Household Consumption Insurance in India *

Arpita Chatterjee† Anand Chopra‡ Urvi Neelakantan§
Gautham Udupa¶

Version: September 2025

Abstract

How well are Indian households insured against income shocks, and how does this vary across the income distribution? We estimate a panel unobserved components model using monthly data from India's Consumer Pyramids Household Survey (CPHS), covering over 90,000 households from 2016-2019. We find that permanent income shocks in India are more volatile than in developed countries, with annualized standard deviation more than twice than that found in the US. Transitory shocks also exhibit substantial higher persistence. Despite this riskier income environment, households achieve considerable consumption insurance via self-insurance: the marginal propensity to consume (MPC) out of permanent income shocks averages 0.22, while the MPC out of transitory shocks is only 0.03. Strikingly, the MPC out of permanent shocks follows a concave pattern across the income distribution — low among the poorest households, rising through middle-income groups, then flattening among the wealthy. We rationalize these findings using a life-cycle model that incorporates consumption commitments proportional to permanent income and a government-provided consumption floor. The model suggests that current transfer policies may have limited efficacy for fiscal stimulus, as the lowest-income households — who would typically have the highest MPCs — are constrained by the design of means-tested programs.

*This research presents views of the authors and not that of the Federal Reserve Board or the Federal Reserve System. First version: September 2025.

†Board of Governors of the Federal Reserve System; chatterjee.econ@gmail.com
‡University of Liverpool
§Federal Reserve Bank of Richmond
¶Centre for Advanced Financial Research and Learning (CAFRAL)

1 Introduction

The degree to which consumption changes in response to income fluctuations—partial insurance in the terminology of Blundell, Pistaferri, and Preston (2008)—determines both the welfare costs of income volatility and the effectiveness of fiscal interventions. However, little is known about consumption insurance in developing economies, where income volatility is higher and formal insurance mechanisms more limited.

India provides a compelling laboratory for studying consumption insurance. With 1.4 billion people, substantial income inequality, and limited access to formal financial markets, India represents the lived experience of a large fraction of the world's population. Indian households face high income volatility from sources ranging from monsoon variability to informal employment, yet they also have access to extensive kinship networks and an expanding system of government transfers through programs like the Public Distribution System.

This paper uses high-frequency panel data from India's Consumer Pyramids Household Survey to provide the first comprehensive estimates of heterogeneous consumption insurance using a Kalman filter based maximum likelihood method. We estimate the semi-structural model of Blundell et al. (2008) using the quasi maximum likelihood methodology developed in Chatterjee, Morley, and Singh (2021). Our paper provides estimates of time-varying permanent and transitory income volatility, persistence of transitory income shocks, and most importantly marginal propensities to consume out of permanent and transitory income shocks for a major emerging economy. Our analysis of over 90,000 households from 2016-2019 yields three main findings that differ markedly from developed country patterns.

First, the income shock environment in India is riskier than in developed countries. We estimate that permanent income shocks are substantially higher and transitory shocks more persistent than in the US data.

Second, despite this riskier environment, Indian households achieve substantial consumption insurance. The average transmission coefficient from permanent income shocks to consumption is 0.34, indicating that households smooth about two-thirds of even persistent income changes. Combined with the average propensity to consume measures from the data, the estimated transmission coefficient implies an estimated marginal propensity to consume around 0.22, considerably lower than that in developed countries. Insurance against transitory shocks is nearly complete, with a transmission coefficient of only 0.03.

Third, the marginal propensity to consume out of permanent income shocks follows a distinctly concave pattern across the income distribution. The poorest households exhibit the lowest MPCs (around 0.15), which rise to peak values of around 0.3 among middle-income households before flattening among the wealthy. This contradicts the standard prediction that liquidity-constrained poor households should have the highest MPCs.

We develop a life-cycle model with incomplete asset markets and a risk-free asset to self-insure against income risk (see Heathcote, Storesletten, and Violante (2009) for a survey). We augment the model in two ways motivated by the Indian institutional environment: consumption commitments that scale with permanent income, and a government-provided consumption floor through means-tested transfer programs (Hubbard, Skinner, and Zeldes, 1995). Consumption commitments capture social and cultural obligations that rise with economic status, generating relative risk aversion decreasing in income and wealth (Guntin, Ottonello, and Perez, 2023). The government-provided consumption floor increases insurance against income shocks for the low-income households. Moreover, the consumption floor creates strong incentives for eligible households to limit asset accumulation given income gains, as it can jeopardize transfer eligibility.

Our calibrated model reproduces the empirical MPC pattern. The model reproduces the "excess consumption smoothing" to transitory and permanent income shocks observed in the data. Moreover, it can reproduce the concave pattern of MPC out of permanent income shocks. The consumption floor suppresses MPCs among the poorest households by making them reluctant to increase consumption when income rises. Consumption commitments proportional to permanent income make wealthy households less responsive to income changes than poor households. Together, these mechanisms generate MPCs that rise with income among lower-income households before flattening out.

These findings carry important policy implications especially for the determination of optimal income and consumption tax progressivity. Our model highlights that the tools used to insure against income risk matter for determining optimal tax progressivity; higher self-insurance through savings would raise optimal tax progressivity whereas insurance through government consumption floor would dampen redistributive benefits of progressive income tax. Moreover, given the importance of consumption taxes in developing countries, our findings highlight that progressive consumption taxes can generate substantial welfare gains when individuals have a strong desire to smooth consumption.

In addition, as Kaplan and Violante (2010) recommend, empirical estimates generated by models such as ours can serve as yardsticks for quantitative macroeconomic models, particularly the heterogeneous-agent, incomplete markets models that have become the workhorse of modern macroeconomics. We hope that the results generated in this paper can therefore provide important targets for the nascent quantitative macroeconomic literature on India.

Our work contributes to three strands of literature. We extend the consumption insurance literature pioneered by Blundell et al. (2008) to a major developing economy, demonstrating how institutional differences shape household responses to income shocks. We contribute to the fiscal policy literature by providing micro-foundations for heterogeneous MPC patterns that affect transfer multipliers (Kaplan and Violante, 2014; Orchard, Ramey, and Wieland, 2025). Finally, we inform debates about social protection design by showing how means-testing can reduce both insurance and macroeconomic effectiveness of transfer programs.

The remainder of the paper proceeds as follows. Section 2 describes the Consumer Pyramids Household Survey data. Section 3 presents our empirical methodology and main results on income processes and consumption insurance. Section 4 develops and calibrates our life-cycle model. In this preliminary version of the paper, we conclude in Section 5 exploring policy implications through proposed counterfactual experiments.

2 Data

We use longitudinal data from India's Consumer Pyramids Household Survey (CPHS), which is conducted by the Center for Monitoring Indian Economy (CMIE). The survey includes a representative sample of over 174,000 households, making it the world's largest household panel survey (University of Pennsylvania Library, 2025). It is a continuous survey that began in 2014 and is conducted in three four-month waves annually. Each household is surveyed once in each wave at an interval of four months; thus, every household is surveyed three times a year. For example, if a household is first surveyed in February of a particular year, it is surveyed again in June and October of that year.[1]

The CPHS sample size and frequency offers several advantages that make it suitable to address the questions that we pose. The frequency of data collection makes it possible to observe income shocks and consumption responses that would be missed at annual or lower frequencies. Higher-frequency data can be particularly valuable for distinguishing between transitory and permanent income shocks, a distinction that is particularly relevant in contexts where seasonal employment and informal work arrangements may also contribute to income volatility. The higher frequency also likely reduces recall bias, as respondents report more recent income and consumption rather than trying to recall numbers from further in the past.

[1] For additional details, see Vyas (2020).

2.1 Sample Selection

Our sample selection process begins with the universe of households in the CPHS data. Starting from the raw merged dataset of 231,371 households, we restrict attention to married households and limit analysis to household heads aged 25-65 years to concentrate on working-age households. We focus on 2016-2019 to ensure a consistent macroeconomic environment and trim the top and bottom 1% of income growth observations to address outliers. Most importantly, we require households to have non-missing income and consumption data for at least 12 of the 16 quarters. This process yields our final analysis sample of 90,817 households. Sample restrictions are described in Table A.2.

2.2 Descriptive Statistics

Table 1 presents comprehensive descriptive statistics for our estimation sample. The data reveal substantial heterogeneity across demographic, income, and consumption dimensions that make the sample well-suited for studying consumption insurance patterns.

2.2.1 Demographic Characteristics

Our analysis sample contains households with a median size of 4 members and typically includes 1 child under age 18. The median age of household heads is 43 years, reflecting our focus on working-age populations. Educational attainment varies considerably: 41% of household heads have 0-7 years of schooling, 37% completed 8-12 years, and 18% have more than 12 years of education.

Geographic distribution shows that 36% of households reside in urban areas, while 6% live in major metropolitan centers ("big cities"). Notably, 20% of households are engaged in agriculture, reflecting the continued importance of this sector in the Indian economy. The unemployment rate in our sample is low at 2%, consistent with the focus on households with regular income reporting.

2.2.2 Income Patterns

The income distribution exhibits substantial inequality and volatility characteristic of developing economies. Median monthly labor earnings are ₹ 12,500 (approximately $150 at 2019 exchange rates) with a standard deviation of ₹ 15,381, indicating considerable dispersion. The income distribution is highly right-skewed: the 10th percentile (₹ 4,644) represents subsistence-level earnings, the median captures lower-middle-class households, and the 90th percentile (₹ 33,437) reflects relatively affluent households by Indian standards.

Income volatility is substantial and underscores the importance of understanding how households respond to income fluctuations. The predominance of labor earnings is striking—for the median household, wage income comprises 99% of total income, with minimal reliance on business income (0%), transfer payments (0.17%), or other sources.

2.2.3 Consumption Behavior

Monthly nondurable consumption patterns reveal both the economic constraints and smoothing behavior of Indian households. The median household spends ₹ 7,592 monthly on nondurable goods, with a standard deviation of ₹ 4,541 that is proportionally much smaller than income dispersion. This suggests some degree of consumption smoothing even in the raw data.

The consumption-to-income ratio has a median of approximately 60%, indicating that typical households save about 40% of their income. However, the large standard deviation (467%) reflects substantial heterogeneity in saving behavior across the income distribution. The quarterly growth rate of nondurable consumption shows a median near zero (-0.16%) with a standard deviation of 27%—markedly lower volatility than income growth, providing preliminary evidence of consumption smoothing.

The composition of consumption reflects the relatively low income levels and spending priorities typical of developing economies. Food represents the largest expenditure category at 58% of nondurable consumption for the median household. Housing costs show a median share of zero, likely reflecting widespread home ownership or residence in family compounds without explicit rental payments. Other significant categories include transportation (2.7%), healthcare (2.3%), and education (0.8%), with the remainder (31%) comprising various other nondurable goods and services.

2.2.4 Sample Characteristics for Consumption Insurance Analysis

Several features of our sample make it well-suited for analyzing consumption insurance. First, the substantial income inequality—spanning from near-subsistence to relatively affluent households—enables examination of heterogeneous consumption responses across the income distribution. Second, the high income volatility provides ample variation in permanent and transitory shocks necessary for identification. Third, the quarterly frequency and multi-year panel structure allow separation of permanent from transitory income components.

The predominance of labor income simplifies the analysis by focusing attention on employment-related shocks rather than complex portfolio responses. The significant agricultural representation (20%) ensures coverage of both formal and informal employment relationships. Finally, the substantial variation in consumption-income ratios across households suggests meaningful differences in saving behavior and liquidity constraints that may drive heterogeneous consumption responses to income shocks.

3 Empirical Model & Results

We follow Blundell et al. (2008), contextualized for salient demographic characteristics of India, to obtain residual income and consumption from the first stage regression. In particular, we include deterministic controls such as year and state fixed effects, age, education categories, number of household members, number of kids, big city and state interacted with year, caste and religion, employment status. Residual income and consumption can be decomposed into a permanent and transitory components:

$$
\begin{aligned}
y_{i,t} &= \tau_{i,t} + \epsilon_{i,t} + \theta\epsilon_{i,t-1}, \qquad && \epsilon_{i,t} \sim i.i.d.(0, \sigma_\epsilon^2),\\
c_{i,t} &= \gamma_\eta \tau_{i,t} + \kappa_{i,t} + \upsilon_{i,t}, \qquad && \upsilon_{i,t} \sim i.i.d.(0, \sigma_\upsilon^2),
\end{aligned}
$$

Here $\tau_{i,t}$ is a common stochastic trend for income and consumption ("permanent income"), $\epsilon_{i,t}$ is a transitory income shock with moving-average parameter $|\theta| < 1$, $\kappa_{i,t}$ is an additional stochastic trend for consumption, and $\upsilon_{i,t}$ is a transitory consumption shock. The trends are specified as random walks:

$$
\begin{aligned}
\tau_{i,t} &= \tau_{i,t-1} + \eta_{i,t}, \qquad && \eta_{i,t} \sim i.i.d.(0, \sigma_\eta^2),\\
\kappa_{i,t} &= \kappa_{i,t-1} + \gamma_\epsilon \epsilon_{i,t} + u_{i,t}, \qquad && u_{i,t} \sim i.i.d.(0, \sigma_u^2),
\end{aligned}
$$

We allow the permanent and transitory income shocks to have time-varying volatility: σ_η^2 and σ_ϵ^2 vary by quarter. We estimate all the parameters of the semi-structural model, $\theta, \gamma_\eta, \gamma_\epsilon, \sigma_u, \sigma_\upsilon, \sigma_{\eta,t}, \sigma_{\epsilon,t}$, following quasi maximum likelihood based approach of Chatterjee et al. (2021).[2]

Sun (2024) has found important differences in household consumption insurance even within India using GMM estimation method of Blundell et al. (2008). Motivated by these findings, we

[2] See the online appendix of Chatterjee et al. (2021) for state space representation of the Blundell et al. (2008) model.

Table 1: Descriptive Statistics - CMIE Consumer Pyramids Household Survey

	Estimation Sample		Full Sample	
	Median	Standard Deviation	Median	Standard Deviation
Panel A: Demographic Variables				
Household size	4.00	1.75	5.00	2.06
Number of children 18 and under	1.00	1.11	1.00	1.16
Age of household head	43.00	10.73	49.0	31.93
Educational Attainment (%)				
0-7 Years	41.33	49.24	35.68	47.91
8-12 Years	36.85	48.24	25.31	43.48
> 12 Years	17.86	38.30	6.46	24.59
Urban household (%)	35.81	47.94	32.96	47.01
Big city (%)	5.56	22.92	5.72	23.23
Agriculture (%)	19.51	39.63	2.51	15.64
Unemployed (%)	1.64	12.70	0.41	6.43
Panel B: Income Measures				
Monthly labor earnings	12,500.00	15,381.03	6,000.00	10,666.60
Log monthly labor earnings	9.43	0.66	9.26	0.74
Growth rate of labor earnings (quarterly, %)	0.00	2,026.78	-3.33	4,523.29
Income percentiles:				
10th percentile	4,643.64		0.00	
50th percentile (median)	12,250.20		8,666.67	
90th percentile	33,436.63		27,150.00	
Income sources (% of total):				
Labor earnings	98.98	26.08	97.70	37.90
Business income	0.00	19.37	0.00	25.69
Transfer income	0.17	9.70	0.28	19.27
Other income	99.65	21.40	99.39	30.70
Panel C: Consumption Measures				
Monthly nondurable consumption	7,591.67	4,541.01	5,923.33	4,961.75
Log monthly nondurable consumption	8.96	0.44	8.89	0.55
Consumption-to-income ratio	59.65	467.13	60.36	1,160.06
Growth rate, nondurable consumption	-0.16	27.46	-2.11	40.08
Consumption shares (% of total):				
Food	58.30	9.94	59.36	10.14
Housing	0.00	7.70	0.00	7.39
Transportation	2.68	2.87	2.41	2.77
Education	0.79	9.09	0.21	8.32
Healthcare	2.31	10.68	2.29	15.85
Other	30.91	18.91	31.04	21.88

Notes: This table presents descriptive statistics for our analysis sample from the CMIE Consumer Pyramids Household Survey covering the period January 2016-December 2019. The sample includes households with heads aged 25-65 years with at least 12 quarters of non-missing income and consumption data. All monetary values are nominal Indian Rupees.

report results both for the overall population and also provide comparisons within two subgroups: rural vs urban and agriculture vs non-agriculture. We also report estimates for the US from the literature (Blundell et al. (2008) and Chatterjee et al. (2021)) to offer a baseline for comparison.

How volatile and persistent are income shocks in India? We begin first by looking at the standard deviation (SD) of the permanent and transitory shocks over time. Figure 1 displays the quarterly estimates of $\sigma_{\eta,t}$ and $\sigma_{\epsilon,t}$. In general, our estimates reveal that the permanent income shock in India is more volatile—and the transitory shock more persistent—than in the US context. Consistent with literature, transitory shock is more volatile than permanent shocks to income.

Figure 1: Standard Deviation of Income Shocks

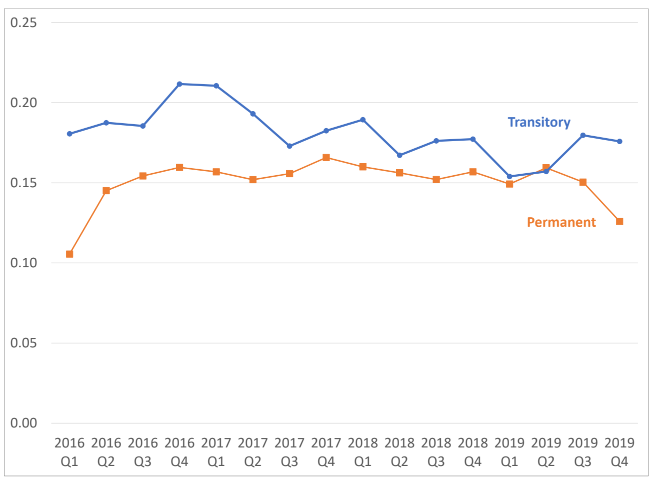

Table 2 reports our estimates of persistence of transitory income shocks, θ, from the pooled data. In addition to the overall numbers being high, there are also notable differences within subgroups. For example, the transitory shock is far more persistent in urban areas than in rural areas, and among households engaged in non-agricultural occupations compared to agricultural ones. For completeness, we also report estimates of σ_u, volatility of the independent trend in consumption in Table 2 and σ_v, volatility of transitory consumption. No notable heterogeneity is observed in these estimates across various subgroups.

What is the estimated transmission coefficient of permanent income shocks to consumption, γ_η that governs household consumption insurance against permanent shocks? Table 3

Table 2: Estimates of Persistence of Transitory Shock

Parameter	Whole Sample	Rural	Urban	Agri	Non-Agri
Persistence of transitory income (θ)	0.2808	0.1927	0.3885	0.2075	0.2980
	(0.0044)	(0.0064)	(0.0051)	(0.0099)	(0.0049)
SD. Consumption Trend (σ_u)	0.092	0.092	0.093	0.096	0.092
	(0.00025)	(0.00046)	(0.000296)	(0.0008)	(0.000262)
SD. Transitory Consumption (σ_v)	0.147	0.139	0.15	0.147	0.147
	(0.00027)	(0.00053)	(0.000316)	(0.00088)	(0.000286)

reports the estimates for the full sample and subgroups. The estimated transmission coefficient out of permanent income shocks is 0.34 for the full sample; consumption changes far less than one-for-one with the income shock implying a very high degree of self insurance. While there is little difference between rural and urban households, we see that the transmission coefficient is significantly higher for agricultural households. Not surprisingly, the transmission coefficient out of transitory shocks is far lower across the board.

Table 3: Estimates of Transmission Coefficients

Parameter	Whole Sample	Rural	Urban	Agri	Non-Agri
Transmission of Permanent Shock (γ_η)	0.3406	0.3399	0.3357	0.3901	0.3345
	(0.0034)	(0.0067)	(0.0037)	(0.0126)	(0.0035)
Transmission of Transitory Shock (γ_c)	0.0279	0.0214	0.0341	0.0174	0.0301
	(0.0015)	(0.0017)	(0.0026)	(0.0028)	(0.0017)

Our primary interest is to assess the pattern of transmission coefficient, average propensity of consume (APC) and ultimately marginal propensity to consume (MPC) across the income distribution to gauge the relationship between income and the degree of household consumption insurance. In Figure 2 we report the estimated transmission coefficients, the APC measures and the implied MPC out of the permanent shock by income decile. Observe that households in the lowest two income deciles have far lower transmission coefficients than the rest, with the transmission coefficient being largely flat for households in the top 80% of the income distribution. APC declines monotonically across the income distribution. As a result, the implied MPCs depict a curious concave pattern with regard to income.

To tease out the mechanisms behind the concave pattern of MPC out of permanent income shocks, we use a life-cycle model calibrated to the Indian data. Before proceeding to present the life-cycle model, here we refer to the key related papers in the literature that estimate the same consumption and income process parameters using the widely used PSID data to compare our

Figure 2: Heterogeneity in Transmission Coefficient, APC and MPC

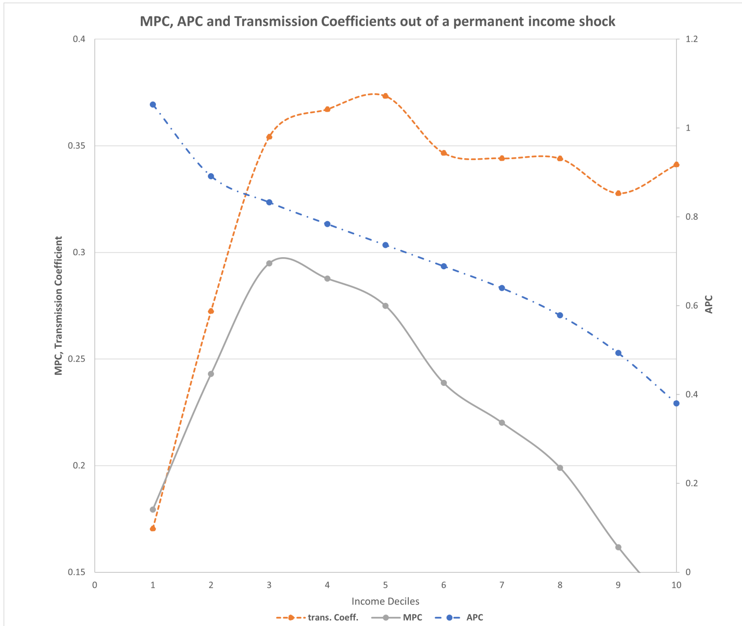

estimates. The semi-structural model we estimate is same as that in Blundell et al. (2008) who use GMM to estimate the parameters, while Chatterjee et al. (2021) propose a maximum likelihood based method of estimation. Our semi-structural model is the same as in Blundell et al. (2008) while our method of estimation follows Chatterjee et al. (2021). In terms of estimates, the persistence of transitory shock is considerably higher, and the transmission of permanent income shocks to consumption is considerably lower in Indian data compared to the PSID using either the GMM (as in Blundell et al. (2008)) or QMLE (as in Chatterjee et al. (2021)). We present this comparison in Table 4, where BPP refers to Blundell et al. (2008) and CMS refers to Chatterjee et al. (2021). For the transmission coefficient estimate, we also report the estimate from Sun (2024) that uses the same CPHS data from India but follows the GMM method of Blundell et al. (2008) for estimation of the semistructural model. Note that the the estimate of Sun (2024) is not statistically different from our QMLE estimate. As noted in the Monte Carlo exercises of Chatterjee et al. (2021), difference between GMM and QMLE methods of estimation for consumption insurance is particularly prominent in small sample sizes, and the large sample size of CPHS probably leads to near convergence between the two methods of estimation.

Table 4: Comparison of Key Estimates

Data	CPHS (India)		PSID (US)	
Parameter	Our estimate	Sun (2024) estimate	BPP	CMS
Transmission of Permanent Shock (γ_η)	0.3406	0.3749	0.6423	0.45
	(0.0034)	(0.0078)	(0.0945)	(0.04)
Persistence of transitory income (θ)	0.2808	NA	0.1132	0.16
	(0.0044)	NA	(0.0247)	(0.02)

4 Life-Cycle Model to Rationalize Empirical Findings

Our empirical analysis shows that (i) households maintain substantial insurance against permanent and transitory income shocks, and (ii) the MPC out of permanent income shocks is a concave function of permanent income. To rationalize these empirical findings, we extend a standard life-cycle model with idiosyncratic income risk and incomplete asset markets (Deaton, 1991; Carroll, 1992; Kaplan and Violante, 2010) by incorporating two key features: (i) household consumption commitments and (ii) a government-provided minimum consumption floor.

4.1 Model Environment

The economy consists of a continuum of households of measure one, indexed by i. Time is discrete; a period corresponds to one year in the model. Households work from age 1 to R periods of life. They retire at age $R+1$ and live until age T. To focus on idiosyncratic income risk, we assume there is no aggregate uncertainty. We also abstract from additional sources of uncertainty during retirement such as survival risk (Yaari, 1965) or health risk (De Nardi, French, and Jones, 2010).

Households receive pre-tax labor income Y_{it}, which during working years comprises three components: (i) a deterministic component (Γ_t) that is common to all households and varies by age, (ii) a permanent income component $\left(\tau_{i,t}\right)$ that follows a random walk process, and (iii) a MA(1) transitory component. We impose a second-degree polynomial in logs for the deterministic component and income shocks follow the same process as in the empirical setup in Section 3. Household income is described as:

$$Y_{i,t} = exp(\Gamma_t) \times exp\left(\tau_{i,t} + \epsilon_{i,t} + \theta\epsilon_{i,t-1}\right) \tag{4.1}$$

$$\Gamma_t = \alpha_0 + \alpha_1 \times t + \alpha_2 \times t^2 \tag{4.2}$$

Households save in a risk-free asset $\left(A_{i,t+1}\right)$ that pays a constant after-tax return of $1+r$. We assume a no-borrowing constraint such that $A_{i,t+1} \geq 0$.

The household's budget constraint is:

$$C_{i,t} + A_{i,t+1} = (1+r)A_{i,t} + Y^{p}_{i,t} + TR_{i,t} \tag{4.3}$$

$$Y^{p}_{i,t} = \chi\left(Y_{i,t}\right)^{1-\psi} \tag{4.4}$$

$$TR_{i,t} = \max\left\{0, \bar{c} - \left((1+r)A_{i,t} + Y^{p}_{i,t}\right)\right\} \tag{4.5}$$

where $Y^{p}_{i,t}$ is post-tax labour income and $TR_{i,t}$ denotes government transfers. Following Benabou (2002), ψ determines the marginal tax rate and χ is an income-scaling parameter. This functional form implies that post-tax labor income is a concave function of pre-tax labor income.

Transfers depend on household earnings and assets. The government guarantees a minimum level of consumption, $\bar{c}$, if cash-on-hand falls below it. Programs such as the Indian Public Distribution System (PDS), where the government provides food-grains at highly subsidized prices, are earnings- and asset-tested (Shrinivas, Baylis, and Crost, 2025). For example, eligibility for the PDS requires owning no motor vehicles or agricultural equipment. Hence, we impose that $C_{i,t} = \bar{c}$ and $A_{i,t+1} = 0$ when households receive positive transfers. Hubbard et al. (1995) and De Nardi et al. (2010) use such a formulation of government transfers to explain low-asset holdings of poor households in the US.

Households maximize time-separable expected utility given by

$$E_0 \sum_{t=1}^{T} \beta^t \frac{\left(C_{i,t} - H_{i,t}\right)^{1-\sigma}}{1-\sigma} \tag{4.6}$$

where $\beta \in (0,1)$ is the discount factor and σ is the coefficient of relative risk aversion if $H_{i,t} = 0$. $H_{i,t}$ represents consumption commitment, which we assume is proportional to the level of permanent income:

$$H_{i,t} = \lambda exp\left(\Gamma_t + \tau_{i,t}\right) \tag{4.7}$$

where λ determines the intensity of commitment. Higher values of λ imply increase the average propensity to consume out of income. Moreover, these preferences generate relative risk aversion that decreases in income and wealth. Relative risk aversion (RRA) is given by:

$$\text{RRA} = -\frac{u''\left(C_{i,t}\right) \times C_{i,t}}{u'\left(C_{i,t}\right)} = \sigma \frac{C_{i,t}}{C_{i,t} - H_{i,t}} \tag{4.8}$$

Since richer households are more likely to be far away from their habit, their relative risk aversion will be lower. This formulation of consumption commitment generates the same decreasing relative risk aversion as habit formation preferences (Constantinides, 1990).[3]

A constant level of consumption commitment (independent of permanent income) would be equivalent to Stone-Geary preferences. Guntin et al. (2023) show that Stone-Geary preferences imply that consumption of low-income households reacts much less to aggregate income changes than high-income households. However, a constant level of consumption commitment changes the average and marginal propensity to consume for low-income households that are close to the threshold, but has limited effect on the response of high-income households. Thus, the our general preference structure subsumes Stone-Geary preferences.

A household's optimization problem during working periods can be written in recursive formulation as:

$$\begin{aligned} V\left(A, \varepsilon, \varepsilon_{-1}, \tau, t\right) &= \max_{A', C} \frac{(C-H)^{1-\sigma}}{1-\sigma} + \beta \mathbf{E}_{\eta', \varepsilon'} V\left(A', \varepsilon', \varepsilon, \tau', t+1\right) \\ &\quad \text{subject to:} \\ C + A' &= (1+r) A + Y^p + TR \\ A' &\geq 0 \end{aligned}$$

Retired households receive pension benefits $\left(B_{i,t}\right)$ that are proportional to their permanent income level in the last working period:

$$B_{i,t} = b \times \exp\left(\Gamma_R + \tau_{i,R}\right) \quad \forall t \in \{R+1, \ldots, T\}$$

[3]Utility with habit formation can rationalize an increasing equity share over the life-cycle in portfolio choice theory (Gomes and Smirnova, 2023; Meeuwis, 2020; Polkovnichenko, 2007).

where b is the replacement rate. The optimization problem during retirement periods becomes:

$$\begin{aligned} V(A,\tau,t) &= \max_{A',C} \frac{(C-H)^{1-\sigma}}{1-\sigma} + \beta \mathbb{1}_{t+1\leq T} V\left(A',\tau,t+1\right) \\ &\quad \text{subject to:} \\ C + A' &= (1+r)A + B + TR \\ H &= \lambda B \\ A' &\geq 0 \end{aligned}$$

4.2 Calibration and Life cycle patterns

We parametrize the model to reproduce certain key features of the Indian economy. First, we estimate some parameters directly from the data or take values from the literature. Second, we calibrate the risk aversion parameter (σ) and the government-provided consumption floor ($\bar{c}$) by targeting moments using simulated data from the model.

Demographics: Households are born at age 25, retire at age $R = 65$, and live until age $T = 85$. Households are born with zero assets and the permanent and transitory components in logs are set to zero.

Preferences: We calibrate $\sigma = 15.1$ to match the empirically observed MPC out of permanent income shocks of 0.23. We set the habit intensity parameter $\lambda = 0.10$, which implies that a household will maintain consumption of at least 10% of their permanent income level.

Discount Factor and Interest rate: We set the risk-free interest rate (r) to 1.5% and the discount factor (β) to 0.92.

Income Process: The coefficients for the deterministic income profile are: $\alpha_0 = 0$, $\alpha_1 = 0.038$ and $\alpha = -0.0011$. Similar to Kaplan and Violante (2010), these coefficients imply a concave income profile that peaks after working 25 years at roughly twice the entry-level income, and then gradually declines to about 77% of the peak value by retirement. We use the estimated values of the variances of the permanent and transitory income shocks and the moving-average parameter in Table 2. The variance of the permanent income level at age 24 is set at the same value as the variance of the permanent income shock. Pension benefits are assumed to be 25% of the permanent income level in the last working period. We assume a low replacement rate to reflect the limited coverage of India's pension system, which became available to all workers in 2009.

Tax system and government transfers: The Indian government's spending on in-kind transfers has been around 0.7% of GDP in recent years (Chakrabarty, 2024). We set the value of $\bar{c}$ to

0.633 such that the average value of transfers over the working life cycle relative to average labor income during working periods is 0.7%. We replicate the Indian tax system by setting the tax progressivity parameter ψ to 0.1550 and the income-shift parameter χ to 1.081, following Chakrabarti, Mishra, and Mohaghegh (2024).

Life cycle patterns: We solve for optimal policy functions using the endogenous grid method of Carroll (2006). We simulate a panel of 100,000 households from ages 25 to 85.[4] Before discussing our results for the MPC out of income shocks, we briefly document the model-implied life-cycle profiles of means and variances of income, consumption and, assets in Figure 3.

Figure 3a shows a concave pre-tax labor income profile. Average consumption grows until retirement due to the precautionary saving motive, declines thereafter due to absence of survival risk during retirement and because the intertemporal saving motive is negative $(\beta R < 1)$. Average asset holdings (right-axis) peak at retirement as households accumulate wealth for retirement consumption and precautionary purposes. Asset holdings then fall, as commonly predicted by life-cycle models. The model-implied wealth-income ratio over the working periods is 6.2, which is a bit higher than the wealth-income ratio of around 5 reported in Kumar (2019). The average savings rate is 26.7% in the model, consistent with gross domestic savings as a percent of GDP of around 30% in India Mohan and Kapur (2015).[5]

Figure 3b shows that the cross-sectional variance of log pre-tax income increases over the life cycle because of the accumulation of permanent income shocks. This is consistent with patterns found for the US (Guvenen, Kaplan, Song, and Weidner, 2022) and other developing countries, such as China (Ding and He, 2018) and Mexico (Puggioni, Calderón, Cebreros Zurita, Fernández Bujanda, Inguanzo González, and Jaume, 2022). Consequently, consumption inequality rises over the life cycle, but more slowly than labor income inequality due to tax progressivity, government transfers and self-insurance.

4.3 MPC coefficients in the model

We now discuss the MPC out of income shocks. The transmission coefficient out of a permanent income shock in the model is:

$$\gamma_t^{\eta} = \frac{\text{Cov}\left(\Delta \log C_{it}, \eta_{it}\right)}{\text{Var}\left(\eta_{it}\right)} \tag{4.9}$$

[4] The model is solved using 101 exponentially-spaced grid points for assets. We discretize the grid of the permanent component using a Markov chain with 21 equally spaced points. The grid for the discretized transitory component contains 11 grid points.

[5] Savings rate displays a hump-shaped profile over the life cycle, in line with the finding of Bairoliya, Chanda, and Fang (2021).

Figure 3: Life-cycle profiles for means and variances of baseline model

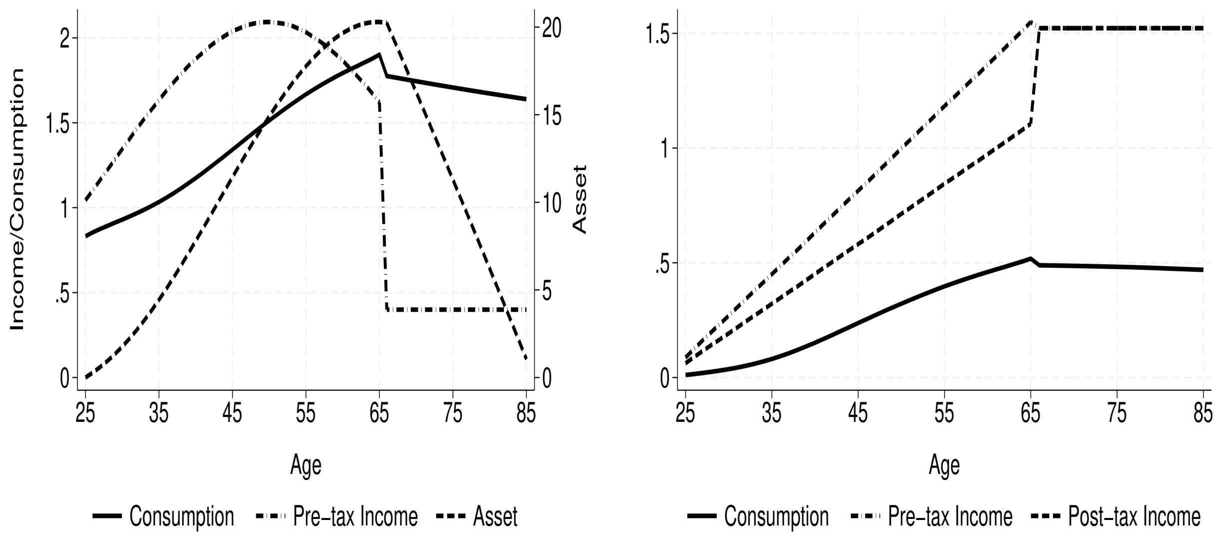

(a) Life-cycle Means

(b) Life-cycle inequality

The MPC out of a permanent income shock equals the transmission coefficient multiplied by the average propensity to consume (APC):

$$\mathrm{MPC}_t^{\eta} = \frac{\mathrm{Cov}\left(\Delta \log C_{it}, \eta_{it}\right)}{\mathrm{Var}\left(\eta_{it}\right)} \times \frac{\sum_i C_{it}/N}{\sum_i Y_{it}/N} = \gamma_t^{\eta} \times \mathrm{APC}_t \tag{4.10}$$

The model-implied MPC out a permanent income shock matches the empirically observed value of 0.23. We compute the MPC out of a permanent income shock by quintiles of lagged permanent income level, where these moments are averaged over the life cycle.

Figure 4 presents the transmission coefficient and MPC out of a permanent income shock (left-axis) and APC (right-axis) across quintiles of lagged permanent income level. It shows that the model can replicate the empirically observed concavity of the MPC function. This concavity results from an increasing transmission coefficient and a declining APC with permanent income. High permanent income households consume a smaller fraction of their income due to diminishing marginal utility. However, conditioning on a given asset level, consumers with higher lagged permanent income have a higher ratio of current assets to expected future income. Hence, upon receiving a income gain, their need for precautionary savings rises less than others, which leads to higher consumption and a higher transmission coefficient. Commault (2024) shows this feature is present in standard life-cycle models.

Figure 4: APC and insurance coefficients out of a permanent income shock

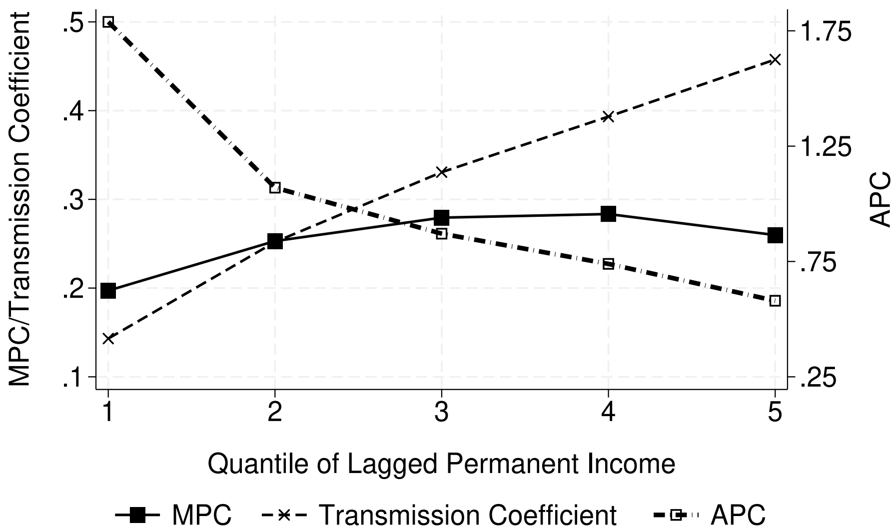

We can define MPC out of transitory income shocks in a similar way as the MPC out of permanent income shocks. We find substantial insurance against temporary shocks. The average MPC out of a transitory shock is 0.02, which is consistent with our empirical findings. Appendix Figure B.1 shows that MPC out of a transitory shock varies little with permanent income. Thus, we primarily focus on explaining the drivers of the MPC out of permanent income shocks in the rest of the paper.

The combination of consumption commitment and government transfers increases the slope of the transmission coefficient function with respect to permanent income, while their effect on increasing APC is limited to low permanent income households. To see this clearly, we analyze counterfactual economies: (i) without consumption commitment, (ii) without government-provided consumption floor, and (iii) a *standard* life-cycle model without either feature. We re-calibrate the risk aversion parameter across models to ensure the MPC out of a permanent income shock remains 0.23. Moreover, we re-calibrate $\bar{c}$ in the model without consumption commitment to maintain the ratio of government transfers to pre-tax income over the life cycle at 0.7%. Appendix Table B.1 shows the parameter values for the three counterfactual models. Appendix Figure B.2 presents the mean consumption, mean asset holdings, mean savings rate and consumption inequality for all working periods across the four models.

Figure 5a shows the MPC out of a permanent income shock in the baseline model and three counterfactual economies. First, let us consider a *standard* life-cycle model without either feature (blue line). The MPC out of a permanent income shock decreases with permanent income. Although the transmission coefficient increases from around 0.235 to 0.41 going from the lowest to the highest permanent income consumers, APC falls more sharply from around 1.46 to 0.54. Second, adding consumption commitment $\lambda = 0.10$ to the utility function implies that relative risk aversion declines with permanent income level. This reduces the pass-through of a permanent income shock into consumption relative to the *standard* model for low permanent income households, but increases it for high permanent income households (grey line). However, because the APC function does not change much relative to the *standard* model, the MPC out of a permanent income shock continues to decrease with permanent income in a model with consumption commitment.

Third, government transfers increase consumption and APC for all households relative to a standard model (red line). Moreover, they dampen the transmission coefficient of a permanent income shock for low permanent income households. The reduction in the transmission coefficient for low permanent income households is stronger than the increase in APC, which implies that the MPC out of a permanent income shock falls for low permanent income households relative to a standard model. Because government transfers are earnings- and asset-tested, they have no effect on high permanent income households. However, the relative risk aversion parameter needed to match the MPC out of a permanent income shock falls from 13.55 in a standard model to 7.95 in a model with government transfers. This increases the transmission coefficient of a permanent income shock for high permanent income households. Taken together, the MPC function becomes a concave function of lagged permanent income.

Finally, combining consumption commitment with a model with government transfers changes relative risk aversion across the permanent income distribution. As relative risk aversion falls with permanent income, it steepens the transmission coefficient out of a permanent income shock relative to a model with government transfers. There is limited effect on APC relative to a model with government transfers. Combined, this implies that the MPC out of a permanent income shock is lower for low permanent income households than for high permanent income households.

To limit the role of outliers, we use aggregate APC $\left(\frac{\sum_i C_{it}/N}{\sum_i Y_{it}/N}\right)$ to compute MPC in equation (4.10) instead of averaged household APCs $\left(\frac{1}{N}\sum_i \frac{C_{it}}{Y_{it}}\right)$. Appendix Figure B.3 shows higher variation in MPC out of a permanent income shock by permanent income level, but our conclusions about the importance of the mechanisms in explaining the empirical findings remain robust. Lastly, Appendix Figure B.4 shows that all models predict that MPC out of a transitory income shock vary little with permanent income. There is around 1 percentage point difference in the

Figure 5: Comparing MPC across models

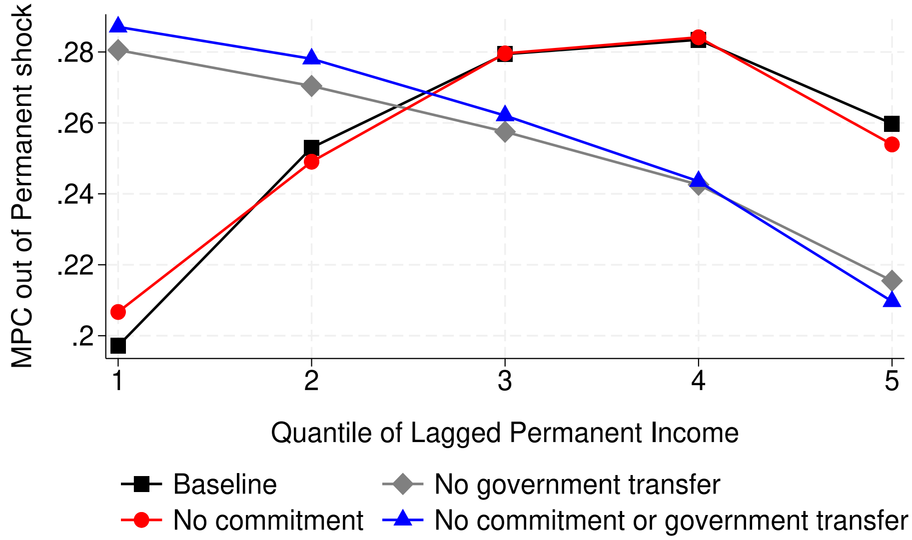

(a) MPC out of a permanent income shock

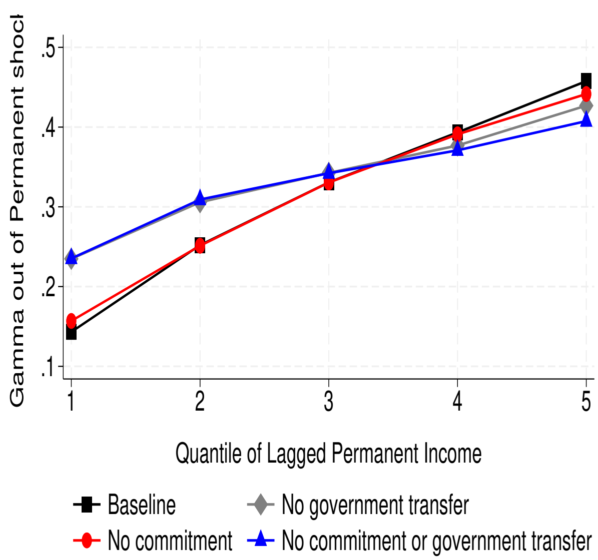

(b) Transmission Coefficient

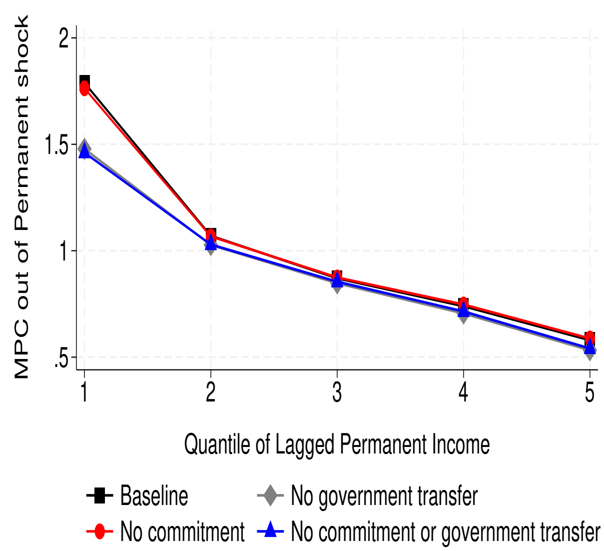

(c) APC

MPC out of a transitory income shock between the highest and lowest lagged permanent income household.

Figure 6: MPC to transitory and permanent shocks over life cycle across models

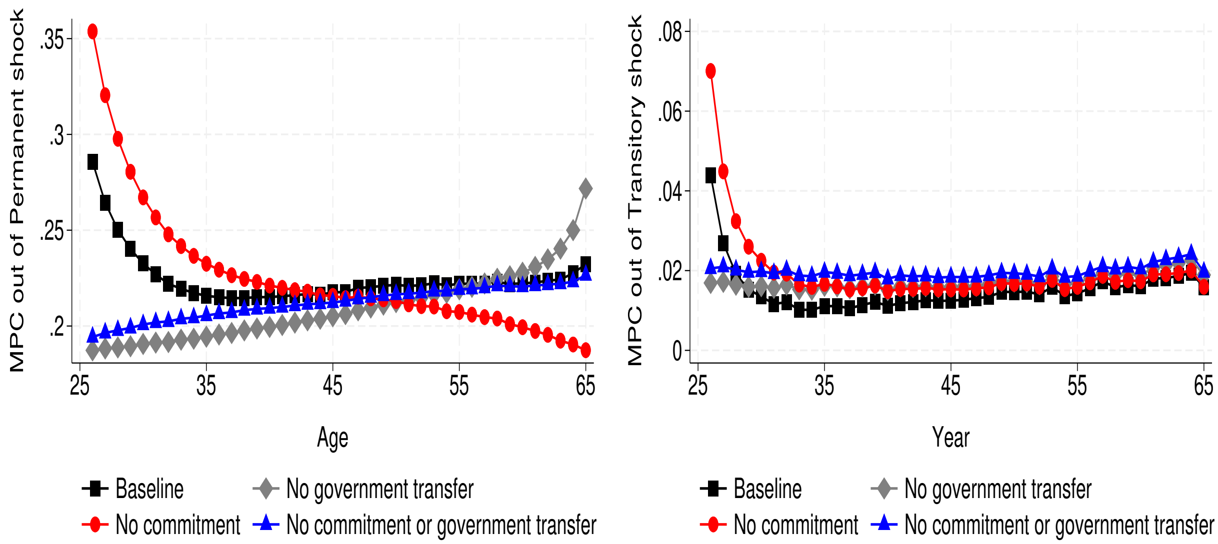

(a) Permanent Shock

(b) Transitory Shock

4.4 Discussion on Other Mechanisms

While the above results suggest an important role for consumption commitment and government transfers in accounting for the empirical findings, we discuss the role of other mechanisms and modeling assumptions in rationalizing the findings.

Borrowing constraints: The baseline model assumed a zero borrowing limit. Relaxing borrowing constraints reduces the pass-through of income shocks into consumption for low permanent income households. On the other hand, relaxed constraints will increase consumption relative to income for borrowers, which will increase their MPC. Hence, the overall effect for low permanent income households is ambiguous. Furthermore, borrowing constraints matter during initial working periods when asset holdings are low, as shown by Kaplan and Violante (2010). Figure 6 shows declining MPC out of permanent and transitory shocks over the life cycle for the baseline model and model without commitment (but including government transfers). In a model without government transfers (with or without commitment), MPCs out of permanent or transitory income shocks are weakly rising over the life cycle. Moreover, MPCs are higher during initial working periods in models with government transfers relative to those without. These are testable predictions of the model, and we plan to calibrate borrowing constraints in future work by comparing MPCs out of income shocks over the life cycle in the data and the model.

Illiquid Assets: Kaplan and Violante (2014) show that when households face costs to adjust their illiquid assets, MPC out of income shocks can be high for households with low liquid but sizable illiquid asset holdings (wealthy hand-to-mouth) compared to households with higher liquid asset holdings. This might help explain high MPC out of permanent shocks for households in the middle of the income distribution, but MPCs would decline at the top of the income distribution as it becomes easier to pay adjustment costs. Moreover, such a model would also predict high average MPC out of transitory shocks, which is inconsistent with our empirical findings.

Durables and Non-durables: Models with non-homothetic preferences over durables and non-durable goods, as in Browning and Crossley (2009) and Andreolli and Surico (2021), predict that MPC rises with income. The intuition is that households in response to a negative income shock delay durable good spending. Conversely, they increase durable good spending in response to a positive income shock. Though this model can help reconcile a concave MPC function for permanent shocks, it will also predict high average MPC out of transitory income shocks, which is inconsistent with our empirical results.

Bequests or medical risk during retirement: A large class of life cycle models argue for the importance of medical or mortality risk during retirement or bequest motives to help explain high asset holdings and saving rates at older ages (Carroll, 2000; De Nardi, 2004; De Nardi et al., 2010; Lockwood, 2018). These models predict that richer households save more, leading to a reduction in MPC out of permanent income shocks as income rises, which would be inconsistent with our empirical results. Though our model under-predicts wealth-income ratios, but given our focus on explaining the pattern of MPC over the permanent income distribution rather than drivers of wealth accumulation, we abstract away from all these dimensions.

5 Discussion

In the previous sections, we have shown that the MPC out of a permanent income shock is low on average, but displays substantial heterogeneity across the income distribution. In particular, the MPC out of a permanent income shock is smaller for the lowest income households than for middle-income households. In this section, we discuss the policy implications of our findings and directions for future research.

The concavity of the MPC function out of permanent income shocks has implications for optimal tax progressivity. Heathcote, Storesletten, and Violante (2017) shows that income tax progressivity and uninsurable risk are inversely related — optimal tax progressivity will be higher in economies with higher uninsurable risk. However, the tools used to insure against income risk matter for determining optimal tax progressivity. On the one hand, self-insuring using saving implies that redistribution through progressive taxes will increase welfare. In contrast, insurance achieved through government-provided consumption implies that the redistributive benefits of tax progressivity are substantially dampened. Thus, estimating the welfare cost of uninsurable risk and optimal tax progressivity given multiple channels for self-insurance can be a fruitful area of research.

Our results also have implications for the determination of consumption taxes. Consumption taxes are widely used in developing countries. For example, consumption taxes contribute to around 60% of total tax revenue (Mukherjee, 2025). da Costa and Santos (2023) shows that in a static model taxing income or consumption progressively is equivalent. However, in a dynamic model they have different implications for labour supply and consumption smoothing. Progressive consumption taxes can generate substantial welfare gains when individuals have a strong desire to smooth consumption, which our empirical estimates indicate is present in the Indian context. Thus, deriving optimal consumption tax progressivity for a developing country can be a policy-relevant research question.

References

ANDREOLLI, M. AND P. SURICO (2021): "Less is more: Consumer spending and the size of economic stimulus payments," *CEPR Discussion Paper No. DP 15918*.

BAIROLIYA, N., A. CHANDA, AND J. FANG (2021): "Consumption Smoothing and Household Savings in India: Role of Demographics and Durables," *Working Paper*.

BENABOU, R. (2002): "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?" *Econometrica*, 70, 481–517.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98, 1887–1921.

BROWNING, M. AND T. F. CROSSLEY (2009): "Shocks, stocks, and socks: Smoothing consumption over a temporary income loss," *Journal of the European Economic Association*, 7, 1169–1192.

CARROLL, C. D. (1992): "The buffer-stock theory of saving: Some macroeconomic evidence," *Brookings papers on economic activity*, 1992, 61–156.

——— (2000): "Why Do the Rich Save So Much?" in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. by J. B. Slemrod, Cambridge, MA: Harvard University Press.

——— (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics letters*, 91, 312–320.

CHAKRABARTI, A. S., A. MISHRA, AND M. MOHAGHEGH (2024): "Inequality and income mobility: the case of targeted and universal interventions in India," *The Journal of Economic Inequality*, 22, 1–27.

CHAKRABARTY, T. (2024): "Demand for Grants 2024-25 Analysis: Food and Public Distribution," *PRS Legislative Research, New Delhi.*

CHATTERJEE, A., J. MORLEY, AND A. SINGH (2021): "Estimating Household Consumption Insurance," *Journal of Applied Econometrics*, 36, 628–635.

COMMAULT, J. (2024): "Heterogeneity in MPC Beyond Liquidity Constraints: The Role of Permanent Earnings," *Working Paper.*

CONSTANTINIDES, G. M. (1990): "Habit formation: A resolution of the equity premium puzzle," *Journal of political Economy*, 98, 519–543.

DA COSTA, C. E. AND M. R. SANTOS (2023): "Progressive consumption taxes," *Journal of Public Economics*, 220, 104854.

DE NARDI, M. (2004): "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 71, 743–768.

DE NARDI, M., E. FRENCH, AND J. B. JONES (2010): "Why do the elderly save? The role of medical expenses," *Journal of political economy*, 118, 39–75.

DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.

DING, H. AND H. HE (2018): "A tale of transition: An empirical analysis of economic inequality in urban China, 1986–2009," *Review of Economic Dynamics*, 29, 106–137.

GOMES, F. AND O. SMIRNOVA (2023): "Stock Market Participation and Portfolio Shares Over the Life Cycle," *Available at SSRN 3808350.*

GUNTIN, R., P. OTTONELLO, AND D. J. PEREZ (2023): "The Micro Anatomy of Macro Consumption Adjustments," *American Economic Review*, 113, 2201–31.

GUVENEN, F., G. KAPLAN, J. SONG, AND J. WEIDNER (2022): "Lifetime earnings in the United States over six decades," *American Economic Journal: Applied Economics*, 14, 446–479.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2009): "Quantitative macroeconomics with heterogeneous households," *Annu. Rev. Econ.*, 1, 319–354.

——— (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132, 1693–1754.

HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1995): "Precautionary saving and social insurance," *Journal of political Economy*, 103, 360–399.

KAPLAN, G. AND G. L. VIOLANTE (2010): "How much consumption insurance beyond self-insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.

——— (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

KUMAR, R. (2019): "The Evolution of Wealth-Income Ratios in India 1860–2012," *WID.world Working Paper No. 2019/07*.

LOCKWOOD, L. M. (2018): "Incidental bequests and the choice to self-insure late-life risks," *American Economic Review*, 108, 2513–2550.

MEEUWIS, M. (2020): "Wealth Fluctuations and Risk Preferences: Evidence from US Investor Portfolios," *Available at SSRN 3653324*.

MOHAN, R. AND M. KAPUR (2015): "Pressing the Indian Growth Accelerator: Policy Imperatives," *IMF Working Paper WP/15/53*.

MUKHERJEE, S. (2025): "Do Taxes on Commodities and Services Bear a Disproportionate Burden in India? An Assessment from 2011–2012 to 2022–2023," *Journal of Development Policy and Practice*, 10, 120–144.

ORCHARD, J. D., V. A. RAMEY, AND J. F. WIELAND (2025): "Micro MPCs and macro counterfactuals: the case of the 2008 rebates," *The Quarterly Journal of Economics*, qjaf015.

POLKOVNICHENKO, V. (2007): "Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk," *Review of Financial Studies*, 20, 83–124.

PUGGIONI, D., M. CALDERÓN, A. CEBREROS ZURITA, L. FERNÁNDEZ BUJANDA, J. A. INGUANZO GONZÁLEZ, AND D. JAUME (2022): "Inequality, income dynamics, and worker transitions: The case of Mexico," *Quantitative Economics*, 13, 1669–1705.

SHRINIVAS, A., K. BAYLIS, AND B. CROST (2025): "Food Transfers and Child Nutrition: Evidence from India's Public Distribution System," *American Economic Journal: Applied Economics*, 17, 161–207.

SUN, J. (2024): "The Heterogeneous Response of Consumption to Income Shocks: Semi-structural and Quasi-experimental Estimates From India," .

UNIVERSITY OF PENNSYLVANIA LIBRARY (2025): "Consumer Pyramids DX Microdata," Accessed: 2025-01-27.

VYAS, M. (2020): *Survey Execution*, Centre for Monitoring Indian Economy Pvt. Ltd.

YAARI, M. E. (1965): "Uncertain lifetime, life insurance, and the theory of the consumer," *The Review of Economic Studies*, 32, 137–150.

A Data Appendix

Table A.1: Variables Used and Cleaning Steps

Variable	Cleaning Steps or Description
Wage Earnings	Monthly nominal earnings, avg. within HH-quarter
Non-durable consumption	Monthly nominal nondurable consumption, avg. within HH-quarter
Age	Age of first person as listed, min by HH-wave-year
Family members	Count of family members by HH-wave-year
State	State of origin, but state when defining big city
Whether kids	Dummy = 1 if zero kids, at HH-wave-year
Number of kids	Count of HH members with age<18, at HH-wave-year
Religion	As reported
Caste category	As reported
Education	Dummy = 1 if education of first member < 8th standard
	Dummy = 2 if education of first member b/w 8-12th standard
	Dummy = 3 if education of first member > 12th standard
Big city	8 districts of Delhi with 1,787 unique HHs
	Gautam Buddha Nagar with 821 unique HHs
	Mumbai and its suburban district with 1,777 unique HHs
	Thane district with 5,077 unique HHs
	District of Chennai with 801 unique HHs
	Bangalore urban and rural districts with 1,738 unique HHs
	Kolkata district with 778 unique HHs
	Hyderabad district with 1,198 unique HHs
	Gurgaon district with 592 unique HHs
Other income	Number of earning members other than HoH and spouse
Employment status	Dummy for working
Agri household	Dummy for HoH being employed in agri
Rural	Rural dummy at HH-wave-year
Earnings decile	Decile of the average HH-wave level income b/w 2016-19
Married	Max of whether a spouse exists within HH-wave-year

Notes: This table provides the steps undertaken in constructing each variable used in our analysis. In particular, we explain how the quarterly data on household earnings and consumption, and other variables are constructed.

Table A.2: Sampling Restrictions

Step	**Unique Households**
Raw CMIE people of India data	234,892
Raw CMIE income/consumption data	233,583
Merged POI, income/cons data	231,371
Sample Restrictions	
Drop HH-wave-year with married dummy = 0	195,714
Drop if age < 25 or age > 65	184,863
Keep years between 2016 and 2019	155,268
Trim top and bottom 1% of income growth	154,986
Changed agri HH status b/w 2016 to 2019	154,986
Drop if missing values in > 11 of 16 quarters	90,817
Final Sample	90,817

B Model Appendix

Figure B.1: APC and insurance coefficients out of a permanent income shock

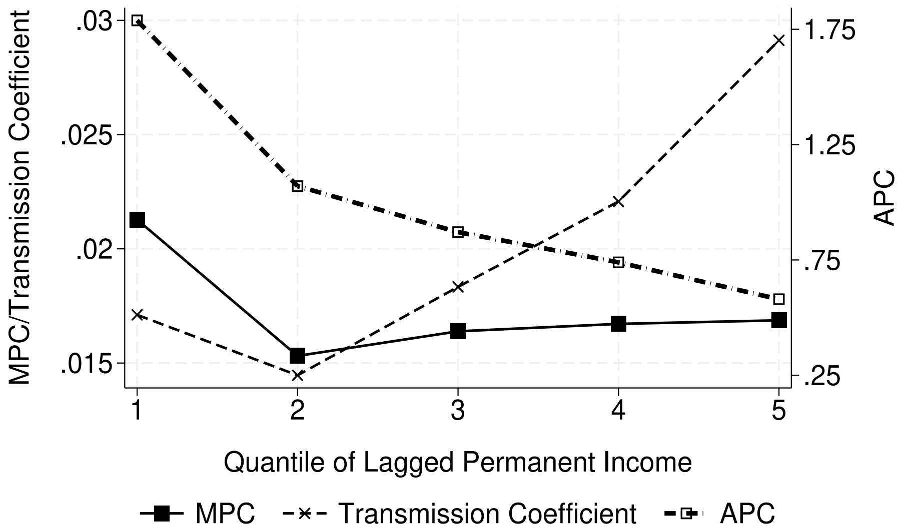

Table B.1: Calibrated Parameters Across Models

Model	σ	$\bar{c}$
Baseline	15.1	0.633
Model without consumption commitment	7.95	0.615
Model without $\bar{c}$	35.5	–
Model without consumption commitment or $\bar{c}$	13.55	–

Figure B.2: Life-cycle profiles for means and variances across models model

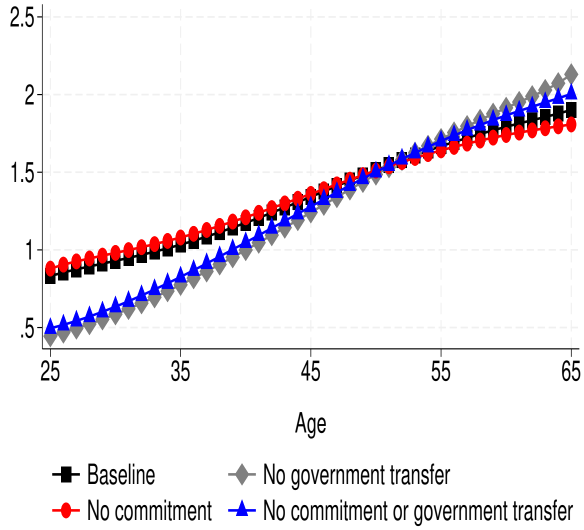

(a) Mean Consumption

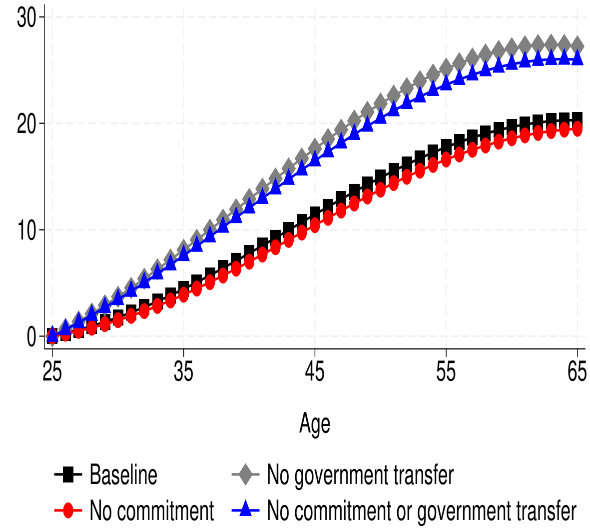

(b) Mean Asset

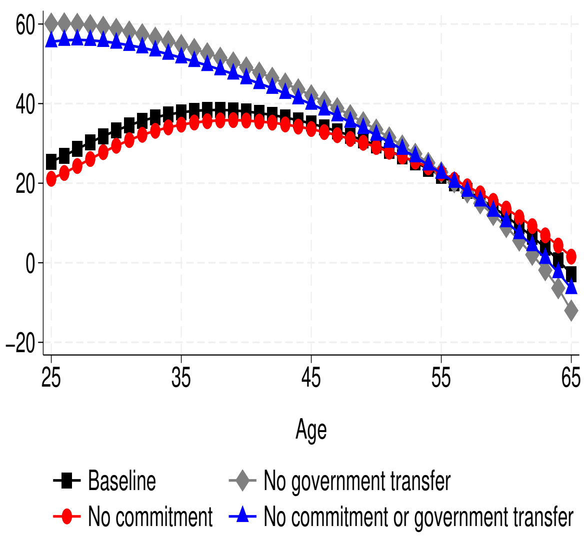

(c) Mean Savings Rate

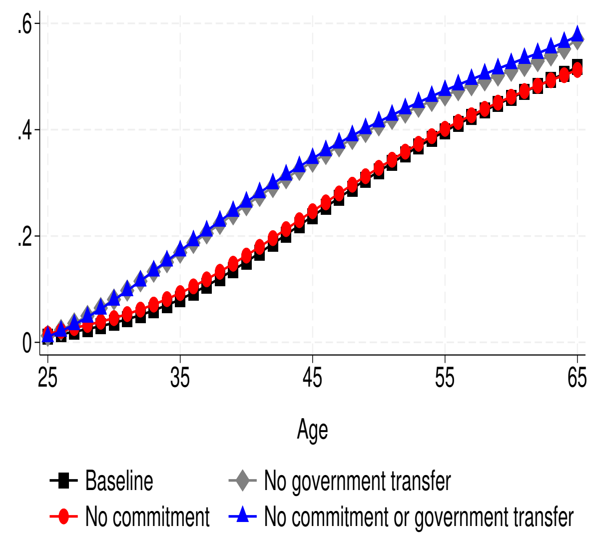

(d) Consumption Inequality

Figure B.3: MPC out of a permanent income shock using individual APC

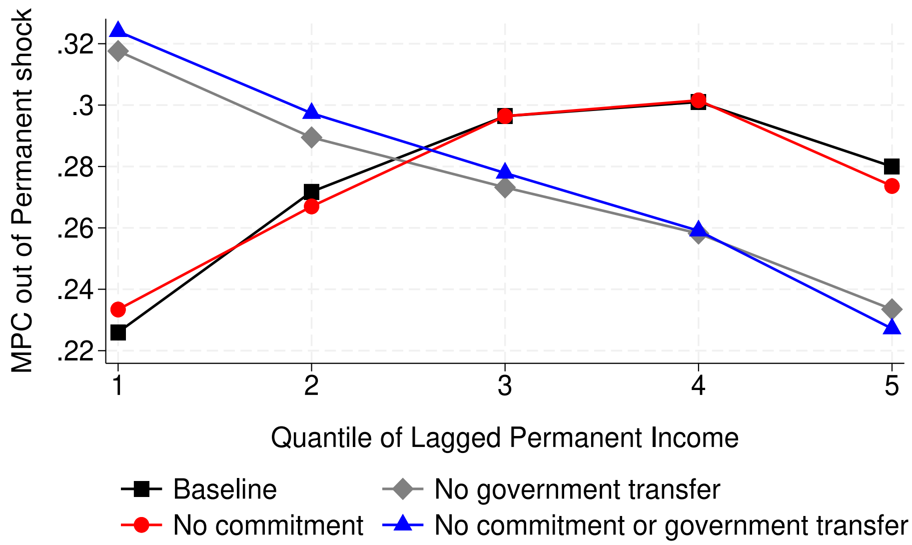

Figure B.4: MPC out of a transitory income shock across models