employment is the percentage of household members age 15 or older who are employed. Because the people module is observed approximately once per four-month wave, we carry the employment value forward for no more than three months. No seasonal adjustment, CPI deflation, or EMI subtraction is applied to the released expenditure outcomes used in those regressions.

A specific pattern deserves note: some CPHS households report a value of exactly zero in every month of a four-month survey wave, which is not economically plausible for consumption and rarely plausible for income, and more likely reflects nonresponse recorded as zero. We quantified the pattern among household-months with Accepted response status, the only months our analyses use. For consumption it is negligible: 13 of 3,625,045 fully observed household-waves are zero in all four months, and zero months are 0.007 percent of observed months. For income it is small but real: 3,212 household-waves (0.089 percent) are zero in all four months, with the incidence concentrated in 2014 through 2016 and declining to near zero by 2024; single zero months spike in 2020, consistent with genuine lockdown income losses rather than coding. These observations are excluded from estimation mechanically, because the log transformation turns zeros into missing values and the 2019 income means used elsewhere condition on positive income. We flag the pattern for data-cleaning purposes: a future revision will recode full-wave zeros as missing at the cleaning stage rather than relying on downstream transformations to exclude them.

Table 15: Cross-Survey Variable Concordance

Object	AIDIS construction	CPHS construction	Harmonized use
Asset position	Item-level rupee stocks at June 30, 2018	Wave-14 ownership indicators; no rupee amounts	Ownership margins are common predictors; amounts are predicted in CPHS
Debt position	Loan-level outstanding amount at the interview date, source, purpose, and repayment	Borrowing indicators but no complete balance-sheet amount	Indicators enter prediction; AIDIS amounts determine the training targets
Income	Not observed; predicted as a positive monthly amount	Mean of observed 2019 monthly income; total income is the baseline	Monthly income enters the HtM threshold
Consumption	Purchases, self-production, in-kind receipts, and one-twelfth of annual durables	Total expenditure net of EMIs plus self-production for the bridge	Monthly resource measure used as a predictor; released nominal expenditure is used in local projections
Demographics	Household roster and characteristics at interview	Continuous variables averaged, binary variables set to any report, and categorical variables set to the first observed 2019 value	Common household size, age composition, head characteristics, and social categories
Employment	Member-level arrangement and industry mapped to NSS work categories	Member-level employment and industry mapped to the same categories	Common prediction covariate; the panel outcome is the employed share among members age 15 or older
Geography	NSS state and district codes	CMIE state and district names	Verified district aliases where possible; otherwise state-level pooling for prediction

Notes: CPHS binary cross-sectional indicators equal one when the household reports the item in any observed month. AIDIS positive-value indicators equal one when the recorded rupee amount is positive. This concordance aligns economic objects but does not eliminate survey-instrument differences.

A.4 Validation checks and remaining differences

The internal AIDIS checks agree closely across files. All 94,023 households with positive owned-and-possessed land in the household block have a member-level ownership record, and 99.6 percent also have a plot-level record. The residual plots are principally homesteads with a median area of 0.02 acres. These checks validate the linkage of the underlying blocks; they do not validate the market value reported for every asset.

The largest cross-survey discrepancy lies in consumption measurement. Mean monthly consumption is Rs. 10,048 in AIDIS and Rs. 12,092 in CPHS. The gap is concentrated in rural households; the weighted urban means are approximately equal. AIDIS consumption is also rounded: 48.9 percent of positive values are multiples of Rs. 100, compared with 0.2 percent of CPHS household means. The implemented multiple-imputation jitter exercise changes income-prediction attenuation by approximately 0.2 percent and leaves the HtM classification nearly unchanged.

Geographic support is incomplete. The crosswalk matches 498 of 683 AIDIS districts exactly or through a verified alias. Nine small states and union territories, representing 0.77 percent of AIDIS households, do not appear in CPHS. For prediction we map those observations to a neighboring covered state while preserving the original state for descriptive statistics; a common-states robustness exercise drops them.

Finally, we compare AIDIS asset, debt, and investment aggregates with Reserve Bank of India series and CPHS consumption with the National Sample Survey benchmark. These are order-of-magnitude diagnostics. They do not identify a cross-survey scaling factor, and we do not rescale either survey to match an external aggregate. The evidence supports ranks, ownership margins, and group shares under stated definitions more strongly than exact cross-survey level equivalence, predicted balance-sheet tails, or comparisons involving very small groups.

B Classification and Prediction Appendix

This appendix reports the construction detail and robustness exercises for classification and prediction. Every description refers to estimators and parameter choices as implemented, and every table corresponds to a computed result used in the paper.

B.1 Complete balance-sheet map

The benchmark liquid-asset aggregate contains financial assets admitted by the corrected AIDIS item map. Benchmark liquid debt contains borrowing for household expenditure, education, medical costs, and litigation, together with Kisan Credit Card balances. The KCC amount is removed from the purpose-based totals before the aggregates are added, so the same debt is not counted twice. Benchmark illiquid assets contain the housing and real-estate components admitted by the KVW mapping, together with corrected illiquid financial claims. Benchmark illiquid debt is housing debt.

Four item-level decisions are imposed in every specification. Life- insurance sum assured is excluded from wealth. Cooperative-society shares are illiquid. Kisan Credit Card balances are liquid debt. Post-office saving certificates are outside the benchmark, enter liquid assets in the physical- asset exercise, and enter illiquid wealth in the broad exercise. Gold, art, transport equipment, livestock, and machinery are outside both benchmark asset aggregates.

The active definition ladder is:

1. *KVW benchmark*: benchmark liquid assets and debt; benchmark illiquid assets and housing debt.
2. *Physical-liquid*: gold, transport equipment, livestock, and machinery are added to liquid assets; the benchmark debt and illiquid aggregates are retained.
3. *Gold-liquid*: gold is added to liquid assets; the other benchmark aggregates are retained.
4. *Gold-illiquid*: gold is added to illiquid assets; the benchmark liquid position is retained.
5. *Broad*: benchmark liquid assets are net of broad liquid debt, and broad illiquid assets are net of broad illiquid debt.

Each net position is built from mutually exclusive components. Appendix A.2 records the concordance of economic objects and reference periods across AIDIS and CPHS.

B.2 Classification robustness

The headline rule uses the union of the zero-kink and borrower branches in equation (1). We also store each branch separately. For a general credit limit $\bar{m}_i$, the zero-kink branch is invariant to the credit-limit assumption, while the borrower branch moves with $\bar{m}_i$. Current robustness runs compare limits equal to zero, one month of income, and two months of income; an India-specific institutional limit combines account overdraft and Kisan Credit Card rules; and a pledge-augmented limit adds a fraction of gold and provident-fund wealth. These schemes are not pooled into one range because they represent different economic access to credit.

The net-worth robustness replaces m_i in the liquidity rule with $n_i = m_i + a_i$:

$$H_i^{NW} = \mathbf{1}\left\{0 \leq n_i \leq \frac{y_i}{2}\right\} \ \vee \ \mathbf{1}\left\{n_i \leq 0 \text{ and } n_i \leq \frac{y_i}{2} - \bar{m}_i\right\}. \tag{6}$$

This comparison is deliberately secondary. It removes the portfolio distinction that defines wealthy hand-to-mouth households.

The real-estate screen is applied item by item. Forest, water bodies, miscellaneous land, animal sheds, barns, work in progress, and wells are excluded. Agricultural land enters when its total value reaches Rs. 1 lakh, and the dwelling bundle (structure plus residential land) enters at the same floor. The screen removes 2.42 percent of recorded real-estate value but affects 32.7 percent of households, illustrating why small-valued items can move the P/W split while leaving total HtM unchanged.

The Rs. 1 lakh screen, the institutional credit limit, and the pledge-augmented limit are parameters in implemented sensitivity exercises. They are not household-level observations of lender eligibility, approval, or available credit. We therefore interpret Table 16 as a set of economic counterfactuals rather than as a confidence interval for the benchmark share.

Table 16: Total HtM Prevalence under Credit-Limit Alternatives

Implemented credit-limit rule	Total HtM share
Zero credit	61.83
One month of total income	59.59
Two months of total income	55.76
India-specific institutional limit	59.00
Pledge-augmented limit	51.98

Notes: Entries are weighted percentages of AIDIS households. All rows use the benchmark liquid and illiquid balance-sheet map, monthly income predicted by the headline gradient-boosting model, $\tau = 2$, and the union of the zero-kink and borrower branches. The institutional rule combines the account-overdraft and Kisan Credit Card rules described above. The pledge-augmented rule adds the encoded fraction of gold and provident-fund wealth to potential credit.
Source: Authors' calculations from AIDIS and CPHS.

B.3 Prediction models

B.3.1 Income

The income ladder begins with two model-free benchmarks. Consumption bins assign households to one hundred consumption groups and impute the group's mean consumption-to-income mapping. The rank map transfers the CPHS income distribution according to a household's consumption rank. The parametric benchmarks are OLS, LASSO, and elastic net. The implemented nonlinear models are random forest and histogram gradient boosting. We do not estimate stand-alone decision-tree, neural-network, or XGBoost specifications.

The headline histogram-gradient-boosting model uses 500 boosting iterations, a learning rate of 0.06, early stopping, and the common numerical and categorical predictor set. A no-consumption version removes all consumption variables. A reweighted version estimates CPHS-to-AIDIS density weights and fits the same model under those weights. Quantile-gradient-boosting models estimate the 5th, 25th, 50th, 75th, and 95th conditional quantiles and draw an income by interpolation. The total-income target is in logs; frequent, narrow, and self-production income use the inverse hyperbolic sine.

B.3.2 Wealth

Every wealth target is modeled in two parts. The first part predicts $\Pr(T_i > 0 \mid X_i)$. The second predicts the amount among positive holders. The combined prediction is

$$\widehat{T}_i(\tau_T) = \mathbf{1}\{\widehat{p}_i \geq \tau_T\}\widehat{A}_i, \tag{7}$$

where τ_T is calibrated separately for each target in AIDIS. The calibration chooses the weighted cutoff at which the predicted positive share matches the observed positive share. Tree-model probabilities are out of fold for this calculation. The cutoff and amount model are then refitted or frozen as appropriate and deployed in CPHS; no CPHS wealth moment is targeted.

The model ladder consists of consumption/wealth bins, logit plus log-OLS with Duan smearing, two-part elastic net, random forest, histogram gradient boosting, a quantile-gradient-boosting draw, deterministic quantile mapping, and quantile mapping with a participation map. Quantile mapping uses only the AIDIS training sample. Among predicted positive holders, it retains the gradient-boosting rank and replaces the amount with the observed positive amount at the same weighted AIDIS quantile. The participation map is applied only at deployment: it maps the CPHS distribution of ownership probabilities onto the AIDIS distribution while preserving CPHS ranks.

B.4 Wealth validation at the classification margin

Table 17 substitutes each predicted balance sheet for the observed balance sheet of held-out AIDIS households. This experiment isolates prediction error from differences between the AIDIS and CPHS populations. Conditional-mean gradient boosting understates the HtM side of the liquidity condition by 21.3 percentage points. Quantile mapping slightly overshoots the observed share but provides the highest household- level agreement. The stochastic draw comes closest to the observed group shares but assigns more individual households to a different group. The participation map is the identity on AIDIS, so the participation-mapped route coincides with the quantile-mapped row in this table; it differs only at deployment.

Table 17: Classification Validation on Held-Out AIDIS Households

Balance sheet	HtM condition	P	W	N	HtM agree.	Class agree.
Observed	58.8	8.0	50.7	41.2	—	—
Gradient-boosting mean	37.5	4.9	32.6	62.5	68.5	67.6
Quantile-mapped GB	66.6	9.5	57.2	33.4	74.3	73.3
Quantile-GB draw	60.1	7.8	52.3	39.9	62.4	61.5

Notes: Entries are weighted percentages. The sample is the common AIDIS holdout. Income is the headline gradient-boosting prediction. “Liquid agreement” is agreement on equation (1); “three-way agreement” is agreement on P-HtM, W-HtM, and N-HtM.
Source: Authors’ calculations from AIDIS and CPHS.

Table 18 performs the corresponding distribution check for gold, the input to the collateral-capacity condition, and documents why the piecewise leg is the selected gold prediction. Every earlier leg assigns literally zero non-owners in both surveys, the footprint of the leaked ownership indicator described in Section 6.3; the conditional-mean leg additionally compresses the distribution, and the draw leg restores AIDIS zeros but explodes the upper tail and loses the zeros in deployment. The piecewise leg alone reproduces the observed AIDIS distribution on every statistic, and its deployed CPHS distribution preserves the 17.5 percent non-ownership margin while the owners' median shifts up with the CPHS sample's covariate ranks.

Table 18: Gold Holdings: Observed, Predicted in AIDIS, and Deployed in CPHS

Leg	Sample	% zero	Min	Max	Median	Mean	Median|>0
Observed	AIDIS	17.5	0	9,000,000	25,000	61,266	38,500
GB mean	AIDIS	0.0	2	1,072,970	60,084	89,650	60,084
GB mean	CPHS	0.0	1	945,950	99,317	119,628	99,317
Quantile map	AIDIS	0.0	1	9,000,000	35,000	69,135	35,000
Quantile map	CPHS	0.0	1	3,849,366	60,000	95,327	60,000
Draw	AIDIS	17.5	0	534,786,089	24,994	121,713	35,373
Draw	CPHS	1.2	0	444,813,366	50,432	122,306	51,248
Piecewise (selected)	AIDIS	17.5	0	8,752,668	24,355	56,949	35,000
Piecewise (selected)	CPHS	17.5	0	3,534,532	50,000	75,486	60,000

Notes: Weighted statistics of the gold value in rupees; minima and maxima are simple extremes. Predicted legs are classify-then-amount at each leg's calibrated threshold. The quantile-map rows coincide with the participation-mapped rows because the participation map was inert for gold under the leaked threshold. The piecewise leg refits ownership without the leaked indicator, rank-maps the deployed ownership scores onto the AIDIS distribution, and rank-maps positive amounts onto the observed positive AIDIS distribution.
Source: Authors' calculations from AIDIS and CPHS.

B.5 Direct CPHS classification

The deployed CPHS classification by prediction method is reported in Table 5 in the main text. The P-HtM row decomposes cleanly. The observed AIDIS holdout share is 8.0 percent; the quantile-mapped prediction on that same holdout gives 9.5 percent, so model error does not produce the collapse. Deployed on CPHS, every amount-only route gives roughly 0.1 percent because the deployed ownership probabilities lack the low-participation mass observed in AIDIS. Mapping the participation margin restores 3.7 percent, stable between 3.6 and 4.3 percent across the calibrated ownership thresholds. The remaining gap to the AIDIS level is consistent with CPHS's documented undercoverage of the poorest, asset-less households.

B.6 Within-survey status extrapolation

The status extrapolation estimates multinomial logit and gradient- boosting models separately within AIDIS and CPHS. Each survey receives a harmonized-only specification and a full-native specification. The CPHS training label is its own direct-route classification, not the transported AIDIS label. The selected full-native gradient-boosting model predicts the 2019 CPHS class and carries that fixed household label to the local- projection panel as a robustness classification; the headline panel labels are the participation-mapped direct-route labels themselves (Section 6.3). The extrapolation training labels remain the conditional-mean direct route. Alternative labels based on the quantile-draw direct route are estimated separately because the conditional-mean label has a much smaller W-HtM cell.

Table 19 reports the diagnostics for the model actually used to attach status to the monthly panel. Overall accuracy is high because N-HtM is the dominant label in the direct-route training data. Balanced accuracy and class-specific recall reveal the weaker margin: the model recovers only 45 percent of P-HtM labels.

Table 19: Within-CPHS Status-Extrapolation Diagnostics

Holdout diagnostic	Value
Overall accuracy	0.948
Balanced accuracy	0.737
Recall: N-HtM	0.990
Recall: P-HtM	0.450
Recall: W-HtM	0.770

Notes: Diagnostics use the common 20 percent CPHS household holdout and the selected full-native gradient-boosting specification. The training outcome is the 2019 CPHS direct-route label constructed from observed income and conditional-mean predicted balance sheets. The table evaluates reproduction of that generated label; it does not validate the unobserved household balance sheet.
Source: Authors' calculations from AIDIS and CPHS.

C Monetary-Policy Appendix

C.1 Estimation sample and outcome construction

The consumption estimation begins with accepted CPHS household-month responses and retains one observation per household and month. It uses nominal total expenditure, food expenditure, and appliance expenditure as released. Each outcome is transformed to 100 times its natural logarithm when positive. Zero or negative values are set to missing by the log transformation. This restriction is consequential for appliances: the outcome is lumpy, many household-months are zero, and the effective sample is much smaller than for total or food expenditure. We therefore treat the appliance response as a diagnostic rather than a headline result.

The income estimation applies the same transformation to positive total household income and wage income. CPHS negative missing-value codes are removed before the logarithm is taken. The employment measure starts from members age 15 or older, averages the paid-employment indicator within household and month, and expresses the result in percentage points. Because the people file is observed approximately once per four-month survey wave, we fill the monthly panel and carry the last observed employment rate forward for no more than three months. All local projections are unweighted and include household fixed effects.

C.2 Classification routes and estimators

For consumption, we estimate the household fixed-effects local projection for the transported classification, the conditional-mean direct classification, the quantile-mapped direct classification, the participation-mapped direct classification, and the full-native CPHS extrapolation. Where predicted gold is available, the quantile- and participation-mapped routes also split W-HtM households into C-W-HtM and NC-W-HtM groups. Every route is measured in 2019 and held fixed over the panel. For income and employment, we estimate the quantile-mapped and participation-mapped three-way and gold-capacity classifications.

The second estimator follows the synthetic-series equivalence described by Almuzara and Sancibrián (2024). For each group and horizon, we first average household cumulative changes into a group-by-month series. It then estimates a time-series local projection with an AIC-selected lag length from one to six, capped by the horizon, and heteroskedasticity-robust inference. Under a balanced panel, its point estimate would coincide with the household fixed-effects estimator; differences can arise because household composition changes with the horizon. The synthetic-series estimator reports group levels but not a covariance-correct C-W-HtM minus NC-W-HtM standard error.

C.3 Additional group and shock contrasts

Table 20 reports the three-way contrasts under the participation-mapped classification. The early P-HtM minus W-HtM difference is the only precise entry: -1.90 log points per unit surprise at month four. The P-HtM minus N-HtM difference has the same sign and a larger magnitude at that horizon but a standard error four times as large, because the N-HtM comparison does not share the tight covariance structure of the within-regression P–W contrast. By month seven, all three differences are close to zero, and no stable full ordering emerges over the horizon. The values are coefficients per unit target surprise; multiplying by 0.090 gives the response to a one-standard-deviation monthly target surprise.

Table 20: Three-Way Consumption Contrasts at Selected Horizons

Contrast	Month 4	Month 7	Month 12
P-HtM minus W-HtM	−1.90 (0.89)	−0.35 (1.25)	−0.58 (1.07)
W-HtM minus N-HtM	−2.90 (3.39)	0.56 (4.84)	−2.35 (6.48)
P-HtM minus N-HtM	−4.80 (3.68)	0.21 (5.70)	−2.93 (6.75)

Notes: Entries are differences in cumulative total- consumption responses to the one-day target factor, in log points per unit surprise. Standard errors are in parentheses and use the covariance between coefficients from the joint household fixed-effects regression. The classification uses observed CPHS income and participation-mapped predicted wealth, measured in 2019 and held fixed.

The path factor does not generate the gold-capacity pattern found for the target factor. The unscaled C-W-HtM minus NC-W-HtM total-consumption differences under the corrected capacity classification are 0.02 (standard error 0.81) on impact, 3.59 (2.12) at month three, 2.32 (2.45) at month six, 3.64 (1.75) at month nine, and 1.65 (1.74) at month twelve. Because the target and path factors enter jointly, these are conditional responses to path news holding the target surprise fixed.

For total consumption, we also implement three robustness specifications: it drops March 2020 through June 2021; replaces the target and path factors with their first principal component; and uses the two-day target and path factors. These variants are run for total consumption only. The income, wage, and employment estimations use only the joint one-day target and path specification. No reported estimate uses Driscoll–Kraay standard errors.

The analysis does not estimate a time-varying HtM classification, a status-versus-characteristics horse race, simultaneous confidence bands across horizons, or standard errors that propagate the estimated wealth and gold models into the local projection. These omissions delimit the inference: the reported confidence bands are pointwise, and all local-projection standard errors condition on the predicted household classification.

D Long-Horizon and Pooled-Group Local Projections

This appendix reports three extensions of the local-projection evidence in Section 9. All estimates use the headline specification of that section: household fixed effects, month-of-year indicators, three lags of the outcome difference and of both shock factors, standard errors clustered by calendar month, and group contrasts with covariance-correct standard errors from the joint regression. The capacity split throughout uses the corrected piecewise

Table 21: Key Contrasts from the Local-Projection Extensions

Exercise	Outcome	Contrast	Peak (month)	End of horizon
Pooled group, 12m	Wages	PNC−N	−7.47 (3.29) [6]	−11.04 (5.13)
Pooled group, 12m	Wages	PNC−CW	−6.06 (2.51) [11]	−6.39 (2.72)
Pooled group, 12m	Consumption	PNC−CW	3.69 (2.53) [2]	−3.28 (3.33)
Pooled group, 12m	Employment (pp)	CW−N	2.12 (0.81) [6]	1.77 (1.03)
Pooled group, 48m	Consumption	PNC−CW	−8.00 (4.32) [13]	−1.24 (2.55)
Pooled group, 48m	Total income	PNC−CW	−6.71 (2.99) [17]	−3.08 (3.19)
Pooled group, 48m	Employment (pp)	PNC−CW	3.49 (1.29) [29]	−0.38 (1.19)
Five-way, 48m	Total income	CW−NCW	7.90 (3.64) [17]	4.44 (4.35)
Five-way, 48m	Employment (pp)	CW−NCW	−4.21 (1.39) [29]	0.30 (1.11)
Income-sorted, 12m	Wages	PNC−N	−13.25 (11.29) [12]	−13.25 (11.29)

Notes: Cumulative responses to the one-day target surprise in log points (employment in percentage points) per unit surprise, with month-clustered, covariance-correct standard errors in parentheses; the peak column reports the horizon with the largest $|t|$-ratio, in brackets. Multiplying by 0.090 gives the response to a one-standard-deviation monthly surprise.
Source: Authors' calculations from AIDIS and CPHS.

gold prediction of Section 6.3. Unless noted, entries are cumulative responses to the one-day target factor in log points per unit surprise; multiplying by 0.090 gives the response to a one-standard-deviation monthly target surprise. Figures are scaled to the one-standard-deviation surprise. Table 21 collects the principal estimates.

D.1 The pooled constrained group

The reading offered in Section 9 is that the population that responds as if constrained is the poor hand-to-mouth plus the wealthy hand-to-mouth without usable collateral. This subsection tests that reading directly rather than assembling it from separate contrasts. Households are regrouped into three categories: the pooled constrained group (P-HtM plus NC-W-HtM, denoted PNC), the collateral-capable wealthy hand-to-mouth (C-W-HtM), and the non-hand-to-mouth, and the local projections are re-estimated on this three-way split under the participation-mapped classification.

Figure 4 reports the twelve-month contrasts. The wage contrasts are the sharpest estimates in the battery. The pooled group's wage income falls relative to the non-hand-to-mouth by −7.47 (standard error 3.29) log points per unit surprise at month six and −11.04 (5.13) at month twelve, and relative to the capable wealthy by −4.20 (1.93) at month six and −6.39 (2.72) at month twelve. The consumption contrast between the pooled group and the capable wealthy is individually imprecise and sign-mixed within the year, −4.51 (4.70) at month seven and −3.28 (3.33) at month twelve; its separation sharpens just past the year mark, in the long-horizon estimates below. Employment moves on the other side of the split: the capable wealthy gain employment relative to the non-hand-to-mouth, 2.12 (0.81) percentage points at month six.

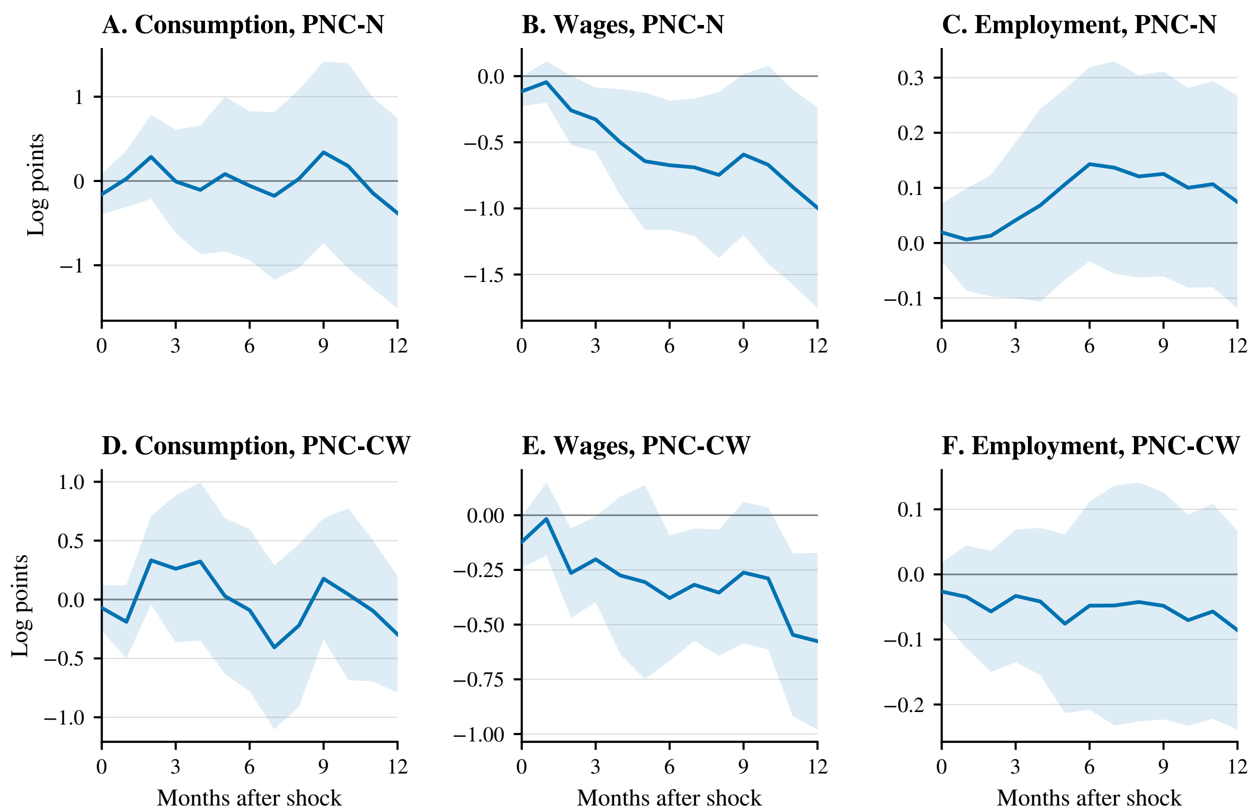

Figure 4: Pooled constrained-group contrasts at the twelve-month horizon

Notes: Cumulative responses to a one-standard-deviation one-day target surprise under the participation-mapped classification with P-HtM and NC-W-HtM pooled (PNC). Top row: PNC minus non-HtM. Bottom row: PNC minus C-W-HtM. Shaded areas are pointwise 90 percent bands with month-clustered, covariance-correct standard errors. *Source:* Authors' calculations from AIDIS and CPHS.

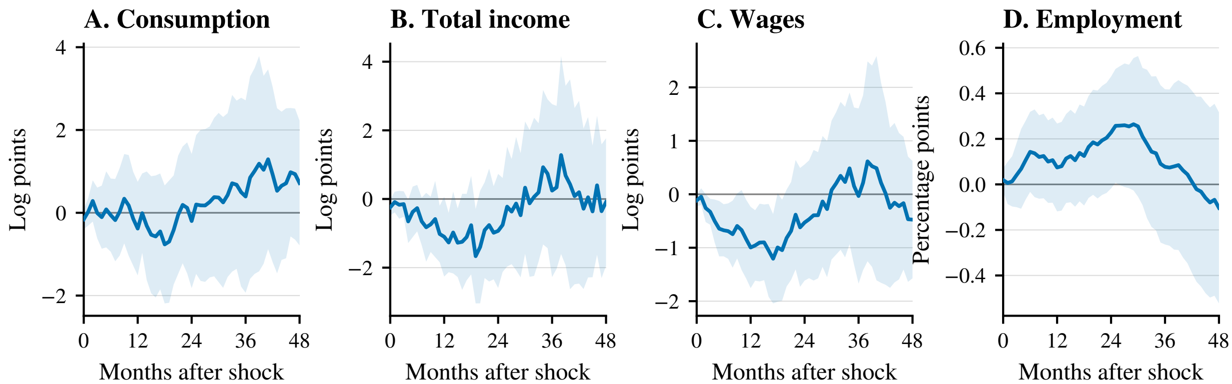

Figure 5: Pooled constrained group versus non-HtM at forty-eight-month horizons
Notes: Cumulative responses to a one-standard-deviation one-day target surprise; PNC minus non-HtM contrast under the participation-mapped classification with P-HtM and NC-W-HtM pooled. Shaded areas are pointwise 90 percent bands. *Source:* Authors' calculations from AIDIS and CPHS.

D.2 Forty-eight-month horizons

Extending the horizon from twelve to forty-eight months asks whether the differential responses are temporary displacements or persistent gaps. Figure 5 traces the pooled group against the non-hand-to-mouth. Two results stand out. First, the income losses of the constrained group are a mid-horizon phenomenon: the PNC minus C-W-HtM total-income contrast bottoms at −6.71 (standard error 2.99) at month seventeen and largely closes by the end of the horizon, −3.08 (3.19) at month forty-eight. Second, the employment margin reverses sign at long horizons: the PNC minus C-W-HtM employment contrast reaches 3.49 (1.29) percentage points around month twenty-nine. Households without usable collateral cannot borrow against wealth, and the rise in their relative employment two to three years out is consistent with self-insurance through labor supply. Consumption completes the sequence: the PNC minus C-W-HtM consumption contrast peaks just past the first year, −8.00 (4.32) at month thirteen, and then decays toward zero, −1.24 (2.55) at month forty-eight. The constrained group's consumption separates as its income gap widens, and both gaps close within four years, consistent with the labor-supply margin financing a recovery of spending. Food consumption sharpens this reading: flat across groups in the first year, it reverses sign at mid horizons, with the pooled group's food spending rising 6.57 (1.92) log points per unit surprise above the capable wealthy at month twenty-two, and the corresponding within-wealthy contrast showing the same reversal (8.40, standard error 2.51, for NC-W-HtM relative to C-W-HtM). The recovery that the labor-supply margin finances is concentrated in necessities.

The same pattern appears inside the wealthy hand-to-mouth group under the five-way classification of the main text. The C-W-HtM minus NC-W-HtM total-income contrast peaks at 7.90 (3.64) log points per unit surprise at month seventeen and fades to 4.44 (4.35) by month forty-eight, and the wage contrast reaches 8.94 (3.27) at month fifteen; the corresponding employment contrast is −4.21 (1.39) percentage points at month twenty-nine, the mirror image of the pooled-group result: it is the households without gold capacity whose relative employment rises. The consumption contrast between the same groups peaks just past the

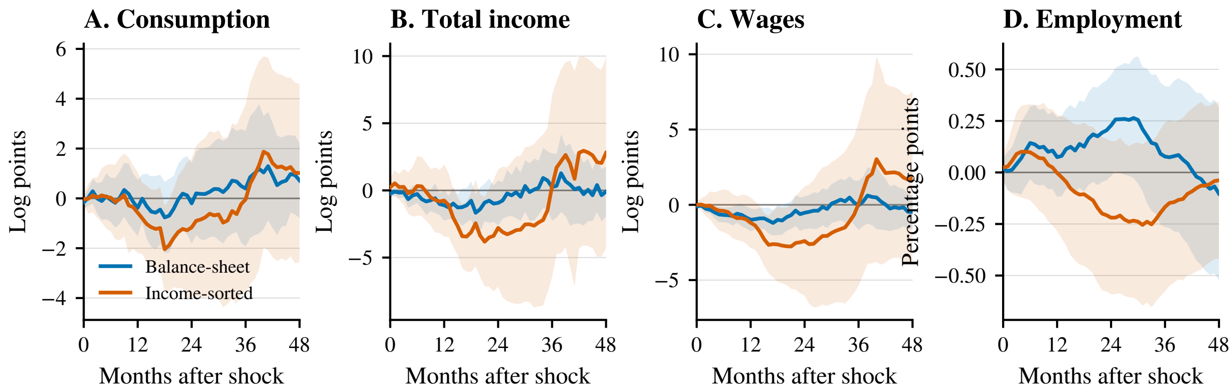

Figure 6: Balance-sheet classification versus income sorting

Notes: PNC minus non-HtM contrasts for consumption, total income, wages, and employment under the balance-sheet classification and under groups formed by sorting the observed 2019 CPHS income distribution at the same population shares (27.7/31.9/40.4), each with its own pointwise 90 percent band. Cumulative responses to a one-standard-deviation one-day target surprise. *Source:* Authors' calculations from AIDIS and CPHS.

first year, 9.48 (5.08) at month thirteen, and decays thereafter.

D.3 An income-sorting placebo

A natural concern is that the balance-sheet classification simply relabels the income distribution. To assess this, we hold the group shares fixed and discard the balance sheet: households are sorted by their observed mean 2019 CPHS income and cut at the population shares of the AIDIS gold-illiquid classification with the constrained groups pooled, 27.7 percent (the PNC analog), 31.9 percent (the C-W-HtM analog), and 40.4 percent (the non-HtM analog). The local projections are then re-estimated on the income-sorted groups.

Figure 6 overlays the two assignments. The income-sorted contrasts track the balance-sheet contrasts in shape and remain inside the 90 percent bands, which is expected: income and balance-sheet position are correlated, and the bottom of the income distribution contains many genuinely constrained households. The difference is precision. At month twelve the wage contrast against the non-hand-to-mouth is -13.25 (11.29) under income sorting against -11.04 (5.13) under the balance-sheet classification, and the employment contrast between the capable group and the non-hand-to-mouth peaks at 1.81 (1.27) against 2.12 (0.81). Consumption completes the pattern: the income-sorted PNC minus non-HtM contrast is -6.36 (10.91) at month twelve against -4.23 (7.62) under the balance-sheet grouping, similar in level with a markedly larger standard error. Sorting on income alone inflates the standard errors substantially at unchanged group sizes, so the balance-sheet classification isolates the responding population more sharply than income rank does. The classification is not a relabeled income distribution.