MPCs in an Emerging Economy: Evidence from Peru[*][†][‡]

Seungki Hong[§]

October 14, 2022

Abstract

This paper estimates the marginal propensity to consume (MPC) out of transitory income shocks for an emerging economy, Peru, using its nationally representative household survey. The mean quarterly MPC across Peruvian income deciles is 0.204, which translates to a mean annualized MPC of 0.545-0.592 under both model-free and model-based annualization methods. To compare Peruvian and U.S. MPCs reflecting the different reference periods of the underlying surveys, I employ a standard incomplete-market model. Two striking differences emerge. First, the mean annual MPC in Peru is three times as large as that in the U.S. Second, the MPCs are substantially more heterogeneous over income deciles in Peru than in the U.S. The model predicts that precautionary saving behavior drives both the higher mean MPC and stronger MPC heterogeneity in Peru.

***JEL classification*:** D12, D31, E21, F41

***Keywords*:** consumption, MPC, precautionary saving, emerging economy

[†]This paper is a revised version of Chapters 1 and 2 of my Ph.D. dissertation at Columbia University. I am highly indebted to Martin Uribe, Stephanie Schmitt-Grohe, and Andres Drenik for their constant guidance, support, and suggestions. This paper has greatly benefited from the insightful comments of the editor and two anonymous referees. I am also grateful to Hassan Afrouzi, Paul Bouscasse, Christopher Cotton, Nils Gornemann, Emilien Gouin-Bonenfant, Yoon Jo, Jennifer La'O, Colin Sullivan, Ken Teoh, Mengxue Wang, and Michael Woodford for their helpful comments and discussions. All errors are my own.

[‡]Declaration of Interest: I worked on this paper during my Ph.D. studies at Columbia University and my employment at Purdue University. My Ph.D. studies were financially supported jointly by the Korean Government Scholarship from the National Institute for International Education of the Republic of Korea and the Graduate School of Arts and Sciences (GSAS) Fellowship from Columbia University. The Korean Government, Columbia University, and Purdue University do not have any influence on or involvement with the research this paper describes.

[§]Department of Economics, Krannert School of Management, Purdue University. E-mail: hong397@purdue.edu

1 Introduction

Macroeconomists have recently discovered that microlevel household consumption-saving behavior and its heterogeneity significantly matter for aggregate dynamics, the effects of monetary and fiscal policies, and their transmission mechanisms. Several works in this line of research have found that these macroeconomic outcomes crucially depend on households' marginal propensity to consume (MPC) out of transitory income shocks. For this reason, researchers often discipline macroeconomic models using MPC estimates from nationally representative samples (see, for instance, the estimates from Johnson, Parker, and Souleles (2006), Parker, Souleles, Johnson, and McClelland (2013), and Blundell, Pistaferri, and Preston (2008)).

However, these estimates are all based on micro data from developed economies. In the context of emerging economies, there is also a long tradition of estimating MPCs using natural experiments (such as weather shocks) or randomized control trials (RCTs), but these studies restrict their focus to certain regions and targeted groups because of the shocks they exploit. As a consequence, nationally representative MPC estimates for an emerging economy, which are suitable for disciplining a macroeconomic model, have been absent.

This paper fills this gap by estimating MPCs out of transitory income shocks using a Peruvian nationally representative household survey (Encuesta Nacional de Hogares, ENAHO[1]). Specifically, I estimate the MPC within each income decile using Blundell et al. (2008)'s method. Given that the method identifies a group's MPC assuming that MPC is homogeneous within the group, the observations are grouped by an unpredictable component of income (or, equivalently, residual income), as precautionary saving theory predicts MPC heterogeneity over this income.

The estimation yields a mean quarterly MPC across Peruvian income deciles of 20.4%. This quarterly estimate translates to a mean annualized MPC of 59.2% and 54.5% under model-free and model-based annualization methods, respectively. For model-free annualization, I use Auclert (2019)'s MPC annualization formula, which essentially approximates dynamic consumption responses to a transitory shock as exponentially dying out over time. For model-based annualization, I calibrate a standard incomplete-market, life-cycle model by targeting the quarterly MPC estimates and compute true annual MPCs in the model.

Compared to the annual U.S. MPC estimates from the Panel Study of Income Dynamics (PSID), the annualized Peruvian MPCs exhibit two starkly different patterns. First, the Peruvian MPCs are substantially greater overall than the U.S. MPCs. The mean of the annualized Peruvian MPCs across income deciles is approximately seven times as large as the mean of the U.S. estimates (8.3%). Second, the Peruvian MPCs are substantially more heterogeneous over income deciles than the U.S. MPCs.

However, this comparison (model-free comparison hereafter) has a critical problem in that it imposes an asymmetric assumption regarding the time frame. The reference period is a quarter in ENAHO, while it is a year in the PSID. Accordingly, when estimating MPCs using Blundell

[1] Instituto Nacional de Estadística e Informática (2004-2016).

et al. (2008)'s method, Peruvian households are assumed to receive income shocks and make consumption decisions quarterly, while U.S. households are assumed to do so yearly. If we instead assume that both Peruvian and U.S. households are under the same quarterly time frame, my U.S. MPC estimates are subject to a 'time aggregation problem' that Crawley (2020) points out: when Blundell et al. (2008)'s method is applied to annual data while households receive income shocks and make consumption decisions at a higher frequency, the method significantly underestimates consumption sensitivity to transitory shocks. Part of the large gap between U.S. and Peruvian MPCs under the model-free comparison can be attributable to this time aggregation problem.

To compare Peruvian and U.S. MPCs reflecting the different reference periods of the underlying surveys, I employ a standard incomplete-market, life-cycle model. Two model economies are calibrated under the same quarterly time frame: one is fitted to Peru by targeting the Peruvian quarterly MPC estimates, while the other is fitted to the U.S. by targeting the U.S. annual MPC estimates.[2] In particular, the U.S. annual MPC estimates are targeted as follows. First, I simulate quarterly income and consumption series from the model and convert them into annual series by aggregating them over every four quarters. Then, the model counterparts of the annual estimates are obtained by applying to the simulated annual series the same estimation procedure applied to the PSID. Lastly, a calibration is conducted such that the model counterparts are as close as possible to the estimates.

I compare model-predicted annual MPCs between the two model economies. I find that the two main differences observed in the model-free comparison remain robust. First, Peruvian MPCs are substantially greater overall than U.S. MPCs. The cross-country mean MPC gap is somewhat smaller than the gap in the model-free comparison because the time aggregation problem is fixed. Nevertheless, the model-based comparison predicts that the mean annual MPC in Peru is 3.0 times as large as that in the U.S. Second, the model-based comparison also predicts that annual MPCs are substantially more heterogeneous over income deciles in Peru than in the U.S.

I examine the causes of the two main differences through the lens of the model and find that households' precautionary saving behavior is the main driver of both the higher mean MPC and stronger MPC heterogeneity in Peru. In this model, households face idiosyncratic income risk and thus exhibit precautionary saving behavior because i) they fear the realization of a low-income path (Kimball, 1990) and ii) they also worry about being hit by borrowing limits (Huggett, 1993). These precautionary saving motives become weaker as households' cash-on-hand (or, equivalently, total currently available resources) increases. Moreover, households with stronger precautionary saving motives exhibit higher MPC because a positive transitory shock relaxes these precautionary saving motives. Peruvian households have substantially stronger precautionary saving motives than U.S. households because they accumulate far fewer liquid assets in the model. Moreover, when households move down from higher to lower income deciles, precautionary saving motives become stronger more rapidly in Peru than in the U.S.

This paper is related to multiple strands of literature. Methodologically, this paper borrows

[2] In addition, the labor income processes and interest rates are calibrated using Peruvian and U.S. data.

one of the main approaches to obtain nationally representative MPC estimates. Three approaches have been widely accepted in this literature: (i) exploiting a nationwide natural experiment of income shocks, (ii) imposing a theory-guided covariance structure on the joint dynamics of income and consumption, and (iii) directly using answers to survey questions asking how much households would spend out of hypothetical income shocks. Well-known works in each of the approaches include Johnson et al. (2006) and Parker et al. (2013) for the first approach[3], Blundell et al. (2008) and Kaplan, Violante, and Weidner (2014b) for the second one, and Jappelli and Pistaferri (2014) for the third one, among many others. I use the second approach because its data requirements are met by ENAHO.

This paper is also related to the literature estimating MPCs for certain regions and targeted groups of emerging economies using natural experiments or RCTs. For example, Paxson (1992) uses rainfall shocks in rural Thailand and identifies the consumption responses of rice farmers. Haushofer and Shapiro (2016) and Egger, Haushofer, Miguel, Niehaus, and Walker (2019) use randomized cash transfers in a small study district within rural Kenya and identify the consumption responses of poor households in the district. My paper complements these studies by estimating MPCs with a semistructural method and, more importantly, using a nationally representative sample.

As noted above, I interpret the MPC estimates through the lens of a standard incomplete-market, life-cycle model. In this aspect, this paper is indebted to a longstanding literature in which many different versions of the model are developed, analyzed, and applied for a variety of topics. Some well-known works in this literature include Carroll (1997), Huggett (1996), Hubbard, Skinner, and Zeldes (1995), and Kaplan and Violante (2010), among many others. The model is used for multiple purposes in this paper. As one of the purposes, I check the performance of the MPC estimator using model simulation in the spirit of Kaplan and Violante (2010).

There is a burgeoning literature examining how microlevel household consumption-saving behavior and its heterogeneity matter for aggregate dynamics, the effects of monetary and fiscal policies, and their transmission mechanisms. Several studies find that these macroeconomic outcomes are significantly affected by high MPCs and correspondingly strong precautionary saving behavior in the U.S. Well-known examples include Krueger, Mitman, and Perri (2016), Kaplan, Moll, and Violante (2018), Auclert (2019), Berger, Guerrieri, Lorenzoni, and Vavra (2018), McKay, Nakamura, and Steinsson (2016), and Oh and Reis (2012), among many others.

This paper suggests that MPCs and precautionary saving behavior can affect macroeconomic outcomes more significantly in emerging economies than in developed economies. In this regard, this paper supports the importance of recent endeavors to expand the literature to open and emerging economies, such as Auclert, Rognlie, Souchier, and Straub (2021), Guntin, Ottonello, and Perez (2020), Guo, Ottonello, and Perez (2020), Zhou (2021), Ferrante and Gornemann (2021), De Ferra, Mitman, and Romei (2020), Villalvazo (2021), Oskolkov (2022), and Hong (2020). I ex-

[3] Johnson et al. (2006) and Parker et al. (2013) estimate MPCs for U.S. households using tax rebates in 2001 and economic stimulus payments in 2008, respectively. Their samples are nationally representative because most U.S. households were eligible for the tax rebates and the stimulus payments.

pect this paper to contribute to this recent wave of research by providing useful off-the-shelf target moments for the key object, MPC.[4]

The remainder of this paper is organized as follows. Section 2 explains the MPC estimation method and the numerical method of solving and simulating the incomplete-market, life-cycle model. Section 3 describes how the data are processed. Section 4 estimates quarterly Peruvian MPCs, calibrates the model to the Peruvian economy using the estimates, evaluates the performance of the MPC estimator in the model, and annualizes the quarterly MPCs using both model-free and model-based methods. Section 5 estimates annual U.S. MPCs, calibrates the model to the U.S. economy using the estimates, and compares Peruvian and U.S. MPCs. Section 6 examines the main driver of the cross-country MPC differences in the model. Section 7 considers several alternative specifications either in the model simulation or in the MPC estimation. Section 8 provides a discussion on external validity. Section 9 concludes.

2 Methods

2.1 MPC Estimation

I use an extended version of Blundell et al. (2008)'s method to estimate MPCs. In this subsection, I describe Blundell et al. (2008)'s original method and the extension.

The method begins by imposing a structural income process. Let $Y_{i,t}$ be household i's income in age t, $Z_{i,t}$ be its observable characteristics, and $Z'_{i,t}\varphi^y_t$ be the predictable component of log income $\log Y_{i,t}$.[5] Assume that the unpredictable component of log income, $y_{i,t}$ $(:= \log Y_{i,t} - Z'_{i,t}\varphi^y_t)$, is composed of a permanent component $P_{i,t}$ and a transitory component $\epsilon_{i,t}$ as follows.

$$
\begin{gathered}
y_{i,t} = P_{i,t} + \epsilon_{i,t}, \\
P_{i,t} = P_{i,t-1} + \zeta_{i,t}, \\
\zeta_{i,t} \sim_{iid} (0, \sigma^2_{ps}), \quad \epsilon_{i,t} \sim_{iid} (0, \sigma^2_{tr}), \quad \text{and} \quad (\zeta_{i,t})_t \perp (\epsilon_{i,t})_t
\end{gathered}
$$

in which $(x_t)_t$ represents time series $(\cdots, x_{t-1}, x_t, x_{t+1}, \cdots)$.

Similarly, let $C_{i,t}$ be household i's consumption in age t, $Z'_{i,t}\varphi^c_t$ be the predictable component of log consumption $\log C_{i,t}$, and $c_{i,t}$ $(:= \log C_{i,t} - Z'_{i,t}\varphi^c_t)$ be its unpredictable component. Blundell et al. (2008)'s partial insurance parameter to transitory income shocks, ψ_G for a group G is defined as follows.

$$
\psi_G := \frac{cov[\Delta c_{i,t}, \epsilon_{i,t} | (i,t) \in G]}{cov[\Delta y_{i,t}, \epsilon_{i,t} | (i,t) \in G]}. \tag{1}
$$

[4] Among the papers mentioned here, some papers already use the MPC estimates from an earlier version of this paper. Auclert et al. (2021), Oskolkov (2022), and Hong (2020) calibrate their models by targeting the estimates. Zhou (2021) uses the estimates to validate the model by checking whether MPCs, which are untargeted moments in his model, are close to the estimates.

[5] In implementation, I include as observable characteristics education, ethnicity, employment status, region, cohort, household size, number of children, urban or rural area, the existence of members other than heads and spouses earning income, and the existence of persons who do not live with but are financially supported by the household. Among these characteristics, education, ethnicity, employment status, and region are allowed to have time-varying effects.

Parameter ψ_G is an elasticity of consumption with regard to income when an income change is caused by a transitory income shock.

Blundell et al. (2008) further assume that households do not foresee future shocks so that $c_{i,t}$ is independent of $(\zeta_{i,t+j}, \epsilon_{i,t+j})_{j\geq 1}$. When this assumption holds and grouping G is independent of $(\zeta_{i,t+j}, \epsilon_{i,t+j})_{j\geq 0}$, ψ_G can be identified by

$$\psi_G = \frac{cov[\Delta c_{i,t}, \Delta y_{i,t+1} | (i,t) \in G]}{cov[\Delta y_{i,t}, \Delta y_{i,t+1} | (i,t) \in G]}. \tag{2}$$

Intuitively, ψ_G is obtained by running an IV regression in which $\Delta y_{i,t+1} = \zeta_{i,t+1} + \epsilon_{i,t+1} - \epsilon_{i,t}$ is used as an instrument.

Equation (2) is the original identification equation of Blundell et al. (2008). However, this equation is not immediately applicable to the data used in this paper, as the Peruvian micro data provide year-over-year growth of quarterly income and consumption and the U.S. micro data provide two-year-over-two-year growth of annual income and consumption. To accommodate these data structures, I extend the identification equation (2) under the same set of assumptions.

Under the imposed income process and the assumption that households do not foresee future shocks, we have

$$cov[\Delta c_{i,t}, \Delta y_{i,t+j} | G] = cov[\Delta y_{i,t}, \Delta y_{i,t+j} | G] = 0, \quad j \geq 2.$$

Therefore, ψ_G can also be expressed as

$$\psi_G = \frac{cov[\Delta c_{i,t} + \Delta c_{i,t-1}, \Delta y_{i,t+1} + \Delta y_{i,t+2} | (i,t) \in G]}{cov[\Delta y_{i,t} + \Delta y_{i,t-1}, \Delta y_{i,t+1} + \Delta y_{i,t+2} | (i,t) \in G]} = \frac{cov[\Delta^2 c_{i,t}, \Delta^2 y_{i,t+2} | (i,t) \in G]}{cov[\Delta^2 y_{i,t}, \Delta^2 y_{i,t+2} | (i,t) \in G]}$$

in which $\Delta^K x_t := x_t - x_{t-K}$ for time series $(x_t)_t$. Extending it further, we can also express ψ_G as

$$\psi_G = \frac{cov[\Delta^K c_{it}, \Delta^K y_{i,t+K} | (i,t) \in G]}{cov[\Delta^K y_{it}, \Delta^K y_{i,t+K} | (i,t) \in G]}, \quad K \geq 1. \tag{3}$$

I use equation (3) to identify ψ_G. Finally, I convert elasticity ψ_G into the MPC by multiplying the mean-consumption-to-mean-income ratio in group G as follows.[6]

$$MPC_G = \psi_G \frac{E[C_{i,t} | (i,t) \in G]}{E[Y_{i,t} | (i,t) \in G]}. \tag{4}$$

In implementation, I define $\kappa_G := \frac{E[C_{i,t}|(i,t)\in G]}{E[Y_{i,t}|(i,t)\in G]}$ and estimate $(\kappa_G, \alpha_G, \psi_G)$ using the following moment conditions and the GMM method.

[6]The consumption-income ratio at time t (*i.e.*, the ratio at the time when the exploited shocks are realized) should be used because $\psi_G = \frac{\partial \log C_{i,t} / \partial \epsilon_{i,t}}{\partial \log Y_{i,t} / \partial \epsilon_{i,t}} \approx \frac{(\partial C_{i,t} / \partial \epsilon_{i,t}) / C_{i,t}}{(\partial Y_{i,t} / \partial \epsilon_{i,t}) / Y_{i,t}}$.

$$
\begin{aligned}
E[\kappa_G Y_{i,t} - C_{i,t} | (i,t) \in G] = 0, \\
E[\Delta^K c_{i,t} - \alpha_G - \psi_G \Delta^K y_{i,t} | (i,t) \in G] = 0, \quad \text{and} \\
E[\Delta^K y_{i,t+K}(\Delta^K c_{i,t} - \alpha_G - \psi_G \Delta^K y_{i,t}) | (i,t) \in G] = 0.
\end{aligned}
\tag{5}
$$

Standard errors are clustered within each household.[7] Using the GMM estimates and their variance-covariance matrix, I obtain the MPC estimate and standard error in group G using equation (4) or, equivalently, $MPC_G = \psi_G \kappa_G$.[8]

A few concerns may arise in applying Blundell et al. (2008)'s method to an emerging economy. First, to justify their partial insurance parameters as good measures of consumption responses to income shocks, Blundell et al. (2008) use an approximated consumption function derived from a standard incomplete market model without borrowing constraints. This justification might not be valid for Peru if a large fraction of households are affected by tight borrowing constraints. In Online Appendix A, I show that this justification can be extended to the case with borrowing constraints by deriving an approximated consumption function from the same model but with the constraints.

Second, the approximated consumption function that Blundell et al. (2008) derive ignores households' precautionary saving motive due to prudence, as defined by Kimball (1990) (or, equivalently, their convexly magnified fear of the realization of a low-income path).[9] The same problem arises in my approximated consumption function derived from a model with borrowing constraints in Online Appendix A. The justification using these approximated consumption functions may not be valid for Peru if households exhibit strong prudent saving. Therefore, instead of resorting to approximated consumption functions, I directly check the performance of the MPC estimator using numerical simulation of a standard incomplete-market model fitted to the Peruvian economy.

2.2 Numerical Simulation

In this paper, I use numerical simulation of a standard incomplete-market, life-cycle model (Carroll (1997), Huggett (1996), Hubbard et al. (1995), and Kaplan and Violante (2010)) for four purposes. First, I examine how well a standard model can match the MPC estimates. Second, I evaluate the performance of the MPC estimator (in terms of the estimates being close to true

[7]Standard errors are clustered because i) the error term $(\Delta^K c_{i,t} - \alpha_G - \psi_G \Delta^K y_{i,t})$ can be autocorrelated and ii) the instrumental variable $\Delta^K y_{i,t+K}$ of observation (i,t) can also be correlated with the error term of observation $(i, t+K)$. For example, in a standard incomplete market model with borrowing constraints, the error term $(\Delta^K c_{i,t} - \alpha_G - \psi_G \Delta^K y_{i,t})$ is greater for household i who becomes constrained between $t-K$ and $t-1$ than for those who do not. Since an asset position changes slowly, this household is also more likely to be constrained between t and $t+K-1$ and thus to have a greater value of the future error term $(\Delta^K c_{i,t+K} - \alpha_G - \psi_G \Delta^K y_{i,t+K})$. Moreover, household i with a lower value of $\Delta^K y_{i,t+K}$ is more likely to be constrained between t and $t+K-1$ and thus to have a greater value of the future error term.

[8]The delta method is used to obtain the standard error of $\psi_G \kappa_G$.

[9]This is because the consumption function is obtained by first-order-Taylor-approximating log marginal utility in Euler equations, while prudence manifests through the second-order terms. See Carroll (1997) and Jappelli and Pistaferri (2017).

MPCs) in a standard model in the spirit of Kaplan and Violante (2010). Third, I also check the performance of Auclert (2019)'s model-free MPC annualization formula (used in subsection 4.4) within a standard model. Fourth, I compare Peruvian and U.S. MPCs taking into account the reference period discrepancy between the underlying surveys through the lens of a standard model.

This subsection outlines the model and the solution method. I consider an overlapping generations economy. A continuum of households are born in each period and live for $T+1$ periods. In each period, households face idiosyncratic income risk and borrowing constraints and can trade a risk-free, one-period bond. At age t_0 ($0 \leq t_0 \leq T$), household i solves the following optimization.

$$E_{t_0} \sum_{t=t_0}^{T} \beta^{t-t_0} \frac{C_{i,t}^{1-\sigma}}{1-\sigma}$$

$$s.t.$$

$$C_{i,t} + A_{i,t} = Y_{i,t} + (1+r)A_{i,t-1}, \quad t_0 \leq t \leq T, \tag{SBC}$$

$$A_{i,t} \geq -\bar{B}, \quad t_0 \leq t \leq T-1, \quad \text{and} \tag{LQC}$$

$$A_{i,T} \geq 0 \tag{TML}$$

in which $C_{i,t}$, $Y_{i,t}$, and $A_{i,t}$ denote the consumption, disposable labor income, and bond holdings of household i, respectively, r denotes the real interest rate on bonds, and $\bar{B}$ denotes a borrowing limit. (SBC), (LQC), and (TML) represent sequential budget constraints, liquidity constraints, and a terminal condition, respectively.

Households' labor income $Y_{i,t}$ is composed of three components in logs, namely, an age-specific deterministic component ω_t, a persistent stochastic component $P_{i,t}$, and a transitory stochastic component $\epsilon_{i,t}$ as follows.

$$\begin{gathered}
Y_{i,t} = \omega_t + y_{i,t}, \quad 0 \leq t \leq T, \\
y_{i,t} = P_{i,t} + \epsilon_{i,t}, \quad 0 \leq t \leq T, \\
P_{i,t} = \rho P_{i,t-1} + \zeta_{i,t}, \quad 1 \leq t \leq T, \\
\zeta_{i,t} \sim_{iid} (0, \sigma_{ps}^2), \quad \epsilon_{i,t} \sim_{iid} (0, \sigma_{tr}^2), \quad P_{i,0} \sim_{iid} (0, \sigma_{P_0}^2), \\
(\zeta_{i,t})_t \perp (\epsilon_{i,t})_t, \quad (\epsilon_{i,t})_t \perp P_{i,0}, \quad \text{and} \quad P_{i,0} \perp (\zeta_{i,t})_t.
\end{gathered}$$

The income process imposed by Blundell et al. (2008)'s method features $\rho = 1$. When calibrating the incomplete-market model, however, I allow ρ to be different than 1 for three reasons. First, the income process fits Peruvian income data much better when ρ is not restricted to 1. Second, as will be shown later in subsection 7.1, ρ is an important determinant of MPC heterogeneity over the unpredictable component of income ($y_{i,t}$) in the model. Third, as shown by Kaplan and Violante (2010) and verified in subsection 4.3, Blundell et al. (2008)'s method can robustly recover true MPCs to transitory shocks even in a model with $\rho < 1$. In subsection 7.1, I examine the case in which ρ is restricted to 1 in the model.

The model is calibrated to Peruvian and U.S. economies in subsections 4.2 and 5.2, respectively. Under each calibration, I solve the model using the method of endogenous grid points developed by Carroll (2006). I use 100 exponentially spaced grid points for assets. The evolution of the persistent income component $P_{i,t}$ is approximated as a Markov chain with age-varying and equally spaced 20 grid points by applying Fella, Gallipoli, and Pan (2019)'s extended version of Rouwenhorst (1995)'s method.[10] The transitory income component $\epsilon_{i,t}$ is approximated with 20 equally spaced grid points by Rouwenhorst (1995)'s original method. Then, I simulate 1,000,000 households in each of the Peruvian and U.S. model economies.[11]

3 Data

3.1 Source

To obtain nationally representative MPC estimates using Blundell et al. (2008)'s method, a micro dataset should satisfy three requirements. First, the dataset should include both income and expenditure. Second, the dataset should have a panel structure such that households appear at least three consecutive times. Third, the sample should be nationally representative. ENAHO is one of the rare datasets, if not the only one, that satisfies all three requirements in an emerging economy. It is the major information source of the quantity indices for the final household expenditure in Peru's national accounts (Instituto Nacional de Estadística e Informática, n.d.); thus, it is nationally representative and includes detailed categories of household expenditure. ENAHO also collects information on detailed sources of household income. ENAHO tracks a subset of annual cross-sectional observations in subsequent years. The panel households are also nationally representative.[12] I use the 2004-2016 waves of ENAHO. These waves provide 11 years of consumption and income growth, as the survey is conducted annually and a panel structure is absent between the 2006 and 2007 waves. Online Appendix B.1 provides more details about ENAHO, including its coverage and nonresponse rates.

To conduct an MPC comparison between emerging and developed economies, I need another micro dataset that satisfies all three requirements discussed above for a developed economy. I choose Kaplan et al. (2014b)'s replication dataset for U.S. households.[13] For the purpose of cross-country comparison, their dataset is relevant for two reasons. First, the sample years are not too different. They use the 1999-2011 waves from the Panel Study of Income Dynamics (PSID), which overlap substantially with my Peruvian sample (waves 2004-2016). Second, they prepare the dataset to estimate Blundell et al. (2008)'s partial insurance parameter to transitory shocks,

[10] Rouwenhorst (1995)'s original method is designed to discretize a stationary process. Fella et al. (2019) extend it to a nonstationary process, such as the process of a persistent income component in a life-cycle model.

[11] The algorithm for the numerical simulation of the model is written in Python. Some Python functions for certain operations, such as linear interpolation, one-step forward iteration of a household distribution, the construction of exponentially spaced grid points, the calculation of a Markov chain's invariant distribution, and the implementation of Rouwenhorst (1995)'s original method, are borrowed from Auclert, Bardoczy, Rognlie, and Straub (2021).

[12] Most panel households appear two or three times, while the maximum number of appearances is six in my sample.

[13] Kaplan, Violante, and Weidner (2014a)

which is the same object upon which I base my MPC estimates.

3.2 Variable Construction

The baseline consumption for both Peruvian and U.S. samples includes nondurable goods and a subset of services, as in many other studies on household consumption, such as Attanasio and Weber (1995) and Kocherlakota and Pistaferri (2009). Following these studies, I exclude health and education expenses due to their durable nature. I also exclude nonpurchased consumption, such as donations, food stamps, in-kind income, and self-production.[14] Nominal consumption is deflated with the Consumer Price Index (CPI) series.

The baseline income for both Peruvian and U.S. samples is composed of disposable labor income and transfers, as in Blundell et al. (2008) and Kaplan et al. (2014b). Capital income is excluded because we do not want to falsely attribute endogenous capital income changes to unexpected income shocks. In ENAHO, labor income and capital income are not distinguishable in self-employment income. Following Diaz-Gimenez, Quadrini, and Rios-Rull (1997) and Krueger and Perri (2006), I split self-employment income into a labor income component and a capital income component using the ratio between unambiguous labor income and unambiguous capital income in the sample.[15] I exclude the imputed components of missing income from Peruvian incomes, as these components might blur the identification of income shocks.[16] Nominal income is again deflated with the CPI series.[17]

In ENAHO, reference periods vary over both expense and income items. More importantly, individual households report more than 97 percent (in value) of expense items and income items, respectively, under reference periods shorter than or equal to the previous three months, on average. Given this feature of the data, I construct quarterly consumption and income by excluding expense and income items with a longer reference period than the previous three months. Expense and income items with a shorter reference period than the previous three months are scaled up to quarterly expense and income, respectively. Since panel households are interviewed annually, we can obtain the year-over-year growth of quarterly consumption and income only from the Peruvian sample. In the PSID, the reference period is fixed to one year, but households are interviewed biannually during the sample years. Therefore, we can obtain two-year-over-two-year growth of annual consumption and income only from the U.S. sample.

Online Appendix B.2 provides more details on the variable construction.

[14]In subsection 7.3.1, I conduct a robustness check by including nonpurchased consumption.

[15]In my ENAHO sample, the ratio of '(unambiguous labor income)/(unambiguous labor income + unambiguous capital income)' is 0.819. This ratio is slightly lower but quite similar to the ratio in the U.S., 0.864, which Diaz-Gimenez et al. (1997) and Krueger and Perri (2006) use.

[16]I cannot do the same for U.S. incomes because the imputed components are not distinguishable in Kaplan et al. (2014b)'s dataset. In subsection 7.3.1, I conduct a robustness check by consistently including the imputed components in Peruvian incomes.

[17]The U.S. incomes and expenses are already deflated with the U.S. CPI series in Kaplan et al. (2018)'s dataset. To deflate Peruvian nominal incomes and expenses, I use the Peruvian CPI series from Banco Central de Reserva del Perú.

3.3 Sample Selection

My sample selection aims to remove problematic observations for the MPC estimation while maximizing the comparability with existing studies, such as Blundell et al. (2008) and Kaplan et al. (2014b).

The sample selection for ENAHO proceeds as follows. First, I start with observations made over at least two consecutive surveys. Second, I drop observations if interviews for the same household in two consecutive surveys are conducted in different months. Third, there are panel observations that are likely to connect two different households by failing to distinguish one that moves out and the other that moves into the same address. Such observations are detected and dropped.[18] Fourth, I drop observations when household heads are replaced. Fifth, I drop observations when interviewers categorize them as an incomplete response. Sixth, I drop observations if household heads are younger than 25 or older than 65. Seventh, I drop observations if any of the observable characteristics used to extract predictable components of income and consumption are missing. Eighth, I drop observations with nonpositive income or consumption. Ninth, I drop observations that have too much value in imputed income components. Similarly, I drop observations that include too much value in expense items or income items with a longer reference period than the previous three months.[19] Tenth, I drop income outliers determined by extreme income growth.[20] The final sample is composed of 47,210 observations, 21,988 pairs of two consecutive observations, and 7,509 triplets of three consecutive observations. Online Appendix B.3 provides more details of the sample selection, including the number of observations dropped in each step.

For the U.S. sample, I adopt Kaplan et al. (2014b)'s sample selection with a few minor revisions. Online Appendix B.3 discusses details of the minor revisions, a remaining difference between the Peruvian and U.S. sample selections, and a robustness check regarding the difference.

3.4 Residual Income Grouping

Blundell et al. (2008)'s method identifies a group's MPC assuming that households in the group have an identical MPC. Therefore, it is important to group households such that households exhibiting very different MPCs belong to different groups.

In this paper, I use the deciles of residual income ($y_{i,t}$) distribution.[21] I do so because the unpredictable component of income $y_{i,t}$ bears income risk and standard incomplete-market models predict MPC heterogeneity over $y_{i,t}$. Specifically, when lower $y_{i,t}$ is realized, households' precautionary saving motives due to income uncertainty become stronger, and thus, they exhibit higher MPC. Importantly, I do not group households by their actual income ($Y_{i,t}$) because the predictable

[18] Online Appendix B.4 provides details of the procedure.

[19] Specifically, for each $(x, y) \in \{$(expense items with a longer reference period than the previous three months, baseline consumption measure), (income items with a longer reference period than the previous three months, baseline income measure), (income imputation, baseline income measure)$\}$, observations are dropped if $x/(x+y) > 0.05$.

[20] Specifically, I define income outliers as households whose income growth is in the range of the extreme 1 percent (0.5 percent at the top and 0.5 percent at the bottom) in the calendar-year subsamples at least once.

[21] Among well-known studies, Berger et al. (2018) also use residual income grouping when estimating MPCs.

component $Z'_{i,t}\varphi^y_t$ captures permanent income heterogeneity (such as income heterogeneity due to education, ethnicity, and region). Such permanent heterogeneity does not induce precautionary saving, and thus MPC heterogeneity, when it is incorporated in these models.[22] [23]

To construct residual income deciles, I sort U.S. observations within each calendar year and Peruvian observations within each calendar quarter in accordance with the reference period of each sample (a year for the U.S. sample and a quarter for the Peruvian sample).[24] Survey weights are used when computing the quantile of each observation.

When estimating MPCs using moment conditions (5), both current log income growth $\Delta^K y_{i,t}$ and future log income growth $\Delta^K y_{i,t+K}$ are used for each observation (i,t). Therefore, three residual incomes $y_{i,t-K}$, $y_{i,t}$, and $y_{i,t+K}$ are available for each observation (i,t). Among them, I use $y_{i,t-K}$ to determine the decile of observation (i,t) to ensure that the grouping is independent of current and future income shocks $(\zeta_{i,t+j}, \epsilon_{i,t+j})_{j\geq 0}$ and therefore does not bias the estimates.

When constructing the baseline income measure of Peruvian households, I use only reported incomes from the past three months. In particular, I exclude i) income items reported under a longer reference period than the previous three months and ii) the imputed values of missing incomes. Moreover, in the sample selection, I drop observations with too much value in excluded parts i) and ii).[25] If the fraction of these excluded parts in a total survey income is correlated with the income level, this sample selection can cause selection bias. Dropping observations that have too much value in expense items with a longer reference period than the previous three months in the sample selection can cause the same issue.

To resolve this concern, when constructing the residual income distribution and determining the quantiles of selected observations in the Peruvian sample, I include the observations dropped due to having too much value in income or expense items with a longer reference period than the previous three months or in imputed incomes. To sort these dropped observations and the selected observations together, I use the residual income of a comprehensive income measure that includes not only the baseline income but also the income items with a longer reference period than the previous three months and the imputed incomes, which are excluded from the baseline measure. Although these excluded parts are bad because they can blur the identification of income shocks, they are useful in determining the income quantiles of the selected observations.

[22] In other words, precautionary saving theory predicts that the monotone MPC heterogeneity over residual income $y_{i,t}$ should be diluted by the predictable component $Z'_{i,t}\varphi^y_t$ when groups are instead formed by actual income ($Y_{i,t}$). In Online Appendix H.2, I examine this prediction by grouping households based on actual income ($Y_{i,t}$).

[23] Notably, many incomplete-market models, including the most canonical models (e.g., Aiyagari (1994), Huggett (1996), and Carroll (1997)), abstract permanent income heterogeneity that the predictable component $Z'_{i,t}\varphi^y_t$ captures in my estimation, such as income heterogeneity due to education, ethnicity, and region. In such models, income (after subtracting age-specific components in the case of life-cycle models) simply corresponds to residual income $y_{i,t}$ in the data, and the income process is often calibrated using residual incomes after controlling for the observable characteristics that they abstract. See Floden and Linde (2001), Heathcote, Perri, and Violante (2010), and Guvenen and Smith (2014), for example.

[24] Because I already remove time fixed effects when controlling for the predictable components (annually for the U.S. sample, quarterly for the Peruvian sample), it should also be fine to sort residual incomes $y_{i,t}$ in a larger observation pool than the pool of the reference period. In subsection 7.3.2, I conduct a robustness check by sorting residual incomes in different observation pools.

[25] See footnote 19 for the exact definition of 'too much value.'

4 Peruvian MPCs

4.1 Quarterly MPC Estimates

I estimate the MPCs of Peruvian households by applying the GMM method with moment conditions (5) to the Peruvian sample. Since the sample provides year-over-year growth of quarterly income and consumption, I set the period t as a quarter and $K = 4$. As a result, I obtain Peruvian quarterly MPC estimates.

The blue line with circle markers (labeled 'data') in Figure 1 represents the MPC estimate within each residual income decile.[26] Two observations are noteworthy. First, the mean quarterly MPC across deciles is 20.4%. Since each decile has the same population share, this mean MPC is a population-weighted average. Thus, we can interpret this number as follows: on average, Peruvian households spend 20.4% out of an unexpected and transitory income increase within a quarter.[27]

Second, MPC estimates are heterogeneous over residual income deciles. The MPC estimates range from 10.6% in the top decile to 28.0% in the bottom decile. The MPC estimates tend to be higher in lower income deciles, although the relationship is not monotonic. In pairwise comparison, some pairs of deciles exhibit statistically significant differences in their MPCs. The top decile is significantly different from the first, fifth, sixth, and seventh deciles at the 95% confidence level and from the fourth and ninth deciles at the 90% confidence level. In these pairs exhibiting sig-

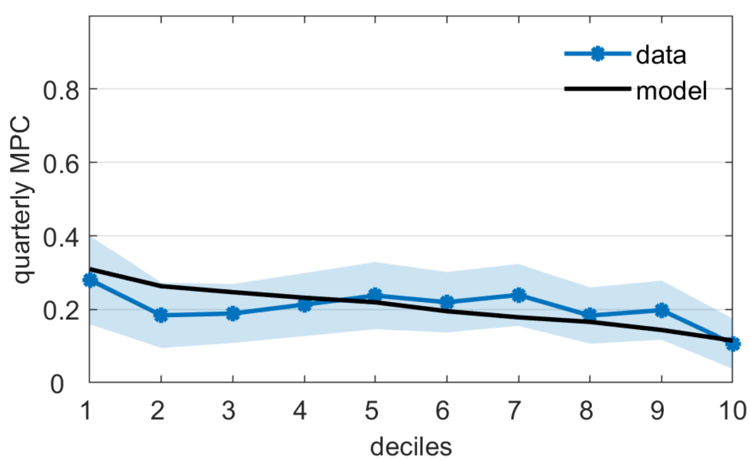

Figure 1: Peruvian quarterly MPCs: data vs model

Notes: Integers on the x-axis represent the deciles of residual income after subtracting a predictable component. The integer 1 denotes the bottom decile. Shaded areas represent 95% confidence intervals.

[26] Online Appendix C reports the estimates and standard errors in a table for interested readers.

[27] Another candidate for a headline number as a national average MPC is an MPC estimate in the ungrouped sample. The Peruvian quarterly MPC estimate in the ungrouped sample is 19.5%. The difference between the mean MPC across residual income deciles and the MPC estimate in the ungrouped sample is that the former allows MPC to be heterogeneous over the deciles in the estimation, while the latter does not. Given the MPC heterogeneity result discussed in the following paragraph, I use the former as my preferred headline number as a national average MPC.

nificantly different MPCs, the lower income deciles always exhibit higher MPCs than the higher income deciles.

4.2 A Standard Model's Fit

How well can a standard model match the MPC estimates? To answer this question, I conduct numerical simulation using a standard incomplete-market, life-cycle model introduced in subsection 2.2.

Table 1 summarizes the calibration of the model fitting to the Peruvian economy. The time unit is set equal to a quarter. Parameters governing the stochastic process for residual income $y_{i,t}$ (ρ, σ_{ps}, σ_{tr}, and σ_{P_0}) are estimated using the GMM method, similarly as in Floden and Linde (2001) and Storesletten, Telmer, and Yaron (2004). For the moment conditions, I use yearly-age-specific variances and covariances of $y_{i,t}$. Two observations are noteworthy. First, ρ is substantially lower than 1. (In terms of an annual rate, $\rho^4 = 0.861$.) This result is driven by the Peruvian data pattern that yearly-age-specific covariances $cov[y_{i,t}, y_{i,t+4}|4a \leq t \leq 4a+3]$'s (in which a is a yearly age) are significantly greater than $cov[y_{i,t}, y_{i,t+4k}|4a \leq t \leq 4a+3]$'s, $k \geq 2$. Note that $cov[y_{i,t}, y_{i,t+4k}|t] = \rho^{4k} var[P_{i,t}|t]$ in the model. Second, $\sigma^2_{P_0}$ (0.296) is very close to $\frac{\sigma^2_{ps}}{1-\rho^2}$ (0.293). This result reflects the Peruvian data pattern that yearly-age-specific variances $var[y_{i,t}|4a \leq t \leq 4a+3]$'s are flat over ages. Note that $var[y_{i,t}|t] = \sum_{s=0}^{t-1} \rho^{2s}\sigma^2_{ps} + \rho^{2t}\sigma^2_{P_0} + \sigma^2_{tr}$ for $1 \leq t \leq T$. Online Appendix E.2.1 provides more details of the estimation procedure and results.

The age-specific deterministic income component ω_t is computed as follows. First, I compute the yearly-age-specific means of the predictable components of income in the data and normalize them by subtracting the unconditional mean. Then, I fit a sixth-order polynomial curve to these normalized yearly-age-specific means. Lastly, I use this fitted curve to interpolate the quarterly-

Table 1: Calibration for the Peruvian economy

	Description	Value	Target / Source
labor income process			
ρ	persistence of the AR(1) component	0.963	ENAHO
σ_{ps}	S.D. of shocks to the AR(1) component	0.146	ENAHO
σ_{tr}	S.D. of shocks to the *i.i.d.* component	0.443	ENAHO
σ_{P_0}	S.D. of initial draw P_0	0.544	ENAHO
ω_t	age-specific deterministic component		ENAHO
other parameters			
T	maximum quarterly age	163	sample age 25-65
$\bar{B}$	borrowing limit	0	ZBL
σ	inverse of IES	2	Kaplan and Violante (2010)
r	quarterized real interest rate	0.003	EMBIG, T-bill
β	discount factor	0.943	MPC estimates

Notes: The time unit is a quarter.

age-specific component ω_t. Online Appendix E.1 provides a figure that plots the data points of normalized yearly-age-specific means and the fitted curve.

The maximum quarterly age T is set equal to 163, consistent with the household ages in the sample, 25-65. Specifically, $t = 0$ corresponds to the first quarter of age 25, and $t = T$ corresponds to the fourth quarter of age 65. In the benchmark economy, I consider a zero borrowing limit (ZBL) by setting $\bar{B} = 0$. In section 6, I consider an alternative economy where the borrowing limit is removed. The parameter σ governs the intertemporal elasticity of substitution, risk aversion, and prudence (in Kimball (1990)'s sense). To be standard, I choose $\sigma = 2$, following Kaplan and Violante (2010).[28]

To calibrate the real interest rate r, we need to consider the characteristics of asset $A_{i,t}$. In the model, the assets are risk-free and perfectly liquid at a quarterly frequency. As noted by Carroll (1997), a government bond is a good proxy for such assets. Based on this consideration, I calibrate r using Peruvian sovereign bond rates. First, I recover the Peruvian sovereign bond rates by adding J.P. Morgan's EMBIG spread with U.S. T-bill rates, following Neumeyer and Perri (2005) and Uribe and Yue (2006).[29] Then, I take an average of the recovered sovereign rates over the sample period, 2004-2016. As a result, I obtain $r = 0.003$. (In terms of an annual rate, $(1+r)^4 - 1 = 0.0125$.)

Lastly, I calibrate β by targeting MPC estimates. Specifically, I find β that minimizes $(MPC^{Q}_{data} - MPC^{Q}_{model})' \cdot \Omega^{Q} \cdot (MPC^{Q}_{data} - MPC^{Q}_{model})$, where MPC^{Q}_{data} is a 10-by-1 vector of the estimated Peruvian quarterly MPC at each residual income decile, MPC^{Q}_{model} is its model counterpart, and Ω^{Q} is a weight matrix. To compute MPC^{Q}_{model}, I simulate quarterly income and consumption of 1,000,000 households over nine quarters and apply to these simulated data the same MPC estimation procedure applied to the Peruvian data.[30] For the weight matrix Ω^{Q}, I use a diagonal matrix whose diagonal elements are $diag((V^{Q})^{-1})$, where V^{Q} is the variance-covariance matrix of MPC^{Q}_{data} to avoid small-sample bias that optimal weight matrix $(V^{Q})^{-1}$ causes, as recommended by Altonji and Segal (1996).[31] As a result, I obtain $\beta = 0.943$.[32]

[28] $\sigma = 2$ is also commonly used in emerging market business cycle studies. See Garcia-Cicco, Pancrazi, and Uribe (2010), for example.

[29] J.P. Morgan's EMBIG spread is obtained from Banco Central de Reserva del Perú. Real U.S. T-bill rates are computed by deflating nominal T-bill rates with expected inflations on U.S. CPIs. Nominal U.S. T-bill rates are obtained from the Federal Reserve Bank of St. Louis. The expected inflations are approximated by an average inflation rate over the current and past three quarters, as in Neumeyer and Perri (2005) and Uribe and Yue (2006). Atkeson and Ohanian (2001) provide empirical support for this approximation. U.S. CPI series are also obtained from the Federal Reserve Bank of St. Louis.

[30] Specifically, I group households based on their residual income in quarter 0, compute each group's ψ_G by equation (3) using differences of residual income and consumption between quarters 0 and 4 and between quarters 4 and 8, and convert it to the MPC estimate by equation (4) using (unresidualized) income and consumption in quarter 4. As in the sample selection, simulated households becoming older than 65 at or before quarter 8 are dropped.

[31] To obtain V^{Q}, I conduct joint GMM estimation for the ten deciles. Note that separate GMM estimation for each decile using three moment conditions in equation (5) and joint GMM estimation for all the deciles using thirty moment conditions yield the same result (except that the joint estimation additionally provides the variance-covariance matrix).

[32] Under this calibration, the ratio of aggregate assets to aggregate quarterly labor income is 1.467 in the model. Assuming that the labor income share in Peru is 0.64 (Guerriero, 2019), a back-of-the-envelope calculation suggests that the ratio of aggregate assets to annual GDP is 0.235 ($= 1.467 \times 0.64 \div 4$). The ratio is very small compared to 2.667, the average ratio of aggregate physical capital to annual GDP in Peru during 2004-2016 computed from Feenstra, Inklaar, and Timmer (2015)'s Penn World Table (version 9.1). As noted above, the assets in this model should be best understood as risk-free and liquid wealth. In this sense, I share Carroll (1997)'s view that this model under the calibration of low

The black solid line (labeled 'model') in Figure 1 represents the model-simulated MPC at each residual income decile. The figure shows that this standard incomplete-market, life-cycle model can fit the MPC estimates quite well. The mean of the model-simulated MPCs across deciles is 20.6%. The model-simulated MPCs vary from 11.5% in the top decile to 30.9% in the bottom decile. The model also predicts that lower income deciles exhibit higher MPCs than higher income deciles.

4.3 Performance of the MPC estimator in a Standard Model

Given that we have a model that closely matches the MPC estimates, I use it as a laboratory to check how closely the MPC estimator recovers true MPCs, as in Kaplan and Violante (2010). Kaplan and Violante (2010) show that Blundell et al. (2008)'s method can recover the partial insurance parameter to transitory shocks quite well in a model fitted to the U.S. economy. I conduct a similar exercise but with three important changes. First, my MPC estimator extends Blundell et al. (2008)'s original method by replacing $(\Delta y_{i,t}, \Delta c_{i,t})$ with $(\Delta^K y_{i,t}, \Delta^K c_{i,t})$. Second, my estimator identifies MPC, not the partial insurance parameter, and thus can be compared with the true MPC defined and computed according to its original meaning (the ratio of consumption change to income change generated by a shock) in the model. Third, the model is now fitted to the Peruvian environment in which, as will be shown later, households exhibit substantially higher MPCs than U.S. households.

The original meaning of quarterly MPC to a transitory shock is the ratio of consumption change to income change caused by the shock within a quarter after its realization. In the data, we recover this object by multiplying Blundell et al. (2008)'s partial insurance parameter with a consumption-to-income ratio. In the model, on the other hand, we can directly compute this object as follows. Let $C_t(A_{i,t-1}, P_{i,t}, \epsilon_{i,t})$ be household i's policy function for (unresidualized) consumption and $Y_t(P_{i,t}, \epsilon_{i,t})$ be its (unresidualized) income at age t. Quarterly MPC to a transitory shock for each household i can be computed as follows.

$$MPC^{Q}_{true}(i,t) = \frac{C_t(A_{i,t-1}, P_{i,t}, \epsilon_{i,t}) - C_t(A_{i,t-1}, P_{i,t}, 0)}{Y_t(P_{i,t}, \epsilon_{i,t}) - Y_t(P_{i,t}, 0)}.$$

I name this model statistic 'true quarterly MPC.' In Figure 2a, I plot the group mean of the true quarterly MPCs within each residual income decile (labeled 'true'). Then, I compare them with the model counterparts of the MPC estimates (labeled 'BPP'), which are computed by applying to model-simulated data the same estimation procedure applied to the Peruvian data.[33] The figure

$\beta(1+r)$ provides a good explanation for the behavior of typical consumers, while it probably does not for the origin of the aggregate capital. Kaplan et al. (2018) show that by introducing a two-asset (liquid and illiquid) structure, one can achieve both the correct degree of consumption sensitivity to income shocks and the correct amount of aggregate capital in a model because households can be liquidity-poor while holding a large amount of illiquid assets.

[33] As discussed in footnote 30, when computing the model counterparts of the MPC estimates, simulated households are grouped by their residual incomes in quarter 0, and consumption responses to transitory shocks in quarter 4 are identified. To compare with them, I compute the group means of true quarterly MPCs by grouping simulated households by their residual incomes in quarter 0 and counterfactually turning off their transitory shocks in quarter 4. As in

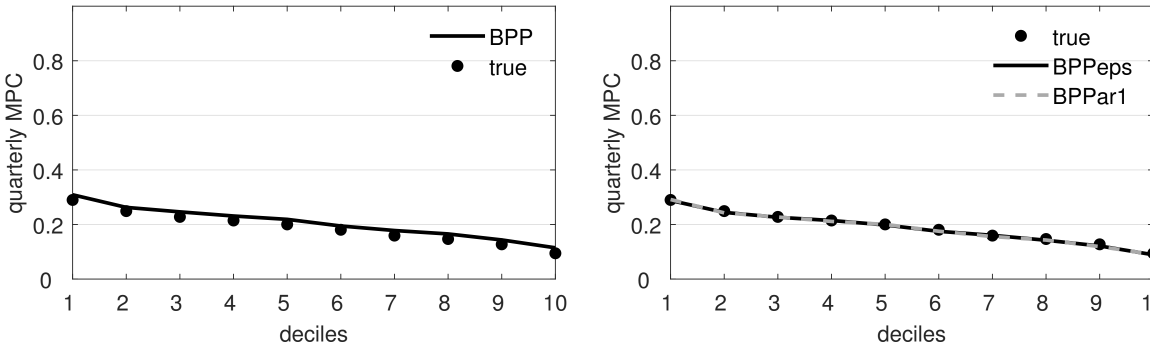

(a) BPP vs TRUE

(b) BPP without misspecification vs TRUE

Figure 2: Performance of the MPC estimator in the model

Notes: In Figure 2a, 'BPP' represents model statistics computed by applying to simulated data the MPC estimator ($\psi_G \frac{E[C_{i,t}|(i,t)\in G]}{E[Y_{i,t}|(i,t)\in G]}$) applied to the actual Peruvian data. In Figure 2b, 'BPPeps' and 'BPPar1' represent model statistics computed in the same way as 'BPP' of Figure 2a except that ψ_G is computed differently such that the effect of misspecification on ρ ($\rho < 1$ in the model, $\rho = 1$ in the estimation) is removed. In both figures, 'true' represents the true MPCs defined and computed according to the original meaning of MPC (the ratio of consumption change to income change generated by a shock) in the model.

shows that the MPC estimator recovers true quarterly MPCs quite well.

Moreover, it turns out that even the slight degree of upward bias that the MPC estimator exhibits in Figure 2a is not a feature of the estimator but rather caused by the income process misspecification that I intentionally allow: the MPC estimation assumes $\rho = 1$, while the model has $\rho < 1$. This misspecification generates a slight bias on ψ_G (identified by $\frac{cov[\Delta^K c_{it}, \Delta^K y_{i,t+K}|(i,t)\in G]}{cov[\Delta^K y_{it}, \Delta^K y_{i,t+K}|(i,t)\in G]}$) and thus on the MPC estimates ($\psi_G \frac{E[C_{i,t}|(i,t)\in G]}{E[Y_{i,t}|(i,t)\in G]}$). When this misspecification effect is removed by computing ψ_G using realized shocks in the model by $\psi_G = \frac{cov[\Delta^K c_{i,t}, \epsilon_{i,t}|(i,t)\in G]}{var[\epsilon_{i,t}|(i,t)\in G]}$, the MPC estimates (labeled 'BPPeps' in Figure 2b) hit the true MPCs precisely.

In the actual data, we do not observe realized shocks $\epsilon_{i,t}$. As Kaplan and Violante (2010) suggest, however, we can revise the identification equation of ψ_G reflecting $\rho < 1$ when the value of ρ is known. Specifically, ψ_G can be identified by $\frac{cov[\Delta^K c_{it}, \tilde{\Delta}^K y_{i,t+K}|(i,t)\in G]}{cov[\tilde{\Delta}^K y_{it}, \tilde{\Delta}^K y_{i,t+K}|(i,t)\in G]}$, where $\tilde{\Delta}^K y_{i,t} = y_{i,t} - \rho^K y_{i,t-K}$. When ψ_G is identified with this revised estimator, the MPC estimates (labeled 'BPPar1' in Figure 2b) again hit the true MPCs precisely.

In short, the MPC estimator recovers true quarterly MPCs quite well even in the presence of the income process misspecification that the MPC estimation assumes $\rho = 1$, while the model has $\rho < 1$. The income misspecification induces a very small degree of upward bias on the MPC estimates, and even this small bias can be corrected by revising the estimator such that $\Delta^K y_{i,t}$ is replaced with $\tilde{\Delta}^K y_{i,t}$ when identifying ψ_G.[34]

footnote 30, simulated households becoming older than 65 at or before quarter 8 are dropped.

[34]Motivated by Figure 2b, I re-estimate MPCs in subsection 7.3.3 using the revised estimator with the value of ρ reported in Table 1. The MPC estimates do not change much.

4.4 Annualization

One of the purposes of this paper is to provide off-the-shelf MPC estimates for researchers. To serve this purpose better, I annualize the quarterly MPC estimates, as annual MPCs can be useful in many circumstances – such as when calibrating a yearly model or when comparing with other annual MPC estimates.[35]

I consider two annualization methods. First, under the assumption that dynamic consumption responses to a transitory shock die out exponentially over time and that the interest rate is zero, Auclert (2019) derives the following approximation.

$$MPC_G^A \approx 1 - (1 - MPC_G^Q)^4 \tag{6}$$

in which MPC_G^A and MPC_G^Q are the annual MPC and quarterly MPC in group G, respectively. I name this approximation 'model-free annualization,' as it does not require a model to be solved. Second, since we already have a model that fits the quarterly MPC estimates well, we can directly compute annual MPC in the model according to its original meaning (the ratio of consumption change to income change caused by a transitory shock within a year after its realization). I name this model statistic 'true annual MPC.' Essentially, I annualize quarterly MPC estimates by computing true annual MPCs in a model that fits the quarterly estimates well. I name this annualization 'model-based annualization.' Online Appendix D provides mathematical details for these annualization methods.[36]

The blue line with circle markers (labeled 'model-free') in Figure 3 plots the annualized MPC at each residual income decile under the model-free annualization.[37] The mean of the annualized MPCs across deciles is 59.2%.[38] The annualized MPCs vary from 36.1% in the top decile to 73.1% in the bottom decile. Inherited from the quarterly estimates, the annualized MPCs also tend to be higher in lower income deciles, although the relationship is not monotonic.

The black solid line (labeled 'model-base') in Figure 3 represents the true annual MPCs in the model (or, equivalently, annualized MPCs under the model-based annualization).[39] The mean of the true annual MPCs across residual income deciles is 54.5%. The true annual MPCs vary from 33.3% in the top decile to 72.6% in the bottom decile. As it does for quarterly MPCs, the model

[35] Many studies estimate MPCs with an annual horizon. See Carroll, Slacalek, Tokuoka, and White (2017) or Crawley and Kuchler (2021) for a nice summary of the existing MPC estimates.

[36] As a way of annualizing MPC, one might consider annualizing ahead of time the quarterly income and consumption data by multiplying both by four and then applying the MPC estimation method to the annualized data. This ahead-of-time annualization yields an annual MPC that is exactly equal to the quarterly MPC and thus essentially ignores dynamic consumption responses to a transitory shock in the subsequent quarters. See Online Appendix D for details.

[37] Online Appendix C reports the estimates and standard errors in a table for interested readers. Standard errors are also converted using equation (6) and the delta method.

[38] The annualized ungrouped MPC estimate is 58.0%.

[39] Annual MPCs in the model are computed using simulated households as follows. First, as in footnotes 30 and 33, I group simulated households by their residual incomes in quarter 0 and drop those becoming older than 65 at or before quarter 8. Then, I counterfactually turn off transitory income shocks in quarter 4 and track how consumption changes in quarters 4, 5, 6, and 7.

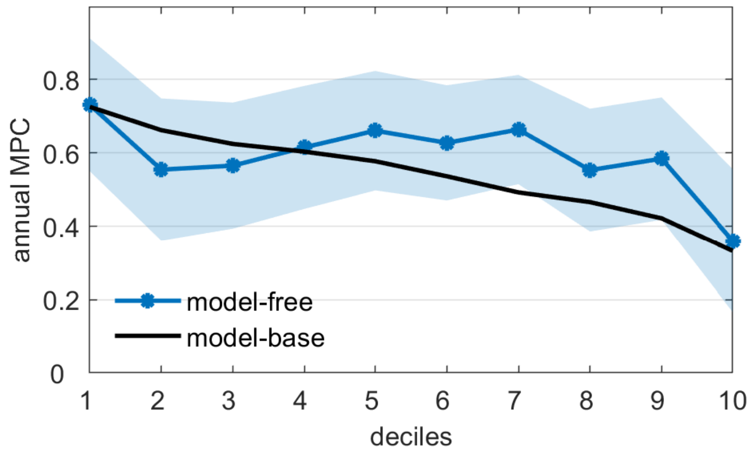

Figure 3: Annualization of Peruvian quarterly MPCs

Notes: In this figure, 'model-free' represents annualized MPC estimates based on Auclert (2019)'s model-free method (6), and 'model-base' represents true annual MPCs in the model. Shaded areas represent 95% confidence intervals.

also predicts that annual MPCs are higher in lower income deciles than in higher income deciles.

Importantly, a comparison between the two annualization methods in Figure 3 suggests that the model-free annualization recovers true annual MPCs quite well.

5 Comparison with U.S. MPCs

5.1 Model-Free Comparison

It is of interest to compare an emerging economy's MPCs with a developed economy's MPCs. To this end, I compare Peruvian MPCs with U.S. MPCs obtained by applying the same estimation method.

As in the estimation of Peruvian MPCs, I estimate MPCs of U.S. households by applying the GMM method with moment conditions (5) to the U.S. sample. Unlike the Peruvian sample, however, the U.S. sample provides two-year-over-two-year growth of annual income and consumption. Therefore, I set the period t as a year and $K = 2$ in the moment conditions (5). As a result, I obtain U.S. annual MPC estimates.

The red dashed line with square markers (labeled 'US') in Figure 4 plots the U.S. annual MPC estimate within each residual income decile.[40] The mean of the MPC estimates across deciles is 8.3%.[41] The annual MPC estimates vary from 2.4% in the ninth decile and 4.0% in the top decile

[40] Online Appendix C reports the estimates and standard errors in a table for interested readers.

[41] The U.S. annual MPC estimate in the ungrouped sample is 7.8%. My estimation result is in the same ballpark as that of existing studies that estimate U.S. MPC or the partial insurance parameter to a transitory shock ψ by applying Blundell et al. (2008)'s method to the PSID data in similar sample periods. Auclert (2019) estimates U.S. MPCs for terciles of gross income (including labor income, capital income, and capital gains) using the 1999-2013 waves. My ungrouped MPC estimate (7.8%) is between his top tercile estimate (1.6%) and bottom tercile estimate (13.2%) and close to his middle tercile estimate (10.2%). Kaplan et al. (2014b) estimate ψ_G for poor HtM (hand-to-mouth) households, wealthy HtM households, and non-HtM households using the 1999-2011 waves. My ungrouped estimate of ψ (18.7%)

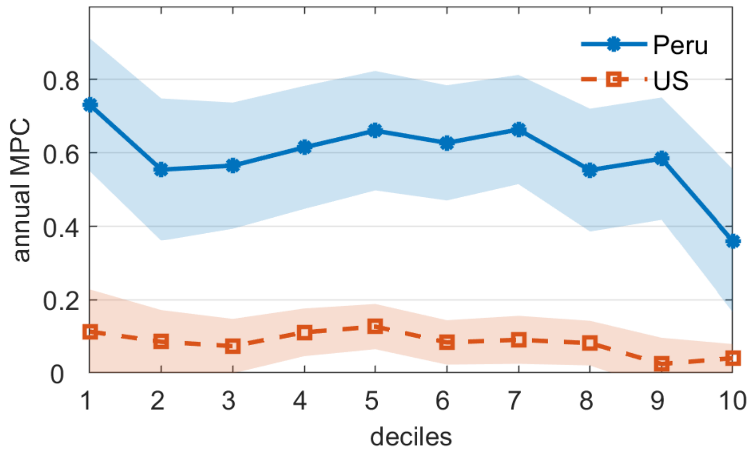

Figure 4: Model-free comparison between Peruvian and U.S. MPCs

Notes: In this figure, 'Peru' represents Peruvian MPC estimates annualized by a model-free method using Auclert (2019)'s conversion formula (6), and 'US' represents annual MPC estimates obtained by applying Blundell et al. (2008)'s method to annual U.S. data. Shaded areas represent 95% confidence intervals.

to 12.6% in the fifth decile and 11.3% in the bottom decile. The U.S. MPC estimates also exhibit a tendency to be higher in lower income deciles, although the relationship is not monotonic. In pairwise comparison, some pairs of deciles exhibit statistically significant differences in their MPCs. The fifth decile is significantly different from the top and ninth deciles at the 95% confidence level, and the fourth decile is significantly different from the top and ninth deciles at the 90% confidence level. In these pairs exhibiting significantly different MPCs, the lower income deciles always exhibit higher MPCs.

In Figure 4, I compare the U.S. annual MPC estimates with the annualized Peruvian MPC estimates by model-free annualization (labeled 'Peru') obtained in subsection 4.4. I name this comparison 'model-free comparison,' as it does not require a model to be solved. The model-free comparison suggests that Peruvian households exhibit substantially higher MPCs overall than U.S. households. The mean annual MPC in Peru (59.2%) is 7.2 times as large as that in the U.S. (8.3%). Moreover, this comparison also suggests that Peruvian MPCs are substantially more heterogeneous over residual income deciles than are U.S. MPCs.

5.2 Time Aggregation Problem

The model-free comparison in subsection 5.1 has a critical problem in that it imposes an asymmetric assumption regarding the time frame. When estimating MPCs, Peruvian households are

is between their non-HtM estimate (12.7%) and their HtM estimates (24.3% for poor HtM, 30.1% for wealthy HtM). However, my ungrouped estimate of ψ (18.7%) is noticeably higher than Blundell et al. (2008)'s original ungrouped estimate (5.3%). They use the 1978-1992 waves of the PSID and impute consumption using the Consumer Expenditure Survey (CEX) due to the PSID's narrow coverage on expense items during their sample period. The consumption imputation might be one possible reason for their estimate being lower than the estimates from more recent samples since the consumption imputation can blur the covariance between $\Delta c_{i,t}$ and $\Delta y_{i,t}$.

assumed to receive income shocks and make consumption decisions quarterly, while U.S. households are assumed to do so yearly. If we instead impose a symmetric assumption that both Peruvian and U.S. households receive income shocks and make consumption decisions quarterly, my U.S. MPC estimates are subject to a 'time aggregation problem' that Crawley (2020) points out: when Blundell et al. (2008)'s method is applied to annual data while households receive income shocks and make consumption decisions at a higher frequency in the underlying model, the method significantly underestimates the partial insurance parameter to transitory shocks. Part of the large gap between Peruvian and U.S. MPC estimates observed under the model-free comparison in Figure 4 can be attributable to the time aggregation problem.

In this subsection, I gauge the time-aggregation-induced bias on my U.S. MPC estimates in the incomplete-market, life-cycle model introduced in subsection 2.2. To this end, the model is calibrated quarterly and fitted to the U.S. economy.

Table 2 summarizes the calibration fitting to the U.S. economy. The time unit is set equal to a quarter. Given that the time unit in the model (quarter) is different from the reference period in the data (year), parameters governing the stochastic process for residual income $y_{i,t}$ (ρ, σ_{ps}, σ_{tr}, and σ_{P_0}) are estimated using the SMM (Simulated Method of Moments) method as follows. First, for a given set of parameters, I simulate quarterly income series and convert them into annual series by aggregating them over every four quarters.[42] After residualizing the simulated annual incomes, I compute their age-specific variances and covariances. I find parameters that minimize the distance between these simulated moments and their data counterparts. Two observations are noteworthy. First, unlike in Peru, ρ is very close to 1 in the U.S. (In terms of an annual rate,

Table 2: Calibration for the U.S. economy

	Description	Value	Target / Source
labor income process			
ρ	persistence of the AR(1) component	0.989	PSID
σ_{ps}	S.D. of shocks to the AR(1) component	0.072	PSID
σ_{tr}	S.D. of shocks to the *i.i.d*. component	0.513	PSID
σ_{P_0}	S.D. of initial draw P_0	0.360	PSID
ω_t	age-specific deterministic component		PSID
other parameters			
T	maximum quarterly age	163	sample age 25-65
$\bar{B}$	borrowing limit	0	ZBL
σ	inverse of IES	2	Kaplan and Violante (2010)
r	quarterized real interest rate	0.001	T-bill
β	discount factor	0.998	MPC estimates

Notes: The time unit is a quarter.

[42]Specifically, I simulate quarterly residual incomes $y_{i,t}$, convert them into quarterly actual incomes $Y_{i,t}$ using age-specific deterministic components ω_t, and then convert them into annual actual incomes by aggregating them over every four quarters.

$\rho^4 = 0.958$.) It reflects the U.S. data pattern that for residual annual income $\hat{y}_{i,a}$ at yearly age a, $cov[\hat{y}_{i,a}, \hat{y}_{i,a+2}|a]$'s are not substantially greater than $cov[\hat{y}_{i,a}, \hat{y}_{i,a+2k}|a]$'s , $k \geq 2$. Second, $\sigma^2_{P_0}$ (0.130) is noticeably smaller than $\frac{\sigma^2_{ps}}{1-\rho^2}$ (0.241). This result reflects the U.S. data pattern that $var[\hat{y}_{i,a}|a]$ increases with age. Online Appendix E.2.2 provides more details of the estimation procedure and results.

The age-specific deterministic component ω_t is again computed by fitting a sixth-order polynomial curve to the normalized yearly-age-specific means of the predictable components and using this curve to interpolate ω_t. Online Appendix E.1 provides a figure that plots the data points of normalized yearly-age-specific means and the fitted curve.

As done in the benchmark Peruvian economy, I set $T = 163$, $\bar{B} = 0$, and $\sigma = 2$. The real interest rate r is calibrated by matching the average real rate on U.S. T-bills during 1998-2010 (which corresponds to the sample period for the 1999-2011 waves of the PSID).[43] As a result, I obtain $r = 0.001$. (In terms of an annual rate, $(1+r)^4 - 1 = 0.004$.)

Lastly, I calibrate β by targeting the U.S. annual MPC estimates. Specifically, I find β that minimizes $(MPC^A_{data} - MPC^A_{model})' \cdot \Omega^A \cdot (MPC^A_{data} - MPC^A_{model})$, where MPC^A_{data} is a 10-by-1 vector of the estimated U.S. annual MPC at each residual income decile and MPC^A_{model} is its model counterpart. To compute MPC^A_{model}, I simulate the quarterly income and consumption of 1,000,000 households over twenty quarters. Then, I convert these quarterly series into annual series (over five years) by aggregating them over every four quarters. I obtain MPC^A_{model} by applying to these simulated annual series the same estimation procedure applied to the PSID.[44] For the weight matrix Ω^A, I again use a diagonal matrix whose diagonal elements are equal to $diag((V^A)^{-1})$, where V^A is the variance-covariance matrix of MPC^A_{data} to avoid the small-sample bias that optimal weight matrix $(V^A)^{-1}$ causes, as suggested by Altonji and Segal (1996). As a result, I obtain $\beta = 0.998$.[45]

In Figure 5, the red dashed line with square markers (labeled 'data') represents the annual U.S. MPC estimate within each residual income decile, which is obtained under the assumption that residual annual income follows an annual income process, and the black dashed line (labeled 'counterpart, model') represents its model counterpart, which is obtained by applying the MPC estimation method to annual income and consumption series simulated from the quarterly model fitted to the U.S. economy. These two graphs track each other closely, showing that this quarterly

[43]See footnote 29 for the construction of the real U.S. T-bill rates.

[44]Specifically, I group households based on their residual annual income in year 0, compute each group's ψ_G by equation (3) using differences of residual annual income and consumption between years 0 and 2 and between years 2 and 4, and convert it to the annual MPC estimate by equation (4) using (unresidualized) income and consumption in year 2. As in the sample selection, simulated households becoming older than 65 at or before the end of year 5 (or, equivalently, quarter 19) are dropped.

[45]Under this calibration, the ratio of aggregate assets to aggregate quarterly labor income is 6.387 in the model. Assuming that the labor income share in the U.S. is 0.74 (Guerriero, 2019), a back-of-the-envelope calculation suggests that the ratio of aggregate assets to annual GDP is 1.182 ($= 6.387 \times 0.74 \div 4$). This ratio is substantially smaller than 3.277, the average ratio of aggregate physical capital to annual GDP in the U.S. during 1998-2010 computed from Feenstra et al. (2015)'s Penn World Table (version 9.1). However, the ratio (1.182) is substantially greater than the same ratio in the Peruvian model economy, 0.235. As noted above, the assets in this model should be best understood as risk-free and liquid assets. According to this interpretation, U.S. households hold a substantially greater amount of liquid assets than Peruvian households.

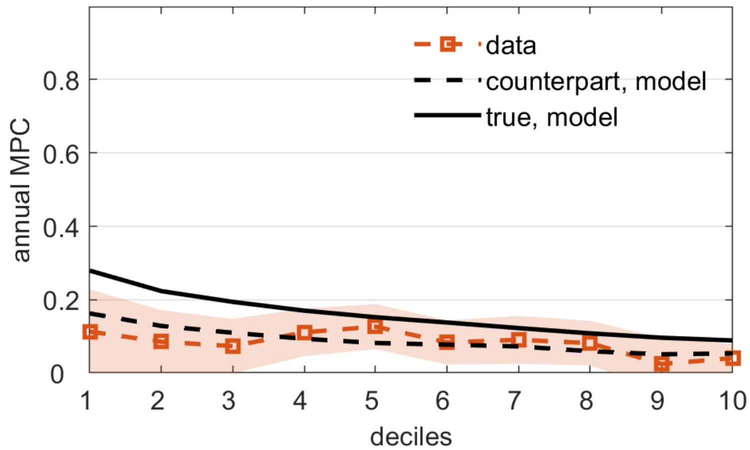

Figure 5: Time aggregation problem in the annual U.S. MPCs

Notes: In this figure, 'data' represents the annual U.S. MPC estimates, 'counterpart, model' represents the model counterparts of the estimates, which are obtained by applying the same MPC estimation method to annual income and consumption series simulated from the quarterly model, and 'true, model' represents true annual MPCs in the model.

model can fit the annual MPC estimates quite well under the same estimation procedure.

The black solid line in Figure 5 (labeled 'true, model') represents true annual MPCs in this model.[46] A comparison between the 'counterpart, model' and the 'true, model' in this figure verifies that a time aggregation problem arises when the annual U.S. MPC estimator is applied to annual income and consumption series simulated from the quarterly model. Specifically, the mean of the true annual MPCs across residual income deciles is 15.7 %, which is 1.8 times as large as the mean of the model counterparts of the annual MPC estimates, 8.9%.

5.3 Model-Based Comparison

Given that we have a Peruvian model economy that fits the quarterly Peruvian MPC estimates well and a U.S. model economy that fits the annual U.S. MPC estimates well, we can compare true MPCs between these two economies. I name this comparison 'model-based comparison.' Note that this comparison is free from the time aggregation problem.

Figure 6 compares the true annual MPC at each residual income decile between the Peruvian and U.S. model economies.[47] This figure shows that the two main patterns observed under the

[46]The true annual MPCs plotted by this graph ('true, model' in Figure 5) are computed as follows. First, I group simulated households in the same way that I group them when computing annual U.S. MPC estimates and their model counterparts. Specifically, I group simulated households based on their residual annual income in year 0 and drop those becoming older than 65 at or before quarter 19, as in footnote 44. Then, I counterfactually turn off transitory income shocks in quarter 8 and track how household consumption changes in quarters 8, 9, 10, and 11.

[47]The true annual MPCs in Figure 6 are computed as in footnote 39. Specifically, I group simulated households by their residual income in quarter 0 and drop those becoming older than 65 at or before quarter 8. Then, I counterfactually turn off transitory shocks in quarter 4 and track how consumption changes in quarters 4, 5, 6, and 7. Note that the graph labeled 'US' in Figure 6 is not identical to the graph labeled 'true, model' in Figure 5, although both plot true annual MPCs. This is because they use different groupings to be comparable with their respective comparison counterparts. The former graph uses quarterly residual income deciles in quarter 0 when MPCs to transitory shocks realized in

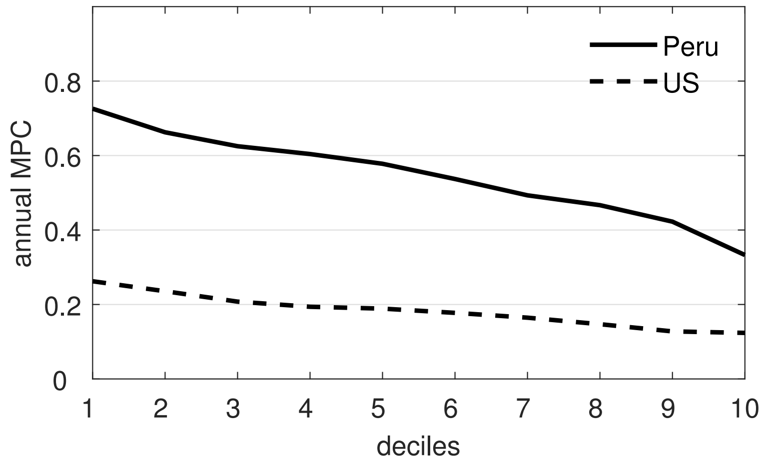

Figure 6: Model-based comparison between Peruvian and U.S. annual MPCs

Notes: 'Peru' and 'US' represent the true annual MPC at each residual income decile in the Peruvian and U.S. model economies, respectively.

model-free comparison in subsection 5.1 remain robust under this model-based comparison. First, Peruvian annual MPCs are substantially greater overall than U.S. annual MPCs. The mean annual MPC is 54.5% in Peru, while it is 18.3% in the U.S. Note that this mean MPC gap is noticeably smaller than the gap in the model-free comparison because the time aggregation problem is fixed. Nevertheless, the model-based comparison predicts that the mean annual MPC in Peru is 3.0 times as large as that in the U.S. Second, as in the model-free comparison, the model-based comparison also predicts that annual MPCs are more heterogeneous over residual income deciles in Peru than in the U.S. Annual MPCs vary from 33.3% in the top decile to 72.6% in the bottom decile of the Peruvian model economy, while they vary from 12.4% in the top decile to 26.2% in the bottom decile of the U.S. model economy. In both economies, annual MPCs are higher in lower income deciles than in higher income deciles.

6 The Role of Precautionary Saving

What drives the mean MPC gap between Peru and the U.S.? Through the lens of the model, this gap can be decomposed into two components.

First, in this standard incomplete-market model, households face idiosyncratic income risk. This income risk induces precautionary saving behavior because i) households fear the realization of a low-income path (Kimball, 1990) and ii) they also worry about being hit by borrowing limits (Huggett, 1993). These precautionary saving motives become weaker as households' cash-on-hand (or, equivalently, total currently available resources $Y_{i,t} + (1+r)A_{i,t-1}$) increases. A positive

quarter 4 are computed. The latter graph, on the other hand, uses annual residual income deciles in year 0 (composed of quarters 0, 1, 2, and 3) when MPCs to transitory shocks realized in quarter 8 are computed. Moreover, I drop simulated households becoming older than 65 at or before quarter 8 when computing the former graph, while I drop those becoming older than 65 at or before quarter 19 when computing the latter graph.

transitory shock relaxes these precautionary saving motives, and thus, households with stronger precautionary saving motives exhibit higher MPCs. Peruvian households can have much stronger precautionary saving motives than U.S. households, as they accumulate far fewer liquid assets[48], and this can explain part of the mean MPC gap between the two economies.

Second, Peruvian households exhibit higher MPCs than U.S. households even in an environment where households can smooth their consumption without income risk and borrowing limits. Consider an economy where income risk and borrowing limits are removed from the model ('deterministic economy,' hereafter). Households know their deterministic income path and smooth their consumption according to the Euler equation

$$C_{i,t}^{-\sigma} = \beta(1+r)C_{i,t+1}^{-\sigma}.$$

When $\beta(1+r)$ is less than 1, households frontload their consumption or, equivalently, intertemporally allocate more resources to $C_{i,t}$ than to $C_{i,t+1}$. Because the value of $\beta(1+r)$ is substantially lower in Peru than in the U.S., Peruvian households frontload their consumption more heavily than U.S. households. In response to a positive income shock, households increase their consumption path proportionally (to the extent that the shock increases the present value of their lifetime wealth), and Peruvian households increase their current consumption more than U.S. households because of their heavier frontloading. Formally, by combining the Euler equation with budget constraints (SBC), we can analytically derive the consumption function of households at age t_0 as follows.

$$C_{i,t} = \{\beta(1+r)\}^{\frac{t-t_0}{\sigma}} \frac{1-[\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]}{1-[\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]^{T-t_0+1}} \left\{(1+r)A_{i,t_0-1} + \sum_{s=t_0}^{T}\left(\frac{1}{1+r}\right)^{s} Y_{i,s}\right\},\ t_0 \leq t \leq T.$$

Therefore, in this deterministic economy, the quarterly MPC to a transitory shock for households at age t_0 is

$$MPC_{DET}^{Q}(t_0) = \frac{1-[\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]}{1-[\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]^{T-t_0+1}}.$$

Note that $MPC_{DET}^{Q}(t_0)$ is equal to the share of the first term in the sum of a geometric series $\sum_{s=0}^{T-t_0}[\{\beta(1+r)\}^{1/\sigma}(1+r)^{-1}]^s$ and thus decreases in $\{\beta(1+r)\}^{1/\sigma}(1+r)^{-1}$. Because $\beta(1+r)$ is substantially smaller in Peru than in the U.S., while $(1+r)$ is very close to 1 in both Peru and the U.S., Peruvian households exhibit higher MPCs than U.S. households in this deterministic economy.[49] Using the analytical consumption function derived above, we can also analytically

[48]See footnote 45.

[49]As discussed above, $MPC_{DET}^{Q}(t_0)$ decreases in $\{\beta(1+r)\}^{1/\sigma}(1+r)^{-1}$, and $\{\beta(1+r)\}^{1/\sigma}$ appears here because it governs the degree of frontloading in consumption smoothing. What about $(1+r)^{-1}$? This term appears here because of the interest rate's wealth effect: when r is higher, households face a relatively cheaper price of future consumption (which is $1/(1+r)^{t-t_0}$ unit of current consumption) and thus can consume more today. Since the value of $(1+r)$ is very close to 1 in both Peru and the U.S., however, this effect is negligible.

determine annual MPC as follows.

$$MPC^A_{DET}(t_0) = \frac{1 - [\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]}{1 - [\{\beta(1+r)\}^{\frac{1}{\sigma}}(1+r)^{-1}]^{T-t_0+1}} \left\{ \sum_{s=0}^{3} \{\beta(1+r)\}^{\frac{s}{\sigma}} \right\}. \tag{7}$$

To quantitatively evaluate each component's contribution to the mean MPC gap between Peru and the U.S., I run counterfactual experiments by removing borrowing constraints and income risk sequentially in each of the Peruvian and U.S. model economies. Figures 7a and 7b compare the benchmark and counterfactual economies for Peru and the U.S., respectively. In these figures, the black solid line (labeled 'ZBL') represents the annual MPC at each residual income decile in the benchmark economy. The black dotted line (labeled 'NBL') represents the MPCs in an economy where the borrowing constraints (LQC) are removed from the ZBL economy or, equivalently, the zero borrowing limit $\bar{B} = 0$ is replaced with natural borrowing limits $\bar{B}^N_t = \sum_{s=1}^{T-t}(\frac{1}{1+r})^s \exp(\omega_{t+s} + y^{min}_{t+s})$, where y^{min}_t is the minimum possible realized value of y_t. The black dashed line (labeled 'DET') represents the MPCs in the deterministic economy, which removes income risk from the NBL economy.[50] Note that in each country, the MPC gap between the NBL and DET economies is generated by a precautionary saving motive to prepare for the realization of a low-income path (or, equivalently, prudence in Kimball (1990)'s term), and the MPC gap between the ZBL and NBL economies is generated by a precautionary saving motive to avoid borrowing constraints.

Figures 7a and 7b show that the precautionary saving motives explain most of the mean MPC gap between Peru and the U.S. The mean MPC gap between the ZBL and DET economies is 37.8%p in Peru, while it is 8.9%p in the U.S. The difference between the two (37.8%p - 8.9%p = 28.9%p)

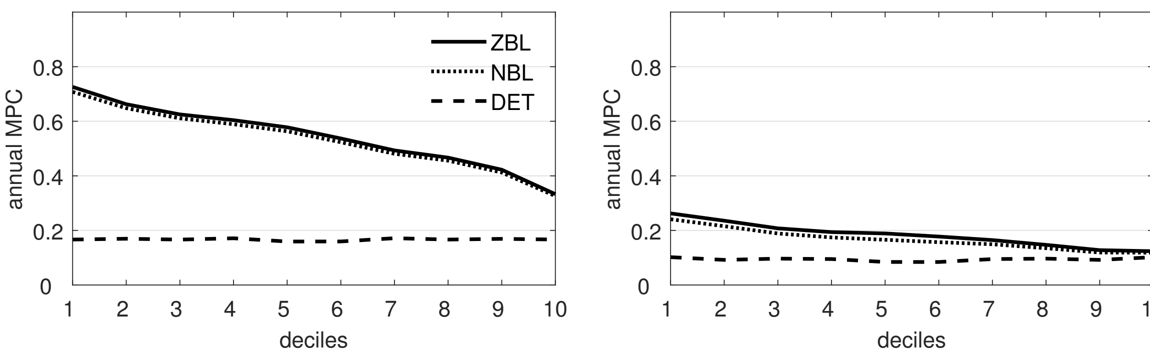

(a) Peru

(b) U.S.

Figure 7: Annual MPCs: ZBL vs NBL vs DET

Notes: Figures 7a and 7b plot the annual MPCs of the benchmark and counterfactual economies for Peru and the U.S., respectively. ZBL represents the benchmark economy, NBL represents an economy in which the zero borrowing constraints are removed from the ZBL economy, and DET represents a deterministic economy in which income risk is removed from the NBL economy.

[50]When computing the annual MPCs in the deterministic economy, I group simulated households by their residual income in quarter 0 and drop those becoming older than 65 at or before quarter 8, as in footnote 47. Then, I use household ages in quarter 4 and equation (7) to compute annual MPCs in this economy.

explains 79.9% of the mean MPC gap between Peru and the U.S. (36.2%p) of the benchmark (ZBL) case.

Figures 7a and 7b also suggest that between the two precautionary saving motives, households' fear of the realization of a low income path (or, equivalently, Kimball (1990)'s prudence) is the main driver of the mean MPC gap between Peru and the U.S. rather than their concern for being hit by zero borrowing limits. Admittedly, however, the role of zero borrowing constraints is tightly restricted by a noneconomic reason in the model. The power of zero borrowing constraints crucially depends on the location of the natural borrowing limits.[51] Therefore, the minimum income level is an important determinant of the power of zero borrowing limits. In my model, the age-specific minimum income levels are determined by the income process discretization and are very close to zero.[52] As a result, zero borrowing limits are very close to natural borrowing limits in my model and thus have a negligible effect on MPCs. In Online Appendix F, I use an alternative income grid that yields greater natural borrowing limits and verify that the zero borrowing constraints can have a stronger effect on MPCs. To have a fair evaluation of the relative importance between the two precautionary saving motives, an evidence-based calibration for the minimum possible income levels needs to be performed.

When the deterministic (DET) economies are compared between Peru and the U.S., Peruvian MPCs are still noticeably higher than U.S. MPCs, as discussed above: the mean annual MPC in Peru (16.7%) is 1.8 times as large as that in the U.S. (9.4%). This MPC gap between the Peruvian and U.S. deterministic economies (16.7% - 9.4% = 7.3%p), which comes from Peruvian households' heavier frontloading behavior, accounts for 20.1% of the mean MPC gap between Peruvian and U.S. benchmark (ZBL) economies (36.2%p).

In addition to the mean MPC gap between Peru and the U.S., Figures 7a and 7b also shed light on what drives the stronger MPC heterogeneity over residual income deciles in Peru. In the deterministic (DET) economies, there is no MPC heterogeneity over the deciles in either Peru or the U.S. This is because MPCs depend only on age in the deterministic economies (as shown by equation (7)) and residual income $y_{i,t}$ is independent of age. On the other hand, households in lower income deciles have stronger precautionary saving motives and thus exhibit higher MPCs in both Peruvian and U.S. model economies. Importantly, Figures 7a and 7b suggest that as households move down from higher to lower income deciles, precautionary saving motives become stronger more rapidly in Peru than in the U.S., generating stronger MPC heterogeneity in Peru.

[51] For example, in a model where natural borrowing limits are formed at zero due to the presence of a nonzero probability on zero income realization, as in Carroll (1997), imposing zero borrowing limits has no effect.

[52] In my Peruvian NBL economy, for example, the natural borrowing limit is 1.4 times the average quarterly labor income ($E[Y_{i,t}]$) at $t = 0$ (the first quarter of age 25) and monotonically decreases to 0.008 times the average quarterly labor income at $t = T - 1$ (the third quarter of age 65). This compares to Kaplan and Violante (2010), in which the natural borrowing limit is 5.8 times the average annual labor income at age 25 and 2.5 times the average annual labor income even at age 50.

7 Alternative Specifications

In this section, I consider alternative specifications either in the model simulation or in the MPC estimation. In subsections 7.1 and 7.2, I change some model parameters and examine how the model predictions change to understand what are the important factors for certain aspects of the simulation results from the benchmark model. In subsection 7.3, I use alternative data treatments when estimating MPCs and check whether the results are robust.

7.1 Permanent Component of Residual Income

In this subsection, I replace the persistent component of residual income $P_{i,t}$, which follows an AR(1) process, with a permanent component that follows a random walk. Parameters governing the stochastic process for residual income $y_{i,t}$ (ρ, σ_{ps}, σ_{tr}, and σ_{P_0}) are estimated in the same way that they are estimated in subsections 4.2 (for Peru) and 5.2 (for the U.S.), except that ρ is now restricted to 1. Online Appendix E.3 provides the estimation results. Two observations are noteworthy. First, as a consequence of assuming $\rho = 1$, the estimated Peruvian income process fails to capture the data pattern that yearly-age-specific covariances $cov[y_{i,t}, y_{i,t+4}|4a \leq t \leq 4a+3]$'s (in which a is a yearly age) are significantly greater than $cov[y_{i,t}, y_{i,t+4k}|4a \leq t \leq 4a+3]$'s, $k \geq 2$ in Peru. Second, $(\sigma^2_{ps}/\sigma^2_{P_0})$ is very small in Peru (0.0015) and much smaller than that in the U.S. (0.0041) This estimation result reflects the data pattern that the yearly-age-specific variances of residual income are flat with age in Peru, while they noticeably increase with age in the U.S.

With the newly calibrated labor income process, I recalibrate β by targeting MPC estimates, as in subsections 4.2 (for Peru) and 5.2 (for the U.S.). Figure 8 plots the model predictions under the recalibration. Two important observations emerge. First, as Figures 8a and 8b show, the model fits the MPC estimates ignoring the heterogeneity over residual income deciles. Nevertheless, the model-predicted MPCs are not statistically significantly different from the MPC estimates when jointly tested by the Wald test in both Peru and the U.S. However, these model-predicted MPCs

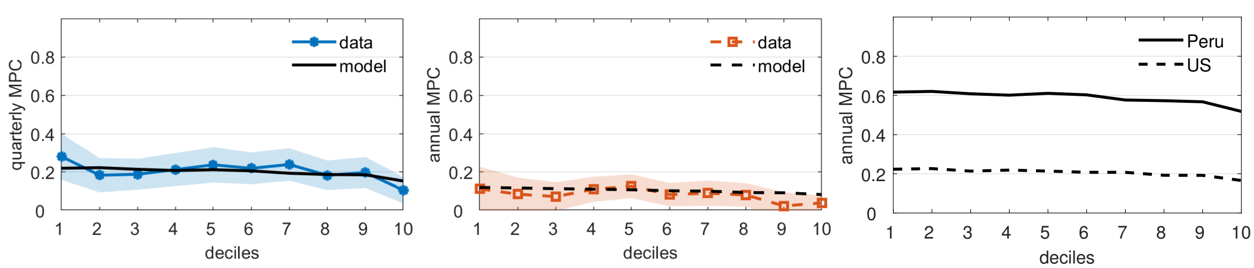

(a) Model's fit in Peru (b) Model's fit in the U.S. (c) Model-based comparison

Figure 8: Model outcomes under the restriction $\rho = 1$

Notes: These figures plot simulation results when the model is recalibrated under the assumption that $\rho = 1$. Figure 8a plots the Peruvian quarterly MPC estimates and their model counterparts. Figure 8b plots the U.S. annual MPC estimates and their model counterparts. Figure 8c compares true annual MPCs between the Peruvian and U.S. model economies.

fail to account for the statistically significant differences found in the pairwise MPC comparison between deciles discussed in subsections 4.1 and 5.1. Second, as Figure 8c shows, the model still predicts a large mean annual MPC gap between Peru and the U.S. (59.0% in Peru, 20.6% in the U.S.), while it no longer predicts stronger MPC heterogeneity in Peru.

This result suggests that a precise estimation of ρ is important not only because the value of ρ matters for explaining income data patterns but also because it determines the degree of MPC heterogeneity in the model.

An economic intuition behind this result is as follows. Under a small value of $\rho < 1$, a lower value of $P_{i,t}$ lowers the current cash-on-hand substantially more than the households' lifetime wealth because the persistent component reverts back toward its mean. Therefore, households with a lower value of $P_{i,t}$ feel poorer (relative to their lifetime wealth), have a stronger precautionary saving motive, and exhibit higher MPC. As ρ approaches 1, however, a lower value of $P_{i,t}$ also significantly decreases lifetime wealth. As a result, households with a lower value of $P_{i,t}$ do not feel much poorer (relative to their lifetime wealth), do not have a much stronger precautionary saving motive, and do not exhibit much higher MPC.[53]

7.2 Switching the Income Process

One natural hypothesis arising from Figure 4 is that Peruvian MPCs might be more heterogeneous simply because their residual incomes are more heterogeneous. In Online Appendix H.1, I plot Figure 4 on the x-axis of group-average residual incomes and verify that the group averages of Peruvian quarterly residual incomes are indeed more heterogeneous than those of U.S. annual residual incomes. This hypothesis is, however, subject to two problems. First, the fact that U.S. annual residual incomes are less dispersed than Peruvian quarterly residual incomes does not necessarily mean that U.S. households face a smaller income risk because annual residual incomes are supposed to be less dispersed than quarterly residual incomes. For example, in the U.S. model economy calibrated in subsection 5.2, the unconditional variance of quarterly residual income is 0.474, while that of annual residual income is 0.272. Second, comparing only the residual income dispersion ignores the relative contribution of persistent and transitory risks, while households respond to them very differently.

Reflecting these two problems, the simple hypothesis above can be refined into the following question: how do the differences in the income process between Peru and the U.S. contribute to their differences in MPCs? In this subsection, I address this question using the model economies. Specifically, I switch the income processes between Peru and U.S. (or, equivalently, the Peruvian income process is replaced with the U.S. one in the Peruvian economy and vice versa) and examine how the MPCs change.

Figure 9a plots annual MPCs in the Peruvian model economy but with the U.S. income process

[53] Carroll (1997) formally shows in his buffer stock model with $\rho = 1$ that there exists a target cash-on-hand-to-permanent-income ratio such that households save when the ratio is below the target and dissave when above the target, implying that the precautionary saving (after being normalized by permanent income) is independent of the permanent income level.

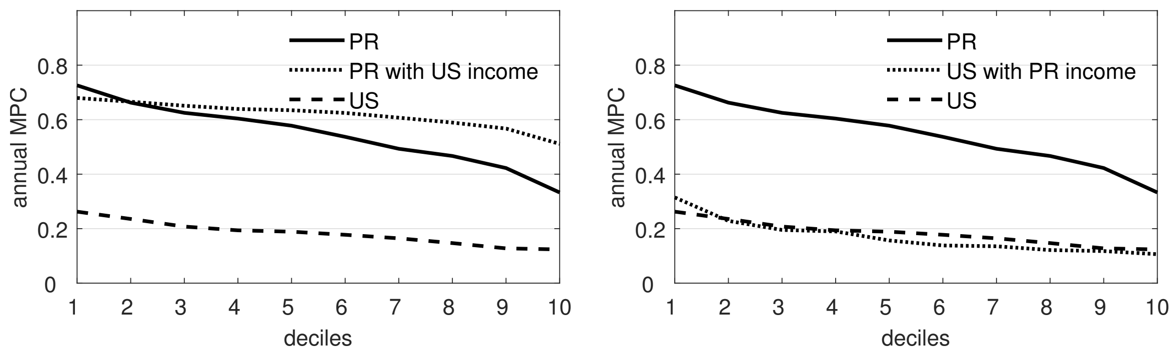

(a) Peru with the U.S. income process

(b) U.S. with the Peruvian income process

Figure 9: Switching the income process

Notes: Figure 9a plots annual MPCs in the Peruvian model economy but with the U.S. income process (labeled 'PR with US income') and compares them with annual MPCs in the benchmark Peruvian and U.S. economies (labeled 'PR' and 'US,' respectively). In the same way, Figure 9b plots annual MPCs in the U.S. model economy but with the Peruvian income process (labeled 'US with PR income') and compares them with annual MPCs in the benchmark Peruvian and U.S. economies (labeled 'PR' and 'US,' respectively).

(labeled 'PR with US income') and compares them with annual MPCs in the benchmark Peruvian and U.S. economies (labeled 'PR' and 'US,' respectively). Specifically, the parameters governing the labor income process, including ρ, σ_{ps}, σ_{tr}, σ_{P_0}, and $\{\omega_t\}_{t=0}^T$, are replaced. Two observations are noteworthy. First, the MPC heterogeneity over residual income deciles becomes weaker. This is because ρ becomes closer to 1 as it changes from 0.963 to 0.989.[54] Second, switching the income process does not narrow the mean MPC gap between Peru and the U.S. (rather, it widens.) This observation implies that income process differences are not the main cause of the stronger precautionary saving behavior in Peru than in the U.S.

Figure 9b plots annual MPCs in the U.S. model economy but with the Peruvian income process (labeled 'US with PR income') and compares them with annual MPCs in the benchmark Peruvian and U.S. economies (labeled 'PR' and 'US,' respectively). This figure reconfirms the main takeaways from Figure 9a. First, switching the income process in the U.S. economy strengthens the MPC heterogeneity over residual income deciles as ρ changes from 0.989 to 0.963, becoming farther from 1. Second, switching the income process does not narrow the mean MPC gap between Peru and the U.S. in the U.S. economy.

Note that in the model, the Peruvian and U.S. economies are distinguished by a finite number of parameters, including those governing the labor income process (namely, ρ, σ_{ps}, σ_{tr}, σ_{P_0}, and $\{\omega_t\}_{t=0}^T$), r, and β. Therefore, the Peruvian model economy with the U.S. income process in Figure 9a can also be interpreted as the U.S. model economy with the Peruvian r and β. In the same way, the U.S. model economy with the Peruvian income process in Figure 9b can be interpreted as the Peruvian model economy with the U.S. r and β. Under this interpretation, another important observation can be made from Figures 9a and 9b: the differences in (r, β) are the main driver of the

[54]See subsection 7.1 for a related discussion.

mean MPC gap between the Peruvian and U.S. model economies. As discussed in section 6, this is not because the Peruvian r and β induce heavier frontloading behavior in household consumption smoothing but because they induce a stronger precautionary saving motive in the model.

7.3 Robustness under Alternative Data Treatments

I estimate MPCs under various alternative data treatments to check whether the main results of this paper are robust. I find that in each case, the following results robustly emerge. i) When annualizing Peruvian MPC estimates, the model-free and model-based methods yield similar outcomes. ii) The annual U.S. MPC estimates have a time aggregation problem in the quarterly model. iii) Under both model-free and model-based comparisons, Peruvian MPCs are substantially higher overall than U.S. MPCs, and iv) MPCs are also more heterogeneous over residual income deciles in Peru than in the U.S. In this subsection, I provide a brief description of each alternative data treatment. Online Appendix G provides a detailed description and the results of each case.

The alternative data treatments I conduct can be grouped into four categories: alternative variable construction, alternative grouping of households, alternative empirical models, and alternative sample selection.

7.3.1 Alternative Variable Construction

I conduct four robustness checks by constructing consumption and income differently. First, the baseline consumption excludes nonpurchased consumption, such as donations, food stamps, in-kind income, and self-production. I conduct a robustness check by including these items in consumption.

Second, the baseline U.S. consumption does not include clothing, recreation, alcohol, and tobacco due to a narrow coverage on expense items in the early waves of the PSID, while the baseline Peruvian consumption includes these items. I conduct a robustness check by consistently excluding these expense items from the Peruvian consumption.

Third, I exclude the imputed components of missing income from the baseline Peruvian income, while I cannot do the same for the U.S. income, as the imputed income components are not distinguishable in Kaplan et al. (2014b)'s dataset. I conduct a robustness check by consistently including the imputed components of missing income in the Peruvian income.

Fourth, the baseline Peruvian income includes two expense items that are also included in their consumption: rental equivalence of housing provided by work (as labor income) and rental equivalence of donated housing (as transfers).[55] On the other hand, the baseline U.S. income does not include any expense items that are included in their consumption. I conduct a robustness check by consistently excluding the two expense items from the Peruvian income.

[55] The rental equivalence of owned housing is categorized as capital income and thus not included in the baseline income of Peruvian households.

7.3.2 Alternative Grouping of Households

When constructing residual income deciles in the baseline estimation, I sort U.S. observations within each calendar year and Peruvian observations within each calendar quarter in accordance with the reference period of each sample (a year for the U.S. sample and a quarter for the Peruvian sample). However, because I already remove time fixed effects when controlling for the predictable components (annually for the U.S. sample, quarterly for the Peruvian sample), it should also be fine to sort residual incomes in a larger observation pool than the pool of the reference period. Therefore, I conduct robustness checks by sorting residual incomes in different observation pools. First, I sort Peruvian observations within each calendar year in accordance with how U.S. observations are sorted in the baseline analysis. Second, I sort observations in the pool of the whole sample years in both the Peruvian and U.S. samples.

7.3.3 Alternative Empirical Models

I conduct three robustness checks by changing the empirical model. First, motivated by Figure 2b, I re-estimate MPCs using a revised estimator of ψ_G that replaces $\Delta^K y_{i,t}$ with $\tilde{\Delta}^K y_{i,t}$, where $\tilde{\Delta} y_{i,t} = y_{i,t} - \tilde{\rho}^K y_{i,t-K}$. For the Peruvian quarterly MPC estimation, $\rho = 0.963$ reported in Table 1 is used as the value of $\tilde{\rho}$. For the U.S. annual MPC estimation, t is a year, and thus, $\tilde{\rho}$ is an annual autocorrelation coefficient. I obtain $\tilde{\rho}$ by estimating the annual income process (assumed in the U.S. annual MPC estimation) using age-specific variances and covariances of annual residual incomes. As a result, I obtain an annual autocorrelation coefficient of 0.958. This value turns out to be very close to 0.989^4, where 0.989 is the value of the quarterly autocorrelation coefficient reported in Table 2.

Second, I revise the empirical model by incorporating subsistence levels. Specifically, I consider the household utility function developed by Stone (1954) and Geary (1950), under which households obtain utility only from consumption beyond a subsistence point, and I revise the MPC estimation equation accordingly.

Third, in the spirit of Crawley (2020), I also consider a continuous-time model as another way to obtain and compare MPC estimates without a time aggregation problem. As in Crawley (2020), moment conditions are derived from a continuous time model and used for the MPC estimation.[56]

7.3.4 Alternative Sample Selection

Lastly, I conduct four robustness checks by revising the sample selection procedure. I consider i) an alternative range of household heads' ages, ii) a revision of the definition of income outliers,

[56] However, the model is not exactly equal to that of Crawley (2020). In particular, Crawley (2020) assumes a random walk consumption function under which consumption responds only to current transitory and permanent shocks, as in Blundell et al. (2008), while I specify a consumption function such that dynamic consumption responses to a transitory income shock decay exponentially over time. My specification is motivated by the observation in subsection 4.4 that Auclert (2019)'s model-free annualization formula (6), which is derived from an assumption that dynamic consumption responses to a transitory shock die out exponentially over time in a quarterly model, provides a good approximation. See Online Appendix G.8 for details.

iii) a sample restriction to male heads, and iv) a stricter rule in detecting panel observations that are likely to connect two different households (in the third step of the sample selection described in subsection 3.3).

8 Discussion on External Validity

The main goal of this paper is to provide an emerging economy's nationally representative MPC estimates, which can be used to discipline a macroeconomic model. I choose Peru because of data availability. Given the current absence of comparable estimates for other emerging economies, a question regarding the generalizability of the findings of this paper naturally arises. In this section, I have a discussion on external validity.

Peru is a typical emerging economy in several important dimensions. First, Peru's national income level in terms of average annual GDP per capita during 2004-2016 in PPP (Purchasing Power Parity)-converted, constant 2017 U.S. dollars is 9,998.9 dollars, which is substantially lower than the average of rich economies (48,650.1 dollars) and close to the average of emerging and poor economies (8,414.6 dollars).[57] Second, Peru also exhibits stylized patterns of emerging market business cycles, as noted by Hong (2020). Namely, output is very volatile (compared to that of developed economies), consumption is more volatile than output, and trade balance is countercyclical. Third, the financial institutions and markets are underdeveloped in Peru compared to those in rich economies, as they typically are in other emerging and poor economies, according to two widely used indicators for financial development, World Bank's Global FINDEX and International Monetary Fund's Financial Development Index.[58] [59]

Some of these typical features of emerging economies might be related to their MPCs. For example, Hong (2020) matches the high levels of Peruvian MPC estimates in his quantitative model using financial frictions households face and finds that the phenomenon of excess consumption volatility (or, equivalently, aggregate consumption volatility being greater than aggregate output volatility) is generated by Peruvian households' high MPCs and correspondingly strong precautionary saving behavior. Under this theory, the coexistence of an underdeveloped financial system and the exhibition of excess consumption volatility, which are typical features of emerging and poor economies, are suggestive of high MPCs, as in Peru.

[57] Data on GDP per capita in PPP-converted, constant 2017 U.S. dollars come from World Bank's World Development Indicators (WDI) database. When computing the group averages, I use Uribe and Schmitt-Grohe (2017)'s categorization of rich, emerging, and poor countries and reflect population weights obtained also from World Bank's WDI database.

[58] The headline index in World Bank's Global FINDEX database is the share of people having an account at a financial institution or at a mobile money service. The average headline index of emerging and poor countries in the 2011 survey is 43.0%, which is substantially lower than the average of rich countries, 91.8%. The index for Peru is 20.5%. IMF's Financial Development Indicator evaluates financial development in six aspects (the depth, accessibility, and efficiency of institutions and markets) and aggregates them. The average headline index of emerging and poor countries during 2004-2016 is 0.403, which is again substantially lower than the average of rich countries, 0.829. The index for Peru is 0.296. When computing the group averages, I use Uribe and Schmitt-Grohe (2017)'s categorization of rich, emerging, and poor countries and reflect population weights obtained from World Bank's WDI database.

[59] For a report on the Global FINDEX database, see Demirguc-Kunt and Klapper (2012), Demirguc-Kunt, Klapper, Singer, and Van Oudheusden (2015), and Demirguc-Kunt, Klapper, Singer, and Ansar (2018). For a report on the Financial Development Indicator database, see Svirydzenka (2016).

Outside of such theory, however, this paper itself does not guarantee any external validity from an empirical point of view. In particular, we cannot rule out the possibility that Peru has a certain feature that is not typical of other emerging and poor economies and that this feature substantially contributes to the high MPC estimates. For example, Peru is highly dollarized in its financial intermediation (regarding both loans and deposits) compared to other emerging economies.[60] If both household income and consumption are affected by a third factor, such as exchange rate fluctuations, because of the high financial dollarization, this can bias my MPC estimates. Fortunately, household incomes in my Peruvian sample are unlikely to be affected by the financial dollarization since wages are barely dollarized in Peru, as pointed out by Contreras, Quispe, and Regalado (2017), and my income measure is labor income. Nevertheless, this example illustrates how a Peru-specific feature can potentially affect the MPC estimation and limit the external validity.

The only way to obtain empirically supported external validity is to conduct similar exercises using more data. To this end, further research needs to be conducted by extending the analyses of this paper to micro data from other emerging and poor economies.

9 Conclusion

This paper estimates MPCs out of transitory income shocks using a nationally representative Peruvian household survey. I compare Peruvian and U.S. MPCs reflecting the different reference periods of the underlying surveys through the lens of a standard incomplete-market, life-cycle model. I find that Peruvian MPCs are substantially higher overall than U.S. MPCs and are also substantially more heterogeneous over income deciles.

In the burgeoning literature examining how microlevel household behavior and its heterogeneity matter for the macroeconomy, researchers have discovered novel mechanisms through which high-MPC households affect aggregate dynamics and policy effects in the context of developed economies. The results of this paper suggest that these mechanisms could play a significantly larger role in emerging economies. In this regard, this paper supports the importance of recent efforts to expand the literature to open and emerging economies. I expect the estimates of this paper to serve as useful target moments for the key object, MPC, in this line of research.

[60] For example, see Figures 3 and 4 of Dalgic (2020).

References

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Altonji, J. G. and L. M. Segal (1996). Small-Sample Bias in GMM Estimation of Covariance Structures. *Journal of Business & Economic Statistics 14*(3), 353–366.

Atkeson, A. and L. E. Ohanian (2001). Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review 25*(1), 2–11.

Attanasio, O. P. and G. Weber (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy 103*(6), 1121–1157.

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

Auclert, A., B. Bardoczy, M. Rognlie, and L. Straub (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica 89*(5), 2375–2408.

Auclert, A., M. Rognlie, M. Souchier, and L. Straub (2021). Exchange Rates and Monetary Policy with Heterogeneous Agents: Sizing up the Real Income Channel. National Bureau of Economic Research Working Paper 28872.

Banco Central de Reserva del Perú (n.d.a). Interest rates: EMBIG (variation in bps) – Spread – EMBIG Peru (bps) (Series Code: PD04709XD). https://estadisticas.bcrp.gob.pe/estadisticas/series/. Original Source: J.P. Morgan.

Banco Central de Reserva del Perú (n.d.b). Metropolitan Lima Price Index (2009 = 100) – Consumer Price Index (CPI) (Series Code: PN01270PM). https://estadisticas.bcrp.gob.pe/estadisticas/series/. Original Source: Instituto Nacional de Estadística e Informática.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018). House Prices and Consumer Spending. *The Review of Economic Studies 85*(3), 1502–1542.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review 98*(5), 1887–1921.

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017). The Distribution of Wealth and the Marginal Propensity to Consume. *Quantitative Economics 8*(3), 977–1020.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics 112*(1), 1–55.

Carroll, C. D. (2006). The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems. *Economics Letters 91*(3), 312–320.

Contreras, A., Z. Quispe, and F. Regalado (2017). Real Dollarization and Monetary Policy in Peru. *IFC Bulletins Chapters 43*.

Crawley, E. (2020). In Search of Lost Time Aggregation. *Economics Letters 189*, 108998.

Crawley, E. and A. Kuchler (2021). Consumption Heterogeneity: Micro Drivers and Macro Implications. Working Paper.

Dalgic, H. (2020). Financial Dollarization in Emerging Markets: An Insurance Arrangement.

Working Paper.

De Ferra, S., K. Mitman, and F. Romei (2020). Household Heterogeneity and the Transmission of Foreign Shocks. *Journal of International Economics 124*, 103303.

Demirguc-Kunt, A., L. Klapper, D. Singer, and S. Ansar (2018). *The Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution*. World Bank Publications.

Demirguc-Kunt, A. and L. F. Klapper (2012). Measuring Financial Inclusion: The Global Findex Database. World Bank Policy Research Working Paper 6025.

Demirguc-Kunt, A., L. F. Klapper, D. Singer, and P. Van Oudheusden (2015). The Global Findex Database 2014: Measuring Financial Inclusion around the World. World Bank Policy Research Working Paper 7255.

Diaz-Gimenez, J., V. Quadrini, and J.-V. Rios-Rull (1997). Dimensions of Inequality: Facts on the US Distribution of Earnings, Income and Wealth. *Federal Reserve Bank of Minneapolis Quarterly Review 21*(2), 3–21.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. W. Walker (2019). General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya. National Bureau of Economic Research Working Paper 26600.

Federal Reserve Bank of St. Louis (n.d.a). 3-Month Treasury Bill: Secondary Market Rate, Percent, Monthly, Not Seasonally Adjusted (Series Code: TB3MS). `https://fred.stlouisfed.org`. Original Source: Board of Governors of the Federal Reserve System (US).

Federal Reserve Bank of St. Louis (n.d.b). Consumer Price Index for All Urban Consumers: All Items in U.S. City Average, Index 1982-1984=100, Monthly, Seasonally Adjusted (Series Code: CPIAUCSL). `https://fred.stlouisfed.org`. Original Source: U.S. Bureau of Labor Statistics.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The Next Generation of the Penn World Table. *American Economic Review 105*(10), 3150–3182. Available for Download at `www.ggdc.net/pwt`.

Fella, G., G. Gallipoli, and J. Pan (2019). Markov-Chain Approximations for Life-Cycle Models. *Review of Economic Dynamics 34*, 183–201.

Ferrante, F. and N. Gornemann (2021). Financial Frictions and the Re-distributive Effects of Exchange Rate Fluctuations. Working Paper.

Floden, M. and J. Linde (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? *Review of Economic Dynamics 4*(2), 406–437.

Garcia-Cicco, J., R. Pancrazi, and M. Uribe (2010). Real Business Cycles in Emerging Countries? *American Economic Review 100*(5), 2510–2531.

Geary, R. C. (1950). A Note on "A Constant-Utility Index of the Cost of Living". *The Review of Economic Studies 18*(1), 65–66.

Guerriero, M. (2019). The Labor Share of Income around the World: Evidence from a Panel Dataset. In *Labor Income Share in Asia*, pp. 39–79. Springer.

Guntin, R., P. Ottonello, and D. Perez (2020). The Micro Anatomy of Macro Consumption Adjustments. National Bureau of Economic Research Working Paper 27917.

Guo, X., P. Ottonello, and D. Perez (2020). Monetary Policy and Redistribution in Open Economies.

Technical report, National Bureau of Economic Research Working Paper 28213.

Guvenen, F. and A. A. Smith (2014). Inferring Labor Income Risk and Partial Insurance from Economic Choices. *Econometrica 82*(6), 2085–2129.

Haushofer, J. and J. Shapiro (2016). The Short-term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya. *The Quarterly Journal of Economics 131*(4), 1973–2042.

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006. *Review of Economic Dynamics 13*(1), 15–51.

Hong, S. (2020). Emerging Market Business Cycles with Heterogeneous Agents. Working Paper.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary Saving and Social Insurance. *Journal of Political Economy 103*(2), 360–399.

Huggett, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control 17*(5-6), 953–969.

Huggett, M. (1996). Wealth Distribution in Life-Cycle Economies. *Journal of Monetary Economics 38*(3), 469–494.

Instituto Nacional de Estadística e Informática (2004-2016). Encuesta Nacional de Hogares. http://iinei.inei.gob.pe/microdatos/.

Instituto Nacional de Estadística e Informática (n.d.). Metodología de Cálculo del Producto Bruto Interno Anual. https://www.inei.gob.pe/media/MenuRecursivo/metodologias/pbi02.pdf.

International Monetary Fund (n.d.). Financial Development Index Database. https://data.imf.org/?sk=F8032E80-B36C-43B1-AC26-493C5B1CD33B.

Jappelli, T. and L. Pistaferri (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–136.

Jappelli, T. and L. Pistaferri (2017). *The Economics of Consumption: Theory and Evidence*. Oxford University Press.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review 96*(5), 1589–1610.

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

Kaplan, G. and G. L. Violante (2010). How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87.

Kaplan, G., G. L. Violante, and J. Weidner (2014a). Replication Data for: "The Wealthy Hand-to-Mouth". https://gregkaplan.me/academic-publications.

Kaplan, G., G. L. Violante, and J. Weidner (2014b). The Wealthy Hand-to-Mouth. *Brookings Papers On Economic Activity*, 77–138.

Kimball, M. S. (1990). Precautionary Saving in the Small and in the Large. *Econometrica 58*(1), 53–73.

Kocherlakota, N. and L. Pistaferri (2009). Asset Pricing Implications of Pareto Optimality with Private Information. *Journal of Political Economy 117*(3), 555–590.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and Household Heterogeneity. In

Handbook of Macroeconomics, Volume 2, pp. 843–921. Elsevier.

Krueger, D. and F. Perri (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *The Review of Economic Studies 73*(1), 163–193.

McKay, A., E. Nakamura, and J. Steinsson (2016). The Power of Forward Guidance Revisited. *American Economic Review 106*(10), 3133–3158.

Neumeyer, P. A. and F. Perri (2005). Business Cycles in Emerging Economies: the Role of Interest Rates. *Journal of Monetary Economics 52*(2), 345–380.

Oh, H. and R. Reis (2012). Targeted Transfers and the Fiscal Response to the Great Recession. *Journal of Monetary Economics 59*, S50–S64.

Oskolkov, A. (2022). Exchange Rate Policy and Heterogeneity in Small Open Economies. Working Paper.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review 103*(6), 2530–2553.

Paxson, C. H. (1992). Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand. *The American Economic Review*, 15–33.

Rouwenhorst, K. G. (1995). Asset Pricing Implications of Equilibrium Business Cycle Models. In T. Cooley (Ed.), *Frontiers of Business Cycle Research*, Chapter 10, pp. 294–330. Princeton University Press.

Stone, R. (1954). Linear Expenditure Systems and Demand Analysis: an Application to the Pattern of British Demand. *The Economic Journal 64*(255), 511–527.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Consumption and Risk Sharing over the Life Cycle. *Journal of Monetary Economics 51*(3), 609–633.

Svirydzenka, K. (2016). *Introducing a New Broad-Based Index of Financial Development*. International Monetary Fund.

Uribe, M. and S. Schmitt-Grohe (2017). *Open Economy Macroeconomics*. Princeton University Press.

Uribe, M. and V. Z. Yue (2006). Country Spreads and Emerging Countries: Who Drives Whom? *Journal of International Economics 69*(1), 6–36.

Villalvazo, S. (2021). Inequality and Asset Prices during Sudden Stops. Working Paper.

World Bank (n.d.a). The Global Findex Database. https://globalfindex.worldbank.org/#data_sec_focus.

World Bank (n.d.b). World Development Indicators – GDP per Capita, PPP (Constant 2017 International $) (Series Code: NY.GDP.PCAP.PP.KD). https://data.worldbank.org/indicator/.

World Bank (n.d.c). World Development Indicators – Population, total (Series Code: SP.POP.TOTL). https://data.worldbank.org/indicator/.

Zhou, H. (2021). Open Economy, Redistribution, and the Aggregate Impact of External Shocks. Working Paper.