Cyclical Inflation Inequality and Monetary Policy

Sophia Chen*

Ting Lan†

Abstract

We document evidence of heterogeneity in the cyclicality of inflation across households. Our estimates reveal a positive covariance between households' marginal propensity to consume (MPC) and their elasticity of consumer prices in response to monetary policy shocks. This cyclical inflation inequality offsets about one-third of the amplification effect from cyclical income inequality, which arises from the positive covariance between households' MPCs and their elasticity of nominal income in response to shocks. The combined effect results in a 7 percent increase in the aggregate MPC compared to a benchmark with equal exposure to inflation and nominal income. These results quantify real income cyclicality——a key moment in heterogeneous-agent models.

JEL Codes: E21, E31, E32, E52

Keywords: inflation inequality, marginal propensity to consume, monetary policy transmission

*International Monetary Fund, email: ychen2@imf.org

†International Monetary Fund, email: tlan@imf.org.

The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

Recent advancements in the Heterogeneous-Agent New Keynesian (HANK) literature underscore the critical role of household heterogeneity in the transmission of monetary policy (Auclert, 2019; Bilbiie, 2020). In these models, the impact of monetary policy hinges on how effectively monetary policy shocks translate into changes in consumption. A key mechanism involves the covariance between the marginal propensity to consume (MPC) and the cyclicality of real income. Bilbiie (2020) demonstrates that the effects of monetary policy shocks are amplified when the real income of high-MPC households reacts strongly to changes in aggregate income (i.e., when the elasticity exceeds 1).

The cyclicality of real income depends on both nominal income cyclicality and inflation cyclicality. This relationship can be illustrated with a simple equation. Let i denote household, Y denote aggregate output, E_i denote nominal income, and P_i denote household-specific CPI. The covariance between households' MPCs and the cyclicality of their real income is given by:

$$cov(MPC_i, \frac{d(E_i/P_i)}{dY}\frac{Y}{E_i/P_i}) = cov(MPC_i, \frac{dE_i}{dY}\frac{Y}{E_i}) - cov(MPC_i, \frac{dP_i}{dY}\frac{Y}{P_i}), \tag{1}$$

where $\frac{d(E_i/P_i)}{dY}\frac{Y}{E_i/P_i}$ is the elasticity of the household's real income to aggregate output. The covariance can be decomposed into two terms as shown in equation (1). The first term captures heterogeneity in the cyclicality of nominal income, while the second captures heterogeneity in the cyclicality of inflation. We refer to these as cyclical income inequality and cyclical inflation inequality, respectively. While substantial exploration of cyclical income inequality exists in the literature, cyclical inflation inequality has received relatively less attention. One exception is Lan et al. (2024), who demonstrate in a Two-Agent New Keynesian (TANK) model that the effects of monetary policy shocks are dampened when the inflation experienced by low-MPC households is more cyclical than that of high-MPC households, thereby establishing a new transmission mechanism in heterogeneous-agent models.

The main contribution of this paper is to provide a systematic empirical estimate of the covariance between households' MPCs and inflation cyclicality, thereby quantifying a key moment in heterogeneous-agent models.

Estimating this covariance empirically presents two challenges. First, households' MPCs are not directly observable. Second, commonly used data often lack the granularity required to simultaneously estimate MPCs and inflation for the same household. We address these challenges by integrating information from multiple datasets. The estimation process has three steps.

In the first step, we calculate the Consumer Price Index (CPI) for household groups with specific characteristics, which include demographics, income, and wealth. We do this by linking household consumption data from the Consumer Expenditure Survey (CES) to item-level price data from the Bureau of Labor Statistics (BLS). In the second step, we estimate MPCs based on detailed household characteristics from the Panel Study of Income Dynamics (PSID), following the approach outlined by

Patterson (2023). In the final step, we assess inflation cyclicality based on group-specific CPI responses to monetary policy shocks and calculate its covariance with MPCs.

We present three main empirical findings. First, we document significant variations in inflation cyclicality across households. Specifically, households with lower educational attainment, lower income, and lower wealth experience more cyclical inflation. For example, households with less than a high school diploma face deflation nearly 5 percentage points greater than those with a college or postgraduate degree, 24 months after a monetary policy tightening of 100 basis points. Additionally, differences in inflation cyclicality are observed among households with varying income and wealth levels.

Second, consistent with prior literature, we find heterogeneous MPCs across households. Estimated as the consumption drop per dollar lost upon unemployment, the average MPC is 0.5. This figure is significantly higher among individuals with lower educational attainment, lower income, lower wealth, and among Black households.

Third, by linking households' MPC and inflation cyclicality, we uncover a substantial covariance between MPCs and inflation responses to monetary policy shocks. Following monetary policy tightening, household groups with higher MPCs tend to experience more deflation. This phenomenon arises because higher-MPC households, on average, allocate a greater share of their expenditures to items with more volatile prices.

The magnitude of cyclical inflation inequality, as estimated, is substantial enough to have a meaningful impact on the aggregate MPC and the efficacy of monetary policy. Using a simple theoretical framework, we demonstrate that the cyclical inflation inequality channel reduces the MPC out of labor income by approximately 6.5 percent. Abstracting from heterogeneity in the MPC out of non-labor income, this translates to a 3.5 percent decrease in the aggregate MPC.

Moreover, we estimate that the cyclical inflation inequality channel dampens the amplification effect from the cyclical income inequality channel—previously established by Patterson (2023)—by one third. Together, these two channels imply an overall amplification effect from real income cyclicality of 7.5 percent.

Our results are significant for several reasons. First, we provide empirical evidence of cyclical inflation inequality across households, which has implications for the distribution of aggregate inflation. Additionally, we quantify the covariance between households' MPC and inflation cyclicality, a key component of real income inequality that governing the transmission of policy shocks in heterogeneous-agent models. This covariance is also an essential moment in a new class of models that directly incorporates inflation heterogeneity. Thus, our empirical estimates represent a vital step in quantitatively evaluating these models. The sizable magnitude of the covariance we find indicates that cyclical inflation inequality is a significant mechanism in the transmission of policy shocks and could directly inform the calibration of optimal monetary policy within these frameworks. Finally, while we focus on monetary policy shocks, the multiplier in our analysis captures the general equilibrium effects

of a broad range of demand shocks. Consequently, our findings have implications not only for monetary policy but also for fiscal policy and the economy's response to business cycle shocks more broadly.

This paper contributes to three strands of the literature. The first strand focuses on monetary policy transmission in heterogeneous-agent models. A substantial body of work has explored how heterogeneous income cyclicality amplifies monetary policy shocks, employing both analytical (Acharya and Dogra, 2020; Bilbiie, 2020) and quantitative models (Bilbiie, 2020; Kaplan et al., 2018; Kaplan and Violante, 2014; McKay et al., 2016). Our paper distinguishes itself by examining a new channel: cyclical inflation inequality. The significance of this channel is highlighted in a TANK model presented in Lan et al., 2024. Our empirical estimates provide essential input for calibrating this class of models.

The second strand addresses differential consumption responses to income changes among various consumers. Empirical research has documented significant heterogeneity in MPCs across individuals and households (Cloyne et al., 2019; Jappelli and Pistaferri, 2014; Johnson et al., 2006; Patterson, 2023). Our paper builds on and extends this work by offering detailed estimates of MPC heterogeneity across an extensive range of household characteristics, emphasizing the covariance between MPCs and inflation cyclicality.

The third strand concerns the differential responses of prices faced by different consumers. Heterogeneity in prices has been documented across households, workers, and geographic areas (Argente and Lee, 2021; Jaravel, 2019; Kaplan and Schulhofer-Wohl, 2017). Heterogeneous responses in consumer prices to macroeconomic shocks is observed following recessions, currency devaluations (Cravino and Levchenko, 2017), and monetary policy shocks (Cravino et al. (2020). Our paper contributes to this literature by documenting new facts about heterogeneity in price cyclicality across the MPC distribution.

The rest of the paper is organized as follows. Section 2 presents a simple conceptual framework to demonstrate the role of cyclical inflation inequality in monetary policy transmission. Section 3 describes the data and empirical methodology. Section 4 presents estimation results. Section 5 discusses the implication for the aggregate multiplier. Section 6 concludes.

2 Conceptual framework

This section provides a conceptual framework linking cyclical inequality to the Keynesian multiplier. The cross-sectional distribution of income and inflation—and, consequently, real income—affects the economy's response to monetary policy shocks across a wide range of models. The impact of these factors depends on the specific setup of each model. Here, we use a simple model to clarify the underlying intuition.

The consumption function of household i is given by $C_i(W_i(E_i, P_i), r)$, where W_i is real income and r is the real interest rate. W_i depends on nominal income E_i and household-specific CPI P_i. Both E_i and P_i are functions of aggregate output Y. We assume that there is no capital or profits, and that all output is consumed. The market-clearing condition equates aggregate output to aggregate consumption:

$Y = C$. The aggregate MPC is $MPC = \frac{dC}{dY}$, which sums across households as:

$$MPC = \sum_i \frac{\partial C_i}{\partial Y} = \sum_i \frac{\partial C_i}{\partial W_i}\frac{\partial W_i}{\partial E_i}\frac{dE_i}{dY} + \sum_i \frac{\partial C_i}{\partial W_i}\frac{\partial W_i}{\partial P_i}\frac{dP_i}{dY}.$$

Let $MPC_i = \frac{\partial C_i}{\partial W_i}$ denote the MPC of household i, $\gamma_i = \frac{dE_i}{dY}\frac{Y}{E_i}$ denote the elasticity of household i's nominal income to aggregate output, and $\kappa_i = \frac{dP_i}{dY}\frac{Y}{P_i}$ denote the elasticity of household i's CPI to aggregate output. We get:

$$\begin{aligned} MPC &= \sum_i MPC_i \frac{E_i}{P_i Y}\gamma_i - \sum_i MPC_i \frac{E_i}{P_i Y}\kappa_i \\ &= [\overline{\gamma}\overline{MPC} + cov(MPC_i, \gamma_i)] - [\overline{\kappa}\overline{MPC} + cov(MPC_i, \kappa_i)], \end{aligned} \tag{2}$$

where $\overline{MPC}$ is the real-income-weighted average MPC in the economy, $\overline{\gamma}$ is the elasticity of the average dollar of nominal income to aggregate output, $\overline{\kappa}$ is the elasticity of the average CPI to aggregate output, $cov(MPC_i, \gamma_i)$ is the real-income-weighted covariance between MPCs and income cyclicality, and $cov(MPC_i, \kappa_i)$ is the real-income-weighted covariance between MPCs and inflation cyclicality.

Equation (2) includes two special cases: (1) all households have a nominal income elasticity equal to the average (i.e., $cov(MPC_i, \gamma_i) = 0$), and (2) all households have a consumer price elasticity equal to the average ($cov(MPC_i, \kappa_i) = 0$). We refer to the aggregate MPC in which both conditions hold as the benchmark MPC:

$$MPC^b = (\overline{\gamma} - \overline{\kappa})\overline{MPC}.$$

The actual MPC deviates from the benchmark when the nominal income of high-MPC households is more exposed to movements in aggregate output, $cov(MPC_i, \gamma_i) > 0$, or when the CPI of high-MPC households are more exposed to movements in aggregate output, $cov(MPC_i, \kappa_i) > 0$.

To see the role of these covariances in the amplification of a monetary policy shock, consider a total derivative of the market-clearing condition:

$$dY = \sum_i \frac{\partial C_i}{\partial Y} dY + \sum_i \frac{\partial C_i}{\partial r} dr.$$

Let Ω denote the multiplier—the equilibrium effect of a change in r on aggregate income. Define $\Omega_D = \frac{\partial C}{\partial r}$ as the change in autonomous expenditure change in response to the change in r. Substituting the aggregate MPC from equation (2) gives:

$$\Omega = \frac{dY}{dr} = \frac{\Omega_D}{1 - MPC} = \frac{\Omega_D}{1 - MPC^b - cov(MPC_i, \gamma_i) + cov(MPC_i, \kappa_i)}. \tag{3}$$

The aggregate multiplier is larger when $cov(MPC_i, \gamma_i) > 0$ and smaller when $cov(MPC_i, \kappa_i) > 0$. These covariances govern whether the presence of heterogeneous agents amplifies or dampens the

aggregate response to a policy shock.

3 Data and methodology

We use several datasets to estimate household-specific inflation and MPC. Household consumption data from the CES is linked to item-level price data from the BLS to estimate household inflation cyclicality. To estimate MPCs, we use income and employment information from the PSID, following the estimation procedure outlined by Patterson (2023). Finally, we match the estimated MPCs and CPI by aligning shared household characteristics across datasets.

3.1 CES data

The CES microdata provide detailed information on households' expenditure across detailed product categories through two modules: the Interview and the Diary. The Interview module reports households' purchases in all categories and covers about 30,000 households annually, with each household interviewed for up to four consecutive quarters. The Diary module collects responses to detailed questions about daily purchases, such as groceries, and covers about 10,000 households per year, with data collected on a weekly basis. Overall, there are questions on 350 distinct expenditure categories in the Interview module and on 250 distinct grocery and related categories in the Diary module.

The majority of households do not report expenditures in all possible categories in a given year. Additionally, a different set of households is surveyed in the Interview and Diary files, meaning that the full consumption profile covering both Diary and Interview module expenditures for any particular household is never observed. This prevents us from computing expenditure shares for each household. Instead, we derive consumption profiles for groups of households based on household characteristics.[1] This method not only offers more comprehensive consumption profiles but also allows us to match the derived inflation to MPC computed from the PSID by household groups

3.2 Measuring CPIs in CES

We calculate CPIs for groups of households with characteristics $x_{i,t}$ as defined in equation (4). This is preferable to deriving CPIs at the individual-household level because, as explained above, the full consumption profile with all product categories is not observed for individual households. Moreover, grouping households using the same set of characteristics as in the MPC estimation equation (4) allows us to link MPC estimates directly to group-specific CPIs and estimate cyclical inflation by MPCs.

For each household group h, we calculate consumption expenditure shares at the detailed item level. We then construct group-specific CPI index $p^h_{j,t}$ using the monthly detailed item level prices and

[1] This approach is similar to Cravino et al. (2020) who derive consumption profiles for household groups based on income percentiles.

expenditure weights, closely following the BLS's CPI methodology.

3.3 PSID data

The PSID began in 1968 with a nationally representative sample of over 18,000 individuals in 5,000 households in the United States. It also includes an oversample of low-income households. We include all samples in the PSID to estimate MPCs. Each household was interviewed annually from 1968 to 1996 and biennially thereafter. The interviews gather information on household demographics, labor market experiences, housing status, wealth, and other information. Here, we highlight key variables used in our analysis, with additional details on data and sample provided in the Online Appendix.

Expenditure: The primary expenditure variable available in most of the PSID sample is food expenditure. Following Patterson (2023), we infer total expenditure from food expenditure. Using total expenditure, rather than just food expenditure, to estimate MPC is crucial, as food expenditure's responsiveness to income changes may not reflect that of overall expenditure. This discrepancy arises because food, as a share of total expenditure, varies across income levels, and food stamps can further distort the income-expenditure relationship. To estimate total expenditure, we use CES data to model the demand for food as a function of durable and nondurable consumption, relative prices, and household characteristics. We then invert this demand function to predict total expenditure from food expenditure and household characteristics in the PSID sample.

Wealth: Our analysis uses the biannual information on household assets and debt starting from 1999. [2] We calculate net wealth following the classification in Aguiar et al. (2024). Total net wealth is the sum of liquid net wealth and illiquid net wealth. The former is the sum of checking and savings balances, money market funds, certificates of deposit, treasury bills, and stocks outside of pension funds, minus debts in the form of credit and store cards, student loans, medical or legal debt, and debt owed to the family members. The latter is the sum of home and other real estate equity, IRA or pension holdings, non-government bonds, insurance equity, and the net value of businesses, farms, and vehicles.

We identify low-wealth households following Kaplan et al. (2014), referred to as hand-to-mounth (H2M) households. These households have liquid wealth that is either below or close to a borrowing limit, or nearly zero. Specifically, they either possess negative liquid wealth with an absolute value exceeding 16.5 percent of their annual earnings or have liquid wealth equivalent to a week or less of earnings.[3] In line with Aguiar et al. (2024), we measure earnings as after-tax income, which includes wages, transfers, net profits from businesses or farms, and net income from assets, minus federal and state income tax liabilities calculated by TAXSIM. For homeowners, we include 6 percent of the home

[2] The wealth module was first introduced in 1984. It was included again in 1989, 1994, 1999, and every wave since then. Information on assets and debt is available in PSID biannually from 1999.

[3] These households may not qualify as H2M based on total net wealth, but they are liquidity constrained.

value as implicit rent, while excluding related property taxes, mortgage interest, and home insurance costs

3.4 Measuring MPCs in PSID

We estimate the MPCs for households with different characteristics, building on an extensive literature beginning with Galí et al. (2007), who estimates consumption decline upon unemployment. We estimate

$$\Delta C_{i,t} = \sum_x (\beta_x \Delta E_{i,t} \times x_{i,t-1} + \alpha_x x_{i,t-1}) + \eta_{i,t} + \epsilon_{i,t}, \tag{4}$$

where i and t denote household and time respectively. C is the total consumption expenditure imputed from CES as explained in Section 3.3. E is labor income. x is a household characteristic as in the estimation of group-specific CPI above. $\eta_{i,t}$ is state-by-year fixed effects. ϵ is the error term. Δ denotes a two-year change because the PSID data are biannual after 1997.

To identify the causal effect of earnings on consumption, we instrument change in earnings using unemployment as an income shock. To do this, we restrict the sample to workers employed in year $t-2$ and instrument $\sum_x \beta_x \Delta E_{i,t} \times x_{i,t-1}$ with $\sum_x \beta_x u_{i,t} \times x_{i,t-1}$, where $u_{i,t}$ is a dummy variable that takes a value of 1 if the worker is unemployment in year t. An advantage of using unemployment as an instrument is that it represents a significant and persistent shock (Gruber, 1997). Consequently, it captures the income variation most relevant for understanding monetary policy shocks, which have also been shown to be persistent (Coibion and Gorodnichenko, 2012).

Using the estimated β_x from equation (4), we calculate MPC for households with characteristic x_i as:

$$\widehat{MPC}_{i,t} = \sum_x \widehat{\beta}_x x_{i,t}. \tag{5}$$

3.5 Defining household groups

Choosing the number of dimensions in household characteristics for estimating group-specific CPIs involves a trade-off between sample size and the comprehensiveness of heterogeneity exploration. For instance, if we slice the sample by one dimension of household characteristics (e.g., income quintile), we ensure a sufficiently large sample in each group. However, this approach may overlook other characteristics that jointly determine household consumption. One way to capture these joint determinants is to create more granular groups based on multiple dimensions. The drawback is that the sample size for each group diminishes as we add more dimensions. Since detailed expenditure data is required to calculate group-specific CPIs, and full consumption profiles for all items are not available for individual households, it is essential to maintain a sufficiently large sample in each group to ensure the accuracy of the CPI estimates.

To address this trade-off, we employ two alternative approaches in our estimation. In the first approach, we prioritize sample size by estimating group-specific CPIs separately for each dimension of characteristics. We calculate the corresponding MPC for each group using a bivariate version of equation (5). To pool all available information for estimating cyclical inflation inequality, we impute CPI and MPC for synthetic bins of households based on combinations of multidimensional characteristics.[4] We consider 80 bins spanning four dimensions: education, income, race, and wealth. Alternatively, we examine 40 bins across three dimensions: education, income, and race.

3.6 Estimating cyclical inflation inequality

The steps outlined above allow us to calculate group-specific CPIs and link them to the MPCs of the groups imputed from the PSID. The imputation is necessary because MPCs are not directly observed in the CES. As explained by Patterson (2023), this approach relies on several assumptions about MPCs. First, MPCs out of unemployment are assumed to be the same as MPCs out of other income shocks, despite differences in signs and magnitudes. Second, MPCs depend solely on the set of household characteristics included in estimation equation (4). Third, conditional on household characteristics, MPCs are stable over the business cycle.

Under these assumptions and utilizing estimation results from Section 3.2, we examine heterogeneity in inflation cyclicality across households. We estimate:

$$p_{i,t+s} - p_{i,t} = \alpha_s + \lambda_s H_i \times Shock_t + \sum_{k=1}^{K} \theta_{s,k} H_i \times Shock_{t-k} + \sum_{j=1}^{J} \kappa_{s,j}(p_{i,t-j} - p_{i,t-j-1}) + \eta_t + \tau_i + \epsilon_{i,t+s}, \quad (6)$$

where i and t denote household group and time (month) respectively. $s >= 0$ denotes the estimation horizon. p_i is the log of group-specific CPI. H_i represent the (single or multidimensional) household characteristics. $Shock_t$ is a measure of monetary policy shock, as discussed below. The coefficient λ_s captures how prices in period $t+s$, relative to period t, respond to a monetary policy shock, varying with household characteristics. We refer to this coefficient as the inflation response to monetary policy shocks. In our model, we control for 24 lags of the shocks (i.e., $K = 24$) and 6 lags of the monthly group-specific inflation (i.e., $J = 6$). We also control for time fixed effects η_t and group fixed effects τ_i.

Our measure of monetary policy shock is derived from the series developed by Bu et al. (2021). This measure has several appealing features for our analysis: it encompasses both conventional and unconventional monetary policies, is largely unpredictable, and contains no significant information

[4] Specifically, the bins consist of synthetic household groups whose MPC and CPI are MPC-weighted averages of the corresponding characteristics from the single-dimensional estimates.

effects. As demonstrated in Bu et al. (2021) and confirmed by our calculations of aggregate CPI from CES microdata, aggregate prices fall significantly in response to tightening monetary shocks identified in this series. This aggregate price response aligns with the predictions of conventional macroeconomic theory, making it a suitable benchmark for understanding inflation heterogeneity.

We next estimate the correlation between MPCs and inflation cyclicality using the following regression:

$$p_{i,t+s} - p_{i,t} = \alpha_s + \lambda_s \widehat{MPC}_i \times Shock_t + \sum_{k=1}^{K} \theta_{s,k} \widehat{MPC}_i \times Shock_{t-k} + \sum_{j=1}^{J} \kappa_{s,j}(p_{i,t-j} - p_{i,t-j-1}) + \eta_t + \tau_i + \epsilon_{i,t+s}, \tag{7}$$

where i and t denote household group and time (month) respectively. $s >= 0$ denotes the estimation horizon. p_i and $Shock_t$ are defined as above. $\widehat{MPC}_i$ is the MPC of the group i estimated from equation (5). The coefficient λ_s captures how prices in period $t+s$, relative to period t, respond to a monetary policy shock, varying with MPCs. We estimate this equation using household groups from the two approaches discussed in Section 3.5. Since $\widehat{MPC}_i$ is a generated regressor in equation (7), we draw inferences using standard errors obtained from a bootstrap procedure.

4 Empirical results

In this section, we first document household heterogeneity in inflation cyclicality and MPCs, followed by results on the covariance between MPCs and inflation cyclicality.

Figure 1 illustrates the heterogeneity in inflation cyclicality. It plots MPCs against group-specific inflation responses to monetary policy shocks over a 24-month horizon, estimated from equation (6) (i.e., $\widehat{\lambda}_{24}$). This coefficient represents changes in prices (in percentage points) in response to a 100-basis-point tightening of monetary policy. The inflation response of the aggregate CPI is normalized to 0. As discussed earlier, the aggregate price response to tightening shocks in the Bu et al. (2021) series is negative, as predicted by conventional theory. Thus, groups with inflation responses lower than the average exhibit higher-than-average inflation cyclicality.

The estimated inflation responses vary across different levels of education, income, and wealth. Households with lower education, lower income, or lower wealth tend to have higher inflation cyclicality. Specifically, in response to a 100-basis-point tightening monetary policy shock, the CPI of households with less than a high school diploma decreases by 3.4 percentage points more than the aggregate CPI, whereas the CPI of households with a college or postgraduate degree decreases by 1.8 percentage points less than the aggregate CPI.

We observe similar trends across the income distribution, with lower-income households exhibiting

higher inflation cyclicality compared to higher-income households. H2M households also have higher inflation cyclicality than non-H2M households, though the difference is smaller compared to the variations observed across education and income. Notably, we do not find significant differences in inflation cyclicality between Black and White households.

Figure 2 provides evidence of MPC heterogeneity. The aggregate MPC, estimated from the full sample of households, is 0.54, which aligns with previous estimates based on similar variations due to the loss of labor income. However, the aggregate MPC conceals substantial variations across education, income, race, and wealth.[5]

An interesting pattern observed in the figure is that nearly all the subsamples presented exhibit MPCs significantly different from the average. Overall, households with lower educational attainment, lower income, lower wealth, and Black households have higher-than-average MPCs. In terms of magnitude, our household-level estimates (based on household income and the race of the household head) are very similar to those estimated by Patterson (2023) at the individual level. Among our estimated categories, households within the 20^{th} to 40^{th} percentiles of the income distribution have the highest MPC, estimated at 1.3. In contrast, households in the top quintile of the income distribution have the lowest MPC, estimated at below 0.3.

While Patterson (2023) does not estimate MPCs by education and wealth, we find the MPC of non-H2M households to be 0.3, which is lower than the MPC of H2M households at 0.7. Interestingly, the MPCs in the lowest quintile of the income distribution do not follow the monotonic relationship observed in higher-income groups. This may reflect the unique composition of this group, which includes many students and young adults. Their relatively low MPCs compared to those in the 20^{th} to 60^{th} percentiles of the income distribution may indicate their future income potential.

Estimates of inflation cyclicality across other dimensions of household characteristics are presented in Appendix Figure A1, which includes age, family size, H2M status based on liquid assets, housing tenure, marital status, and region.[6] We find some evidence of heterogeneity for certain types of households: one-member households and households with a married head of household exhibit higher inflation cyclicality than their counterparts. Estimates of MPC based on these characteristics are shown in Appendix Figure A2. H2M households (based on liquid assets) have higher MPCs than non-H2M households, while households that own a home with a mortgage show lower MPCs compared to renters. However, the combined results from Appendix Figures A1 and A2 do not suggest strong evidence of cyclical inflation inequality across these dimensions of household characteristics. In the following results, we focus on household education, income, race, and wealth (based on the H2M definition by Kaplan et al., 2014).

Figure 3 combines the results shown in Figures 2 and 1, linking MPCs to inflation cyclicality. It

[5]Education and race reflect the characteristics of the household head. Income encompasses earnings from all household members, while wealth includes aggregate household assets and debt.

[6]The definition of H2M based on liquid assets follows Zeldes (1989) and includes households with liquid assets less than two months of labor income.

reveals a positive relationship between inflation cyclicality and MPCs. Households with lower educational attainment have higher MPCs and greater inflation cyclicality than those with higher educational attainment. A similar pattern is observed across income and wealth groups.

Figure 4 plots the coefficient for the inflation response to monetary policy shocks estimated using 80 synthetic bins that span education, income, race, and wealth. Recall that this coefficient represents changes in prices (in percentage points) in response to a 100-basis-point tightening of monetary policy. The results suggest that, on average, households with higher MPCs experience greater CPI deflation compared to those with lower MPCs—confirming the visual impression from Figure 3. For example, the coefficient of 1.6 in the 12-month horizon indicates that households whose MPCs are 25 percent below the average (0.56) experience deflation approximately 0.5 percentage points larger than households whose MPCs are 25 percent above the average. Figure 5 illustrates this relationship between MPC and cyclical inflation by plotting group-specific responses to monetary policy shocks (i.e., $\widehat{\lambda}_{24}$ of equation (6)) against the group MPCs.

Appendix Figure A3 presents results estimated using 40 synthetic bins that span education, income, and race. The result in the 12-month horizon are nearly identical to those in Figure 4, while results over a longer horizon are slightly larger. Appendix Figure A4 shows results estimated based on empirical multi-dimensional bins spanning education, income, and race. Again, the results in the 12-month horizon are very similar to those in Figure 4, but the longer-horizon results are slightly smaller and less precise.

5 Implication for the aggregate multiplier

What are the quantitative implications of cyclical inflation inequality for the aggregate multiplier? The theoretical framework outlined in Section 2 establishes a direct link between cyclical inflation inequality and the multiplier. To apply the empirical estimates from Section 4 within this framework, we need to adjust for the share of labor income in total income. In the derivation of equation (3), we assume that all output is earned by households as labor income. However, empirically, households also receive income from other sources, such as capital income and profits. Let α_l denote the share of output earned by households in the form of labor income. Let MPC and MPC^{nl} denote the MPC out of labor and non-labor income respectively. We show in Appendix C that the benchmark and actual multipliers adjusted for non-labor income are given by:

$$MPC^b = \alpha_l(\overline{\gamma} - \overline{\kappa})\overline{MPC} + (1 - \alpha_l\overline{\gamma} - \alpha_l\overline{\kappa})MPC^{nl}, \tag{8}$$

$$MPC^a = \alpha_l[(\overline{\gamma} - \overline{\kappa})\overline{MPC} + cov(MPC_i, \gamma_i) - cov(MPC_i, \kappa_i)] + [1 - \alpha_l(\overline{\gamma} - \overline{\kappa})]MPC^{nl}. \tag{9}$$

We calibrate several parameters based on the literature. We set the labor share at two-thirds and the MPC out of non-labor income at 0.19. The covariance of MPC with the elasticity of labor income

is set at 0.099, and the aggregate elasticity of labor income to GDP, $\overline{\gamma}$, is set to 1.03. These values closely follow those used by Patterson (2023), who quantifies the impact of cyclical income inequality on the multiplier. Using the same parameters allows us to compare the relative roles of cyclical income inequality and cyclical inflation inequality. Additionally, we set the aggregate elasticity of prices to GDP, $\overline{\kappa}$, to be 0.6.

Table 1 presents the results. Column 1 quantifies the cyclical inflation inequality channel while shutting down the cyclical income inequality channel. Panel A infers the covariance between MPC and inflation cyclicality. Rows 1 and 2 show the mean and variance of MPCs in the sample. Row 3 provides the estimated correlation between MPC and the shock, derived from equation (7) over a 12-month horizon. Row 4 shows the aggregate output elasticity to the shock. Together, these imply a covariance between MPC and inflation elasticity of about 0.02.[7] Panel B illustrates the implications for the aggregate MPC. The benchmark MPC, defined in equation (8), is 0.3. After adjusting for the covariance, the actual MPC is 0.29, which is approximately 3.5 percent smaller. Panel C shows the implications for the multiplier based on the definition in equation (3). The difference in MPC from Panel B leads to a decrease in the multiplier of 1.5 percent.

We compare this with the cyclical income channel in Column 2 of Table 1. Patterson (2023) estimates the covariance between MPC and income cyclicality to be just shy of 0.1. Row 7 indicates an actual MPC of 0.37, or 22 percent larger than the benchmark. The corresponding change in the multiplier is an increase of 10.5 percent.

Column 3 reflects the combined effects of cyclical inflation inequality and cyclical income inequality. With both channels active, the actual MPC is 0.35, or 15 percent larger than the benchmark. The corresponding multiplier is 7 percent larger than the benchmark.

Overall, cyclical inflation heterogeneity significantly affects the transmission of shocks. While cyclical income inequality amplifies the output effect of shocks by just over 10 percent, as shown in previous literature, we demonstrate that cyclical inflation inequality dampens this effect by one-third, resulting in a combined amplification effect of real income inequality of 7.5 percent.

6 Conclusion

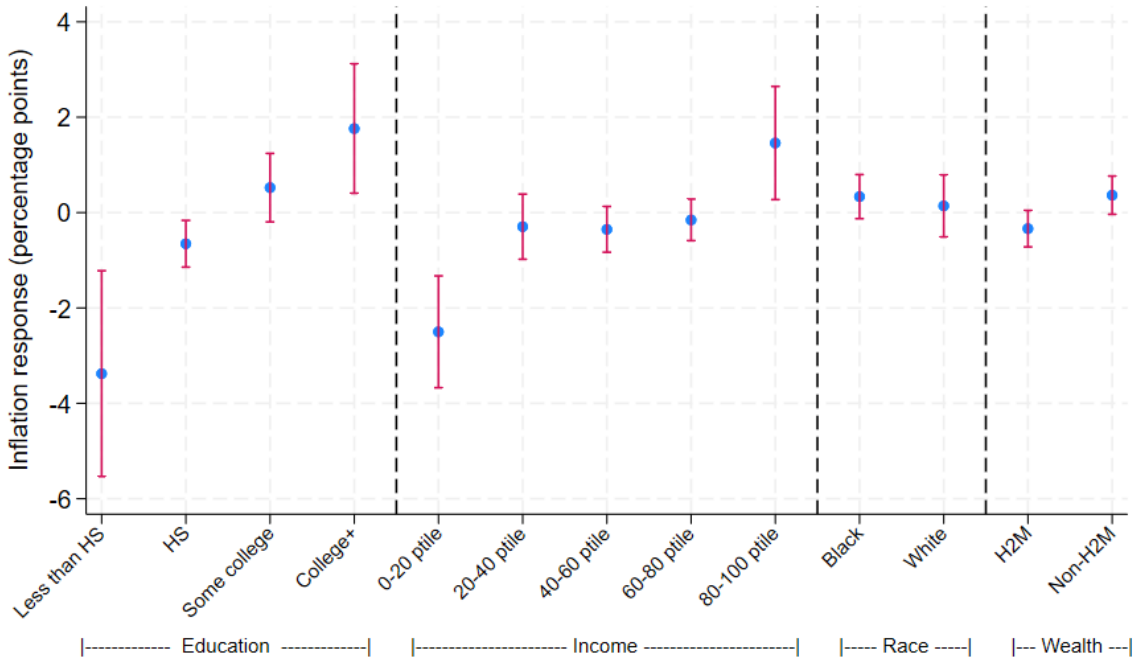
We document evidence of cyclical inflation inequality across households. Drawing on microdata related to household income and consumption, we show that CPI and MPC vary with household characteristics. We estimate a positive covariance between households' MPCs and their elasticity of consumer prices in response to aggregate shocks, indicating that households with a higher MPC experience greater cyclical inflation.

This cyclical inflation inequality is a key component of cyclical real income inequality——a critical

[7] This is calculated as $cov(MPC_i, \kappa_i) = cov(MPC_i, \frac{\eta_i}{\phi}) = \frac{1}{\phi} cov(MPC_i, \eta_i)$, where I is the monetary policy shock. $\eta_i = \frac{dP_i}{dI}\frac{I}{P_i}$ is the price elasticity to the shock and $\phi = \frac{dY}{dI}\frac{I}{Y}$ is the output elasticity to the shock.

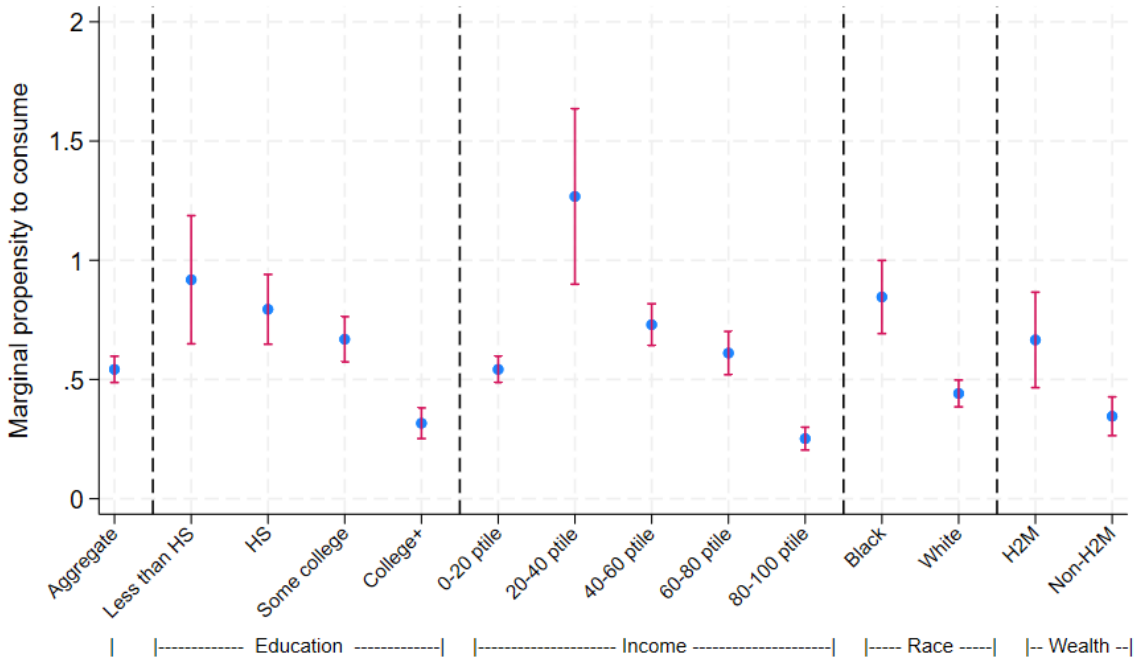
object in heterogeneous-agent models governing the transmission of policy shocks. Our estimates show that cyclical inflation inequality offsets about a third of the amplification effect from cyclical income inequality, as demonstrated in previous literature. The combined effect results in a 7 percent increase in the aggregate MPC compared to a benchmark with equal exposure to inflation and nominal income.

Figure 1: Inflation response to monetary policy shock by household group



Notes: This figure plots the coefficient for group-specific inflation response to monetary policy shock (i.e., $\widehat{\lambda}_{24}$) of equation (6) for selected household characteristics. The bars show 90 and 10 percent confidence intervals. Source: CES, PSID, and authors' calculations.

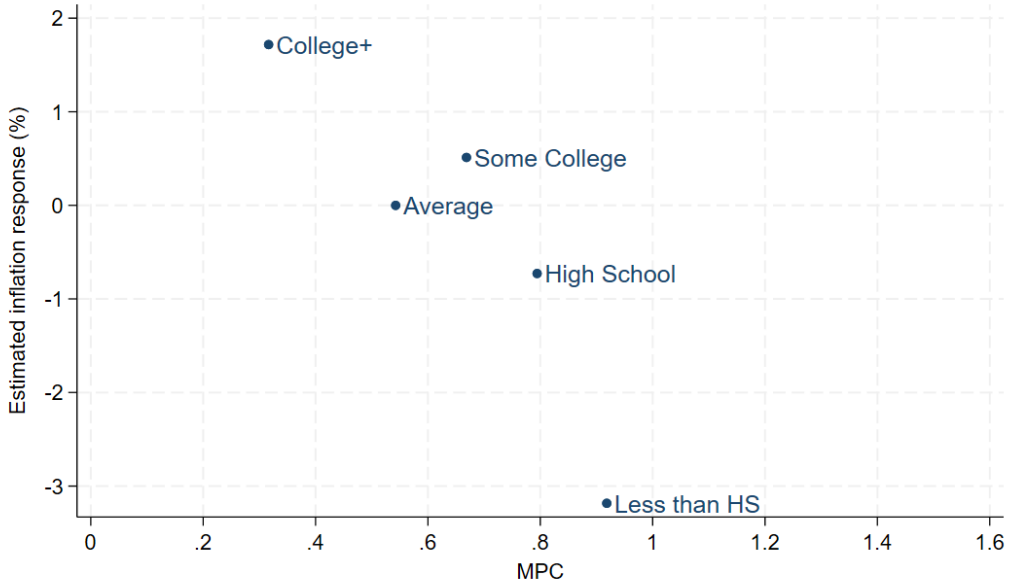
Figure 2: Estimated MPCs by household group



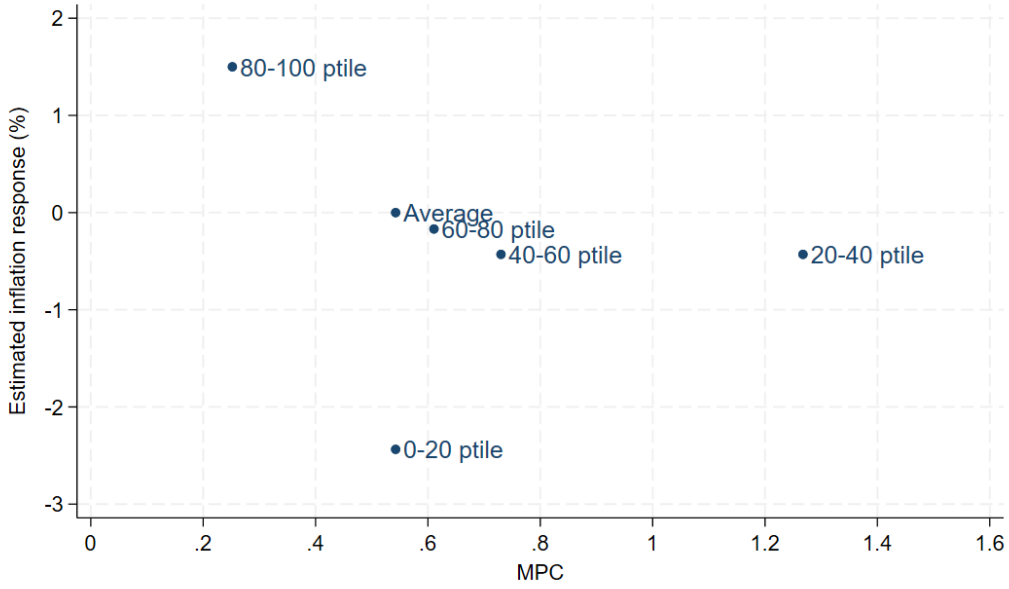
Notes: This figure plots estimated MPCs from equation (5). The bars show 90 and 10 percent confidence intervals. Source: CES, PSID, and authors' calculations.

Figure 3: Inflation response to monetary policy shock and MPC by household group

(a) Education



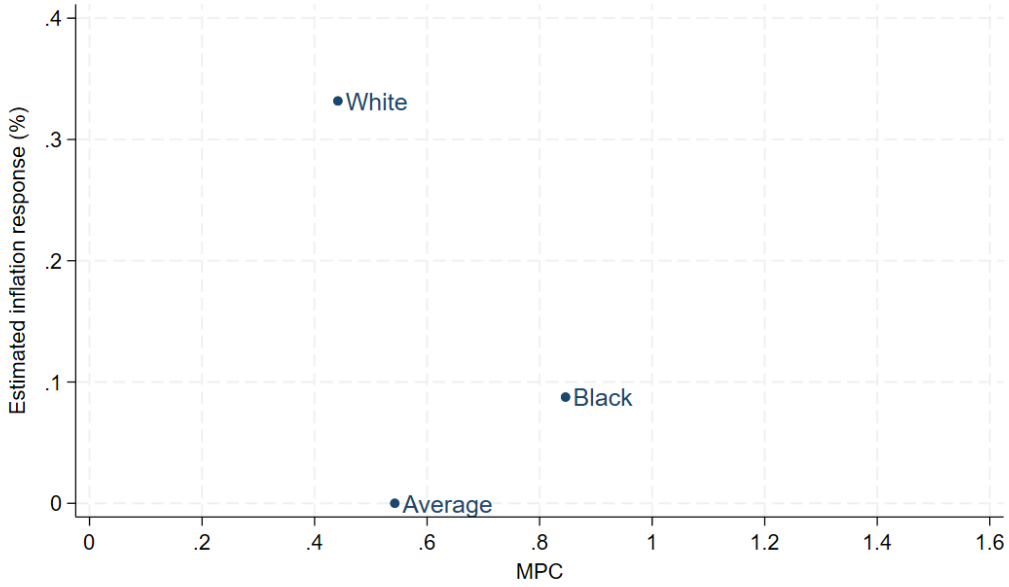
(b) Income



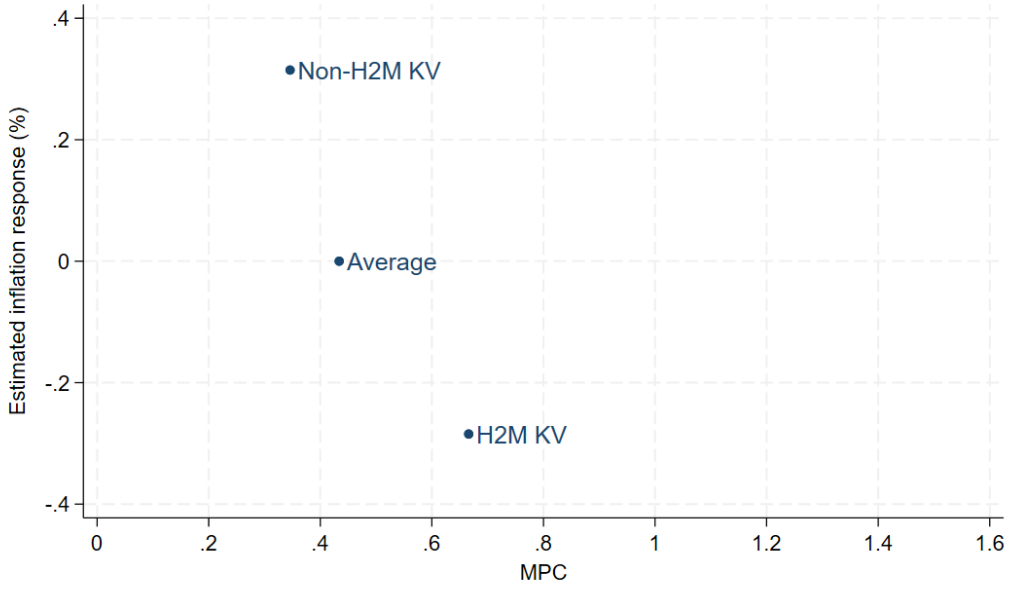
Notes: The y-axis plots the coefficient for group-specific inflation response to monetary policy shock (i.e., $\widehat{\lambda}_{24}$) of equation (6) for selected household groups. The x-axis plots estimated MPCs of the corresponding groups estimated from equation (5). Source: CES, PSID, and authors' calculations.

Figure 3: Inflation response to monetary policy shock and MPC by household group

(c) Race

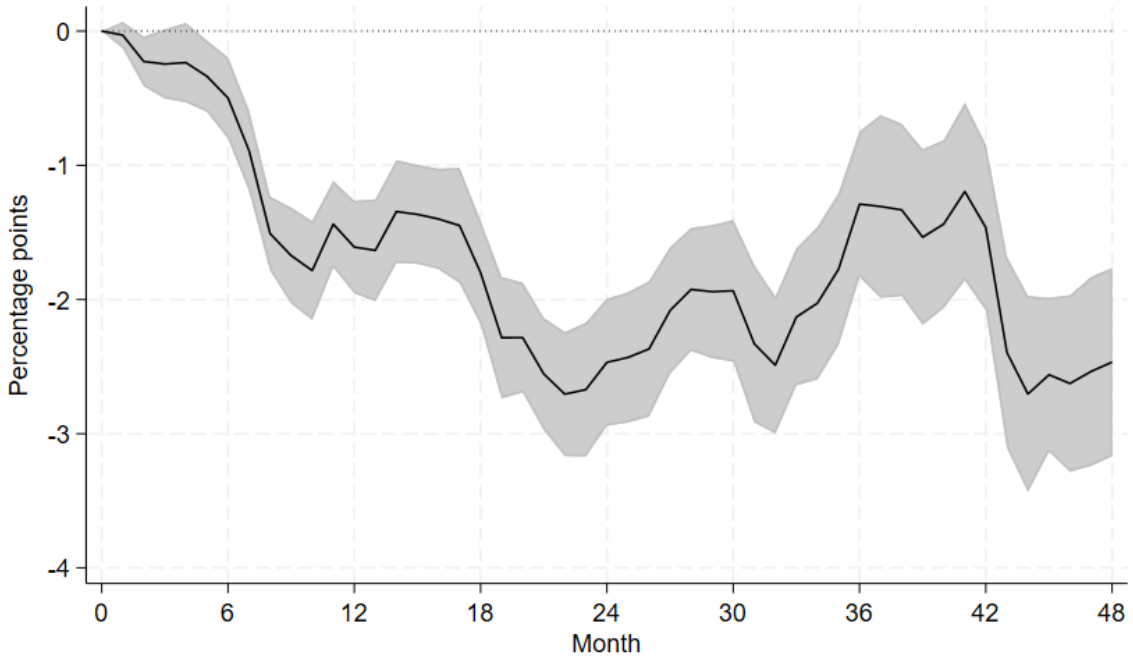


(d) Wealth



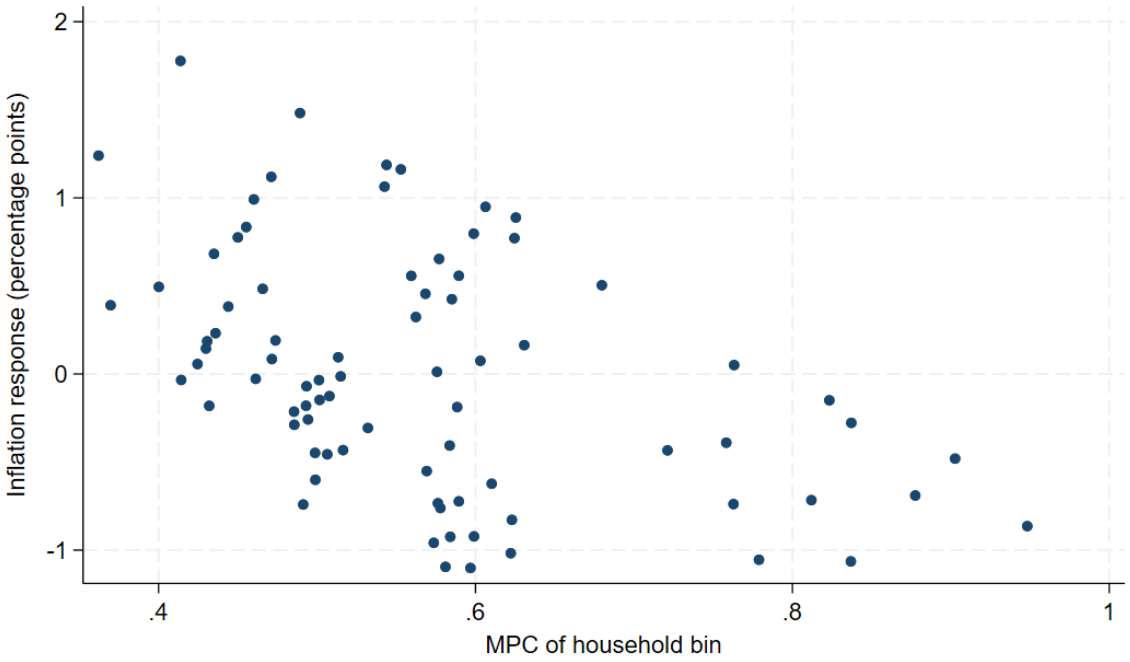
Notes: The y-axis plots the coefficient for group-specific inflation response to monetary policy shock (i.e., $\widehat{\lambda_{24}}$) of equation (6) for selected household characteristics. The x-axis plots estimated MPCs of the corresponding groups estimated from equation (5). Source: CES, PSID, and authors' calculations.

Figure 4: Inflation response to monetary policy shock by MPC



Notes: The figure plots the coefficient for inflation response to monetary policy shock (i.e., $\widehat{\lambda_s}$) of equation (7), estimated using 80 bins. Source: CES, PSID, and authors' calculations.

Figure 5: Inflation response to monetary policy shock and MPC by bins



Notes: The y-axis plots the coefficient for inflation response to monetary policy shock (i.e., $\widehat{\lambda_{24}}$) of equation (6). The x-axis plots the MPCs of the corresponding groups estimated from equation (5). Each dot represents a synthetic household bin defined by the combination of education (less than high school, high school, some college, and college+), income (5 quintile), race (Black and White), and wealth (hand-to-mouth and non-hand-to-mouth). Source: CES, PSID, and authors' calculations.

Table 1: Implications for the multiplier

	Inflation (1)	Earnings (2)	Combined (3)
Panel A. Covariance			
1. average(MPC)	0.560		
2. var(MPC)	0.016		
3. Coefficient on MPC $\times shock$	1.633		
	[1.180, 2.086]		
4. cov(MPC$_i, \kappa_i$)	0.015		0.015
	[0.011, 0.020]		[0.011, 0.020]
5. cov(MPC$_i, \gamma_i$)		0.099	0.099
		[0.040, 0.164]	[0.040, 0.164]
Panel B. Change in aggregate MPC			
6. Benchmark MPC (MPC^b)	0.301	0.301	0.301
7. Actual MPC (MPC^a)	0.290	0.367	0.346
8. Percent change in MPC	-3.4%	22.0%	15.2%
	[-2.5%, -4.4%]	[22.0%, 36.4%]	[3.8%, 27.7%]
Panel C. Change in multiplier			
9. Benchmark multiplier	1.430	1.430	1.430
10. Actual multiplier	1.409	1.580	1.530
11. Percent change in multiplier	-1.5%	10.5%	7.0%
	[-1.1%, -1.8%]	[3.9%, 18.6%]	[1.7%, 13.5%]

Notes: Source: CES, PSID, and authors' calculations.

References

Acharya, Sushant and Keshav Dogra, "Understanding HANK: Insights From a PRANK," *Econometrica*, May 2020, *88* (3), 1113–1158.

Aguiar, Mark, Mark Bils, and Corina Boar, "Who Are the Hand-to-Mouth?," *The Review of Economic Studies*, 05 2024, p. rdae056.

Argente, David and Munseob Lee, "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 2021, *19* (2), 913–952.

Auclert, Adrien, "Monetary policy and the redistribution channel," *American Economic Review*, 2019, *109* (6), 2333–67.

Bilbiie, Florin O., "The New Keynesian cross," *Journal of Monetary Economics*, 2020, *114* (4), 90–108.

Bu, Chunya, John Rogers, and Wenbin Wu, "A unified measure of Fed monetary policy shocks," *Journal of Monetary Economics*, 2021, *118* (C), 331–349.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico, "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *The Review of Economic Studies*, 01 2019, *87* (1), 102–129.

Coibion, Olivier and Yuriy Gorodnichenko, "Why Are Target Interest Rate Changes So Persistent?," *American Economic Journal: Macroeconomics*, May 2012, *4* (4), 126–62.

Cravino, Javier and Andrei A. Levchenko, "The Distributional Consequences of Large Devaluations," *The American Economic Review*, 2017, *107* (11), 3477–3509.

__ **, Ting Lan, and Andrei Levchenko**, "Price stickiness along the income distribution and the effects of monetary policy," *Journal of Monetary Economics*, 2020, *110* (C), 19–32.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, March 2007, *5* (1), 227–270.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, March 1997, *87* (1), 192–205.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 107–136.

Jaravel, Xavier, "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector," *The Quarterly Journal of Economics*, 2019, *134* (2), 715–783.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, July 2014, *82* (4), 1199–1239.

__ **and Sam Schulhofer-Wohl**, "Inflation at the household level," *Journal of Monetary Economics*, 2017, *91* (C), 19–38.

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

__ **, Giovanni L. Violante, and Justin Weidner**, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *45* (1 (Spring), 77–153.

Lan, Ting, Lerong Li, and Minghao Li, "Cyclical inequality in the cost of living and implications for monetary policy," mimeo 2024.

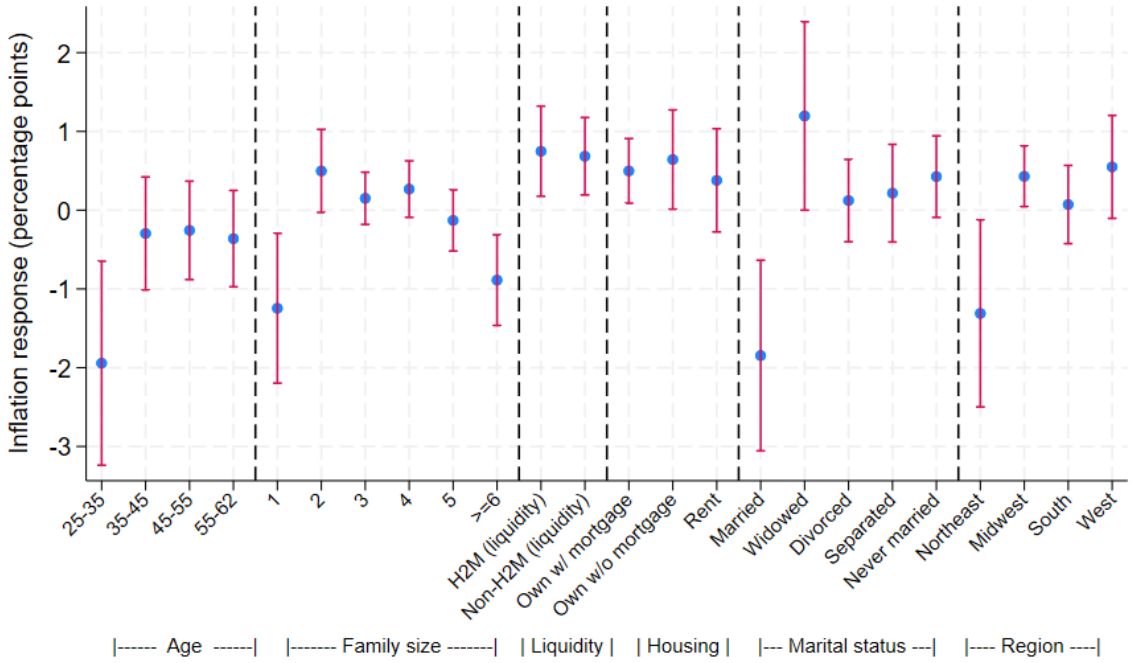
McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, October 2016, *106* (10), 3133–3158.

Patterson, Christina, "The matching multiplier and the amplification of recessions," *American Economic Review*, 2023, *113* (4), 982–1012.

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

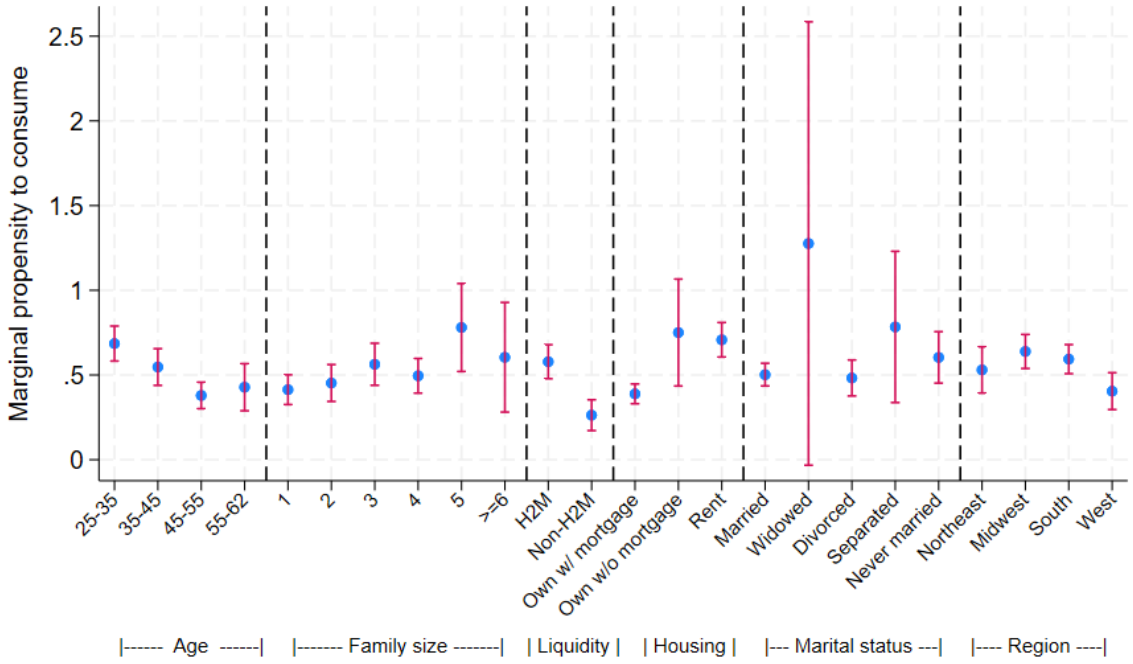
A Additional results

Figure A1: Inflation response to monetary policy shock by household group



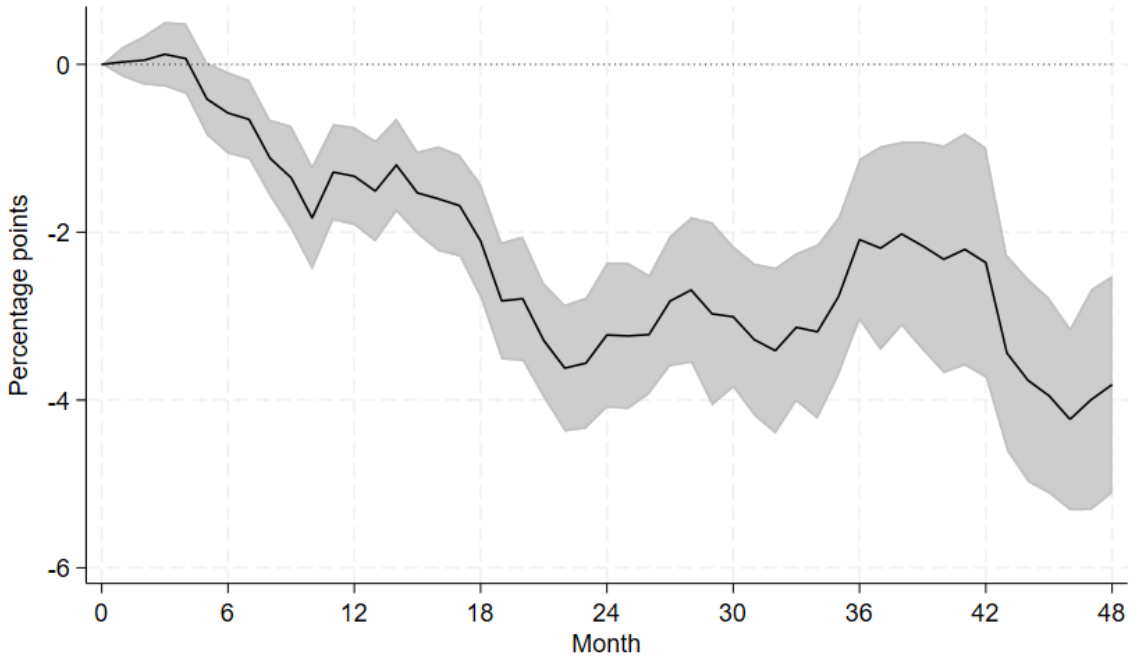
Notes: This figure plots estimated MPCs from equation (5). The bars show 90 and 10 percent confidence intervals. Source: CES, PSID, and authors' calculations.

Figure A2: Estimated MPC by household group



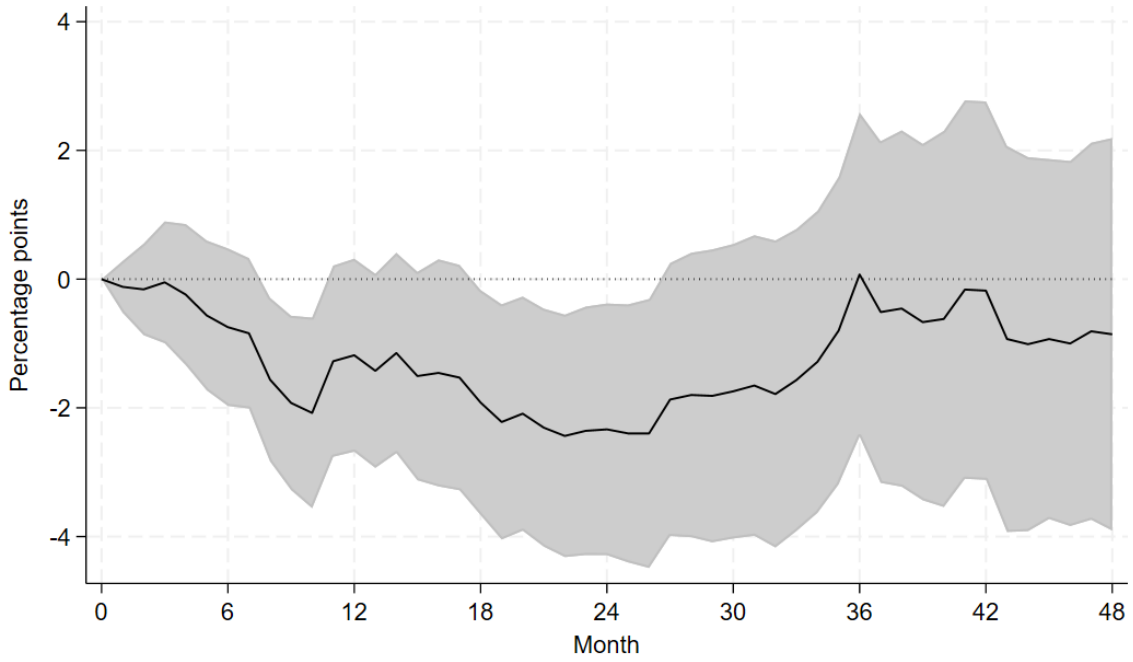
Notes: This figure plots estimated MPCs from equation (5). The bars show 90 and 10 percent confidence intervals. Source: CES, PSID, and authors' calculations.

Figure A3: Inflation response to monetary policy shock and MPC



Notes: The figure plots the coefficient for inflation response to monetary policy shock (i.e., $\widehat{\lambda_s}$) of equation (7), estimated using 40 bins. Source: CES, PSID, and authors' calculations.

Figure A4: Inflation response to monetary policy shock and MPC



Notes: The figure plots the coefficient for inflation response to monetary policy shock (i.e., $\widehat{\lambda_s}$) of equation (7), estimated using 40 bins. Source: CES, PSID, and authors' calculations.

B A simplified framework with real income

Consider an economy with heterogeneous households. The consumption function of household i is given by $C_i(W_i, r)$, where W_i is real income as a functions of aggregate output $W_i(Y)$ and r is real interest rate. The market clearing condition equates aggregate output to aggregate consumption: $Y = C$. The aggregate MPC is $MPC = \frac{dC}{dY}$, which sums across households as:

$$MPC = \sum_i \frac{\partial C_i}{\partial Y} = \sum_i \frac{\partial C_i}{\partial W_i}\frac{dW_i}{dY}.$$

Define $MPC_i = \frac{\partial C_i}{\partial W_i}$ as the MPC of household i, $\chi_i = \frac{dW_i}{dY}\frac{Y}{W_i}$ as the elasticity of household i's real income to aggregate output. We get:

$$MPC = \sum_i MPC_i \frac{W_i}{Y}\chi_i = \overline{\chi}\overline{MPC} + cov(MPC_i, \chi_i).$$

where $\overline{MPC}$ is the real-income-weighted average MPC in the economy, $\overline{\chi}$ is the elasticity of the average dollar of real income to aggregate output. $cov(MPC_i, \chi_i)$ is the real-income-weighted covariance between households' MPCs and real income elasticity to aggregate output.

When all households have a real income elasticity equal to the average so $cov(MPC_i, \gamma_i) = 0$, the benchmark MPC given by:

$$MPC^b = \overline{\chi}\overline{MPC}.$$

The actual MPC deviate from the benchmark when the real income of high-MPC households are more exposed to movements in aggregate output, $cov(MPC_i, \chi_i) > 0$.

A total derivative of the market-clearing condition gives:

$$dY = \sum_i \frac{\partial C_i}{\partial Y}dY + \sum_i \frac{\partial C_i}{\partial r}dr.$$

Let Ω denote the multiplier—the equilibrium effect of a change in r on aggregate income. Define $\Omega_D = \frac{\partial C}{\partial r}$ as the change in autonomous expenditure change in response to the change in r. Substituting the aggregate MPC gives:

$$\Omega = \frac{dY}{dr} = \frac{\Omega_D}{1 - MPC} = \frac{\Omega_D}{1 - MPC^b - cov(MPC_i, \chi_i)}.$$

This result highlights the the role of the heterogeneity in real income in the amplification of shocks. The aggregate multiplier is larger when $cov(MPC_i, \chi_i) > 0$. It is smaller when $cov(MPC_i, \chi_i) > 0$. Bilbiie (2020) thoroughly discusses this mechanism in a TANK model.

The framework in Section 2 reduces to this simplified framework with $W_i = \frac{E_i}{P_i}$.

C Marginal propensity of consumption with non-labor income adjustment

Consider the case where is aggregate output is earned by households in the form of labor and non-labor income: $Y = W + R$, where W and R are real labor and non-labor income respectively. The consumption function of household i is given by $C_i(W_i(E_i, P_i), W_i(K_i, P_i), \theta_i)$, where E_i is nominal labor income, K_i is nominal non-labor income, and P_i is the household-specific price index. Nominal income (E_i and K_i) and P_i are functions of aggregate output Y. The aggregate MPC is given by:

$$\begin{aligned} MPC &= \frac{\partial C}{\partial W}\frac{\partial W}{\partial Y} + \frac{\partial C}{\partial R}\frac{\partial R}{\partial Y} \\ &= \frac{\partial C}{\partial W}\frac{\partial W}{\partial Y} + \frac{\partial C}{\partial R}(1 - \frac{\partial W}{\partial Y}) \end{aligned}$$

Summing across households, we get:

$$MPC = \sum_i \frac{\partial C_i}{\partial W_i}\frac{\partial W_i}{\partial E_i}\frac{dE_i}{dY} + \sum_i \frac{\partial C_i}{\partial W_i}\frac{\partial W_i}{\partial P_i}\frac{dP_i}{dY} + \sum_i \frac{\partial C_i}{\partial R_i} - \sum_i \frac{\partial C_i}{\partial R_i}(\frac{\partial W_i}{\partial E_i}\frac{dE_i}{dY}) + \sum_i \frac{\partial C_i}{\partial R_i}\frac{\partial W_i}{\partial P_i}\frac{dP_i}{dY}$$

Define $\alpha_l = \sum_i \frac{E_i}{P_i Y}$ as the share of labor income in total income, $\gamma_i = \frac{dE_i}{dY}\frac{Y}{E_i}$ as the elasticity of household i's nominal income to aggregate output, $\kappa_i = \frac{dP_i}{dY}\frac{Y}{P_i}$ as the elasticity of household i's CPI to aggregate output, and $\frac{\partial C_i}{\partial R_i}$ as the MPC out of non-labor income, we get:

$$\begin{aligned} MPC = {} & \alpha_l(\overline{\gamma} - \overline{\kappa})\overline{MPC} + \alpha_l[cov(MPC_i, \gamma_i) - cov(MPC_i, \gamma_i)] \\ & + \overline{MPC^{nl}} - \alpha_l(\overline{\gamma} - \overline{\kappa})\overline{MPC^{nl}} - \alpha_l[cov(MPC_i^{nl}, \gamma_i) - cov(MPC_i^{nl}, \kappa_i)]. \end{aligned}$$

Abstracting from heterogeneity in MPC out of non-labor income give the empirical equivalence of the actual multiplier:

$$MPC^a = \alpha_l[(\overline{\gamma} - \overline{\kappa})\overline{MPC} + cov(MPC_i, \gamma_i) - cov(MPC_i, \kappa_i)] + [1 - \alpha_l(\overline{\gamma} - \overline{\kappa})]MPC^{nl},$$

and the corresponding benchmark multiplier:

$$MPC^b = \alpha_l(\overline{\gamma} - \overline{\kappa})\overline{MPC} + [1 - \alpha_l(\overline{\gamma} - \overline{\kappa})]MPC^{nl}.$$