

What Can Explain Excess Smoothness and Sensitivity of State-Level Consumption?

María José Luengo-Prado
Northeastern University

Bent E. Sørensen
*University of Houston
and CEPR*

June 13, 2006

Abstract

This article estimates marginal propensities to consume (MPC) out of current and lagged income for U.S. states using panel data regressions that control for time-specific and state-level fixed effects. The MPCs vary across states, in particular, the MPC out of current income is higher in states where income is more persistent and the MPC out of lagged income is lower in agricultural states. Several models of individual consumer behavior are analyzed and simulated in order to isolate a model which is able to match the estimated MPCs well.

KEYWORDS: Permanent Income, Credit rationing, Precautionary saving, Time-Aggregation, Durable Goods, Risk Sharing.

JEL CLASSIFICATIONS: E21 - Consumption; Saving.

1 Introduction

We estimate marginal propensities to consume (MPC) out of current and lagged income for U.S. states using panel data regressions that control for time-specific and state-level fixed effects. The MPCs vary across states, in particular, the MPC out of current income is higher in states where income is more persistent and the MPC out of lagged income is lower in agricultural states. We then analyze and simulate models of individual consumer behavior and isolate a model which is able to match the estimated MPCs well.

Virtually all current models of aggregate consumer behavior depart from the Friedman (1957)-Hall (1978) Permanent Income Hypothesis (PIH). Hall (1978) showed that if consumers are forward-looking, have rational expectations, evaluate consumption streams utilizing a quadratic utility function, and can freely borrow and lend at a constant interest rate, consumption is a martingale; i.e., a regression of period t consumption growth on any variable known at period $t - 1$ should return an estimate of zero. Regressions using aggregate data, however, consistently return an estimate significantly larger than zero when current growth in consumption is regressed on lagged aggregate income growth—a phenomenon known as “excess sensitivity” (of current consumption to lagged income). The PIH-model also provides closed-form solutions for the predicted growth in consumption as a function of innovations to income when income is described by a general Auto Regressive-Moving Average (ARMA) model. For example, if income is a random walk, consumption is predicted to move one-to-one with income. Empirical work using aggregate data consistently finds a significantly smaller reaction of consumption to income shocks—a phenomenon known as “excess smoothness.”¹

State-level data has several advantages for our study. Compared with purely macro approaches, the existence of 50 states with different income processes (some agricultural, some oil-based, etc.) vastly expands the relevant variation and the number of data points. In addition, by considering state-level variation that is orthogonal to aggregate variation, simultaneity problems are likely to be alleviated and U.S.-wide aggregate constraints will have little impact on the results as argued by Ostergaard, Sørensen, and Yosha (2002) and Dejuan and Luengo-

¹For studies of excess sensitivity, see Flavin (1981), Blinder and Deaton (1985), Campbell and Deaton (1989), and Attanasio and Weber (1993). Building on the results in Hansen and Sargent (1981), excess smoothness has been documented by Deaton (1987), Campbell and Deaton (1989), and Galí (1991).

Prado (2006).² The persistence of state-level income affects the current MPC significantly—with states with higher persistence exhibiting higher MPCs—consistent with a standard PIH model. However, state-level consumption displays strong excess smoothness and sensitivity. Further, we observe that excess sensitivity is higher in states where consumers are likely to face higher income uncertainty. These facts imply that the PIH-model needs to be extended. We will do so in stages in order to highlight the effect of each modification.

A simple and, we think, reasonable, explanation of excess sensitivity is that consumption decisions are made at a higher frequency than the frequency of observation (annual).³ We allow for a bi-annual frequency of income and consumption in the model and explicitly aggregate over time. This helps explain the average level of excess sensitivity. However, excess sensitivity of consumption is significantly higher in agricultural states which needs a separate explanation. Based on the observation that farmers' income is particularly volatile as shown by Carroll and Samwick (1997), we show that the so-called buffer-stock model of savings, pioneered by Deaton (1991) and Carroll (1992), can explain this pattern. This is a model of individual-level behavior that does not allow for a representative agent and we simulate the model for a large number of individuals, aggregate, and find the predicted MPCs by regressing on the simulated data. This approach, therefore, also controls for potential biases that may arise from considering a representative agent instead of aggregating across agents (see Attanasio and Weber 1993).

However, the buffer-stock model cannot fully explain the amount of excess smoothness as already pointed out by Ludvigson and Michaelides (2001). Including durable consumption—in the form of large indivisible durable goods that can only be purchased with a non-negligible down payment—generates significant smoothness as found by Luengo-Prado (2006). Nonetheless, the MPC out of current income is lower than the model with durable goods can explain with

²The state-level data is not plagued by the large amount of idiosyncratic variation in micro data and it is hard to know what patterns in micro-data is likely to survive aggregation. Compared to cross-country data, the data is collected in a consistent manner and most institutional features do not vary across states, making our results less likely to suffer from omitted variable bias. One drawback of our state-level data, compared to macroeconomic data, is that we only observe retail sales and not all components of consumption; however, no U.S. sources of micro data is better along this dimension. For example, the much used Panel Study of Income Dynamics mainly records food consumption rather than the total retail sales available at the state-level. In any event, studies based on state-level data provide a useful complement to other studies.

³An alternative modification that might explain excess sensitivity would be to allow for habit formation. We briefly explored this issue but did not find much evidence of habit formation in our data using common parametric specifications as used by, for instance, Dynan (2000).

“reasonable” parameterizations. To fully match the observed smoothness of consumption, we follow Attanasio and Pavoni (2006) and assume that measured income shocks are imperfect measures of shocks to consumers’ inter-temporal budgets.⁴ The models discussed so far assume that a consumer’s budget constraint is determined by labor income and interest from safe bond holdings. In the real world, asset income comes in many forms that often are not measured well, in particular, capital gains have been large and variable during the sample we consider. Also, individuals share risk with family members and through various ethnic, religious, etc. networks. Therefore, labor income may be an imperfect measure of innovations to a consumer’s present value of wealth, and a model that allows for partial risk sharing across states can explain the low observed MPCs. Adding a simple form of partial risk sharing to the buffer-stock model with durable goods and time-aggregation allows us to match the observed MPCs in a way that simpler models cannot.

The remainder of the paper is organized as follows. Section 2 describes our empirical results. Section 3 compares the ability of suitable parameterized models of consumer behavior to explain the results found in Section 2, while Section 4 summarizes our conclusions.

2 Panel Data Estimation of the Income Process and the MPCs

2.1 Estimating the time series process for income

In forward-looking models, consumption patterns depend on the income processes so we start by estimating time series models for income state-by-state for the years 1964–1998. Ostergaard, Sørensen, and Yosha (2002) show that disposable labor income growth at the state-level is well-modeled as an autoregressive (AR) model of order 1.⁵ A previous version of this paper, that focused on the buffer-stock model and uncertainty, allowed for a more complicated process for state-level income. Specifically, an additive i.i.d. income shock was allowed for beyond the AR-term. However, for most states this transitory aggregate shock was not significant and only for

⁴In the case of perfect Arrow-Debreu markets, consumption should not react to individual specific labor income shocks. However, an outcome where there is not perfect risk sharing (in the sense that labor income risk is not fully insured) can be a constrained optimal outcome if income depends on unobserved effort as rigorously developed by Attanasio and Pavoni (2006).

⁵We tested the level of disposable labor income for unit roots using an Augmented Dickey-Fuller test with one lag and could only reject the unit root for 6 and 1 states at the 5% and 1% level, respectively. Therefore, the growth rate of income is reasonably well modeled as a stationary variable.

a few states did we find the parameter of the transitory shock to be precisely estimated. We, therefore, find it preferable to use the simpler, more common, AR(1) specification. We elaborate on this point because the predicted impact on consumption of the uncertainty (variance) of state-level income depends on whether income shocks are transitory but our sample is too short to convincingly decompose the state-level income processes into permanent and transitory components.

Define the state-specific component of growth in disposable labor income as $\log G_{st} = \Delta \log Y_{st} - \Delta \log Y_t$. In the data, mean growth rates vary by state but in our theoretical discussion we do not explore the effects of state-varying growth rates. We make this choice because in-migration in some states, such as Nevada, is so large that mean growth rates over 20-odd years cannot easily be interpreted as reflecting the income growth prospects of individuals. In the estimation of state-specific processes, we demean the data prior to estimation, and estimate for each state s the process:

$$\log G_{st} = a^s \log G_{s,t-1} + \sigma_{G_s} u_{st},$$

where u_{st} is i.i.d. standard normal and $\sigma_{G_s}^2$ is the variance of the innovation term, $\sigma_{G_s} u_{st}$. We refer to a^s as a measure of persistence—the larger a^s the larger effect an innovation of a given size will have on expected future income.

In Table 1, we report the estimated values of a^s and σ_{G_s} for each of the 50 U.S. states—data sources are described in Appendix A. The persistence of state-specific shocks varies widely across states. The point estimate of a^s is -0.53 for Idaho while the largest value of 0.50 is found for Louisiana, with an average value of 0.07 . If the aggregate effect is not removed average persistence is significantly higher at 0.38 which reflects that the aggregate component of income growth displays more persistence than the state-specific component. This difference in persistence implies that forward-looking consumers will react more to aggregate than to state-specific shocks in most states. Estimating an AR(1) process for aggregate income growth, $\Delta \log Y_t = A_t$, we obtain estimates of $\mu_A = 0.016$, $a^A = 0.42$ and $\sigma_A = 0.02$.

2.2 Panel-data estimation of the MPCs

2.2.1 Econometric Implementation

Let $c_{st} \equiv \Delta \log C_{st}$ denote the growth rate of state-level consumption. In our implementation, we regress c_{st} on income growth, y_{st} , and lagged income growth, $y_{s,t-1}$, respectively. Aggregate policy and aggregate interest rates affect consumption. It is not obvious how to best capture such aggregate effects using exogenous regressors, so we follow Ostergaard, Sørensen, and Yosha (2002) and perform all regressions in terms of the deviations from the average value across states in each time period.⁶ In symbols, we regress $c_{st} - \bar{c}_{.t}$ on $y_{st} - \bar{y}_{.t}$ and $y_{s,t-1} - \bar{y}_{.,t-1}$, respectively, where $\bar{c}_{.t} = \frac{1}{50} \sum_{s=1}^{50} c_{st}$ is the time-specific mean of consumption growth and similarly for the other variables. Removing time-specific means is equivalent to including a dummy variable for each time period. Such time-specific dummy variables are referred to as time-fixed effects in the panel-data literature, and including time-fixed effects implies that we are measuring the effect of state-specific changes in income on state-specific consumption. We also want our results to be robust to permanent differences between the states. For instance, some states may have higher consumption growth due to demographic factors that are hard to control for. We, therefore, also remove state-specific averages; i.e., we use data in the form (for a generic variable x): $z_{st} = x_{st} - \bar{x}_{.t} - \bar{x}_{s.} + \bar{x}_{..}$, where $\bar{x}_{s.} = \frac{1}{T} \sum_{t=1}^T x_{st}$ is the state-specific mean of x and the last term is the overall average across states and time; this is added to keep the mean of z_{st} equal to 0. Using variables in this form is equivalent to including state-specific (and, as before, time-specific) dummy variables. In the language of panel-data econometrics, we include a state-fixed effect (also referred to as a “cross-sectional fixed effect”). We will use the shorter panel-data econometric notation and write our regressions as:

$$c_{st} = \mu_s + \nu_t + \alpha^c y_{st} + u_{st} , \tag{1}$$

where the μ_s terms symbolize the inclusion of cross-sectional fixed effects and the ν_t terms symbolize the inclusion of time-fixed effects. In the above regressions, α^c is measuring the current MPC. If X is a variable that might affect the MPC, we allow the MPC to change with

⁶Empirically, it matters little if the data is adjusted by subtracting average values of the state-level variables or if U.S.-wide aggregate values are subtracted. The method chosen here is the most straightforward in terms of implementation.

X by estimating the regression:

$$c_{st} = \mu_s + \nu_t + \alpha_{st}^c y_{st} + u_{st}, \quad (2)$$

where $\alpha_{st}^c = \alpha^c + \zeta^c (X_{st} - \bar{X}_{.t})$.

In this regression, the current MPC is $\alpha^c + \zeta^c (X_{st} - \bar{X}_{.t})$ where the time-specific average of X_{st} is subtracted in order to remove U.S.-wide aggregate effects.⁷ We subtract the time-specific average $\bar{X}_{.t}$ from the X variable so the ζ -coefficient will not pick up variations in the average (across states) MPC over time. We do not subtract the state-specific average from the X variable. The goal is to study if the MPC varies across states and, indeed, many of the “ X -variables” we utilize are constant over time and would become trivially zero if the state-specific averages were subtracted.

In our implementation, we will often include more than one interaction variable and each of them will be treated as explained above. Our regressions using lagged income are done in the exact same fashion, substituting y_{t-1} for y_t everywhere.

We use the sample period 1970–1998 for which our data is available.⁸ We approximate state-level consumption by state-level retail sales. We transform retail sales and labor income to per capita terms and deflate them using the Consumer Price Index. Details are given in Appendix A.

2.3 Regressors

We turn to the empirical estimation of the MPCs as functions of state-level variables. As interaction terms, we use variables that capture state-level persistence of income shocks and indicators of uncertainty.

Persistence of aggregate shocks. Our measure of the persistence of aggregate shocks in state s is the estimated parameter \hat{a}^s shown in Table 1, column (1).

⁷In order to estimate this model, we regress c_{st} on y_{st} , $(X_{st} - \bar{X}_{.t})(y_{st} - \bar{y}_{.t} - \bar{y}_s + \bar{y}_{..})$, and time- and state-specific dummy variables.

⁸A previous version used a sample period starting in 1976. The results are not very sensitive to the exact sample.

Aggregate state-level uncertainty. The estimated standard error of the innovation to aggregate income $\hat{\sigma}_{G_s}$, Table 1, column (2).

Individual-level income volatility. We use the share of farmers in a state. Farmers are subject to substantially higher transitory income uncertainty than other income groups as documented in table 4 in Carroll and Samwick (1997). Based on this result, we examine if the lagged MPC is lower in states where a relatively large number of consumers can be expected to have high variance of transitory idiosyncratic income. In our implementation, we use the interaction variable “farm share” (number of employed—including proprietors—in farming divided by total employment in the state).

Government sector jobs are less subject to the vagaries of nature and to the state-level business cycle implying that the share of government employees may be a good proxy for states with a low value of individual-level transitory uncertainty—see table 4 in Carroll and Samwick (1997).

Correlation matrix for regressors. Table 2 presents the correlations of our regressors in the estimations of current and lagged MPCs: the share of farmers in total employment, the share of government employment and the interactions of these employment shares, and estimated persistence and standard deviations of the income shocks interacted with income growth (non-interacted persistence and standard deviations are not included as regressors as they would be perfectly collinear with the state-fixed effects). The differences between the top and the bottom part of the table are mainly in the correlations between the regressors with current versus lagged income growth. The agriculture share “interaction” is strongly correlated with the persistence interaction as well as with the parameter for aggregate uncertainty interacted with income growth, which may render the interpretation of the variable uncertain, but the share of government interaction is not highly correlated with other regressors.

2.4 Empirical results of the panel-data estimations

We estimate the model allowing for state- and time-specific variances. More precisely, we estimate the model by feasible Generalized Least Squares (GLS) by first estimating the model using ordinary least squares (OLS) and then estimating a time-specific variance and state-specific vari-

ance for use in the second stage GLS regression. In specifications that involve the state-level persistence and uncertainty parameters, the standard errors will be biased because these parameters are measured with error (they are “generated regressors”). We calculate adjusted standard errors using a Monte Carlo method described in Appendix D.

Current MPC. Columns (1) and (4) of Table 3 show specifications where non-interacted farm and government shares are included. The non-interacted terms are not significant and subsequently dropped. These regressions simply document that the coefficients to the interaction terms do not simply pick up left-out “linear terms.” Column (2) reports the MPC out of current income from a regression of consumption growth on current income growth and current income growth interacted with the four variables that we suspect may affect the MPC. In column (3), we show the results leaving out regressors that are clearly insignificant. Because all the interaction variables have the time-specific mean subtracted, the coefficient to current income growth is the predicted value for a typical state—i.e., a state with average persistence, agricultural share, etc. We find that the MPC out of current income for a typical state is 0.33 (or 0.34 from the column (2) specification). This is clearly lower than the prediction of a PIH model with persistent income shocks such as was found for state-level income. This is, of course, the well-known “excess smoothness” result that we will focus on explaining in Section 3.

The effect of persistence is estimated robustly and clearly significant at the 1% level. The estimated value of the coefficient to persistence of state-specific income is 0.53 from column (23) and even larger from column (2). This corresponds to a large economic impact. For example, the MPC in Iowa (with persistence -0.38) is predicted to be about 0.10 while the MPC in Washington (with persistence 0.40) is predicted to be 0.51.⁹ The effect of persistence is estimated to be slightly larger when other interaction terms are included. The state-specific variance of the innovation to state-level income is significant at the 10% level. We suspect that state-level uncertainty may influence consumption patterns but we are not able to identify this with any certainty.

Lagged MPC. Table 3 also examines the same specification in terms of the MPC out of lagged income (“excess sensitivity”). For the average state, the MPC is estimated to be 0.16

⁹The number for, e.g., Iowa, is obtained as $0.33 + 0.53 \times (-0.38 - 0.07)$, where the term -0.07 corresponds to the subtraction of the average value.

in column (5) which includes all interactions. The effect of lagged income growth is robustly estimated—the coefficient is 0.17 in column (6), which includes only the significant regressors, and significant at the 1% level. Clearly, consumption displays excess sensitivity in our data. The interaction terms involving parameters of the estimated income process are clearly insignificant. This is somewhat encouraging because PIH-type models predict that persistence affects the current MPC and not the lagged MPC and the fact that the data so clearly displays this pattern indicates that our results are not likely to be spurious. We find a clear effect of the agricultural share, with more agricultural states displaying much lower sensitivity to lagged income shocks. We believe that this is due to a higher level of income uncertainty in agricultural states, a belief that is reinforced by the fact that states with large government employment display the opposite pattern. The interaction terms for both these variables have extremely large levels of statistical significance.¹⁰ We next turn to more rigorously evaluating if a combination of standard consumption models can explain these patterns.

3 Alternative Consumption Models

Our analysis will show that a combination of models is indeed required to account for the state-level MPCs that we just presented. In particular, we need forward-looking rational consumers, liquidity constraints and indivisible durables, buffer-stock saving, and risk sharing.

This section briefly presents the different models we consider. We compare a PIH model, a standard buffer-stock model, an augmented buffer-stock with durables and collateral credit constraints, and a simple model with rule-of-thumb consumers (i.e., agents who consume their income every period). We also discuss time-aggregation and risk sharing. Before focusing on the predictions of each model, in Table 5, we present the income process common to all models. Then, we briefly discuss how the models differ, as well as our calibration and aggregation procedures.

¹⁰Large coefficients of opposite sign are sometimes an indicator of multi-collinearity so we remind the reader that the government and agricultural employee shares are not strongly correlated. Also, if we drop one of these regressors the coefficients of the remaining regressors change by little.

3.1 The income process

In all models considered, disposable labor income is assumed to be exogenous to the agent, stochastic, and the only source of uncertainty. We assume the income of agent j in state s follows the model:

$$\begin{aligned} Y_{jt} &= P_{jt} V_{jt}, \\ P_{jt} &= P_{jt-1} A_t G_{st} N_{jt}. \end{aligned} \tag{3}$$

Labor income, Y_{jt} , is the product of permanent income, P_{jt} and an idiosyncratic transitory shock, V_{jt} . A_t can be thought of as growth of permanent income attributable to aggregate productivity growth in the country, while G_{st} reflects growth of permanent income specific to state s . N_{jt} is a permanent idiosyncratic shock. $\log N_{jt}$ and $\log V_{jt}$ are independent and identically normally distributed with variances σ_N^2 and σ_V^2 , and means $-\sigma_N^2/2$ and $-\sigma_V^2/2$, respectively. $\log A_t$ is assumed to be an $AR(1)$ process with persistence a^A , unconditional mean μ_A , and variance σ_A^2 . $\log G_{st}$ is also an $AR(1)$ process with persistence a^s , mean 0, and variance $\sigma_{G_s}^2$.

This income specification is useful since it allows for consumers to share in aggregate and state-specific growth while the variance of their income can be calibrated to be dominated by idiosyncratic permanent or transitory components. The formulation implies that the growth rate of individual labor income follows an ARMA process, $\Delta \log Y_{jt} = \log A_t + \log G_{st} + \log N_{jt} + \log V_{jt} - \log V_{jt-1}$, consistent with microeconomic evidence (e.g., MaCurdy 1982, Abowd and Card 1989). By the law of large numbers, aggregate income growth can be written as $\Delta \log Y_t = \log A_t$, while state-specific income growth is $\Delta \log Y_{st} - \Delta \log Y_t = \log G_{st}$.

3.2 The models

The Permanent Income Hypothesis

The PIH assumes forward-looking rational consumers. The basic set up is as follows. Consumer j maximizes the present discounted value of expected utility from consumption of a nondurable good, C . Let $\beta \leq 1$ be the discount factor and R the interest factor (gross interest rate). S_{jt} is agent j 's holding of a riskless financial asset at the end of period t . Each period,

the funds available to agent j consist of the gross return on assets RS_{jt-1} plus Y_{jt} units of labor income. The agent chooses optimal consumption C_{jt} according to the maximization problem:

$$\begin{aligned} \max_{C_{jt}} \quad & \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t U(C_{jt}) \right\} \\ \text{s.t.} \quad & S_{jt} = RS_{jt-1} + Y_{jt} - C_{jt}. \end{aligned} \quad (4)$$

The utility function is assumed to be quadratic, $\beta R = 1$, and there are no liquidity constraints.

The PIH-model provides closed-form solutions for the predicted growth in consumption as a function of innovations to income. Hansen and Sargent (1981) and Hansen, Roberds, and Sargent (1991) show that if income can be represented by the ARMA process $a(L)Y_{jt} = b(L)\varepsilon_{jt}$, the PIH predicts that:

$$\Delta C_{jt} = \frac{R-1}{R} \varepsilon_{jt} \times \frac{b\left(\frac{1}{R}\right)}{a\left(\frac{1}{R}\right)}, \quad (5)$$

where ε_{jt} is the income innovation. See Deaton (1992) for a textbook exposition. For example, if income growth is an AR(1) process with persistence a :

$$\Delta C_{jt} = \frac{R}{R-a} \varepsilon_{jt}.$$

When $a > 0$, consumption should react more than one-to-one to changes in current income. On the other hand, the predicted MPC out of lagged income is 0 (see Appendix C for more details). Furthermore, because of certainty equivalence, changes in the volatility of income (keeping the mean fixed) do not affect the MPCs. This allows us to study the aggregate implications of the model by analyzing a representative-agent model where the representative consumer receives aggregate income.

The Buffer-Stock Model

The buffer-stock model also considers forward-looking rational consumers. However, in this setup consumers are impatient ($\beta R < 1$) and prudent with a constant relative risk aversion utility function, $U(C_{jt}) = \frac{C_{jt}^{1-\rho}}{1-\rho}$. With $\rho > 0$, the agent is risk-averse and has a precautionary motive for saving.

In the literature, buffer-stock savings behavior has been derived from two different assumptions. Deaton (1991) explicitly imposes a no-borrowing constraint, $S_{jt} > 0$, but assumes agents always receive positive income. Carroll (1992), on the other hand, endogenously generates a no-borrowing constraint by assuming individuals may receive zero income (a transitory disastrous state) with a very small probability. In this case, the agent will optimally never want to borrow to avoid $U'(0) = \infty$. We follow Deaton's specification in this paper.

A closed-form solution to the model does not exist and it must be solved by computational methods (see Appendix D). It is well-known that consumption functions for a buffer-stock consumer are nonlinear, so explicit aggregation is needed to obtain implications for aggregate consumption. In this case, the MPCs depend both on persistence and uncertainty levels. However, Ludvigson and Michaelides (2001) show that an explicitly aggregated buffer-stock model cannot replicate the excess smoothness and excess sensitivity observed in U.S. aggregate data and recur to incomplete information to generate some excesses. The assumption that agents have less information about their own income than econometricians is somewhat controversial and the amount of smoothness generated by this assumption does not seem large enough to explain the very significant smoothness found in state-level data. Other mechanisms that can lower the predicted current MPC are habit formation (see Carroll 2000, Michaelides 2001), or the inclusion of indivisible durables that can only be purchased with non-negligible down payments (see Luengo-Prado 2006). Both mechanisms deliver an optimal sluggish response of consumption to changes in permanent income that can generate some excesses. We describe a buffer-stock model with durables and down payments next.

An Augmented Buffer-Stock Model with Durables and Down Payments

Our buffer-stock model with nondurables, durables and down payment requirements follows Luengo-Prado (2006). Consumer j derives utility from the consumption of a nondurable good C and the services provided by a durable good K (for simplicity we assume that durable services are proportional to the durable stock). The consumer solves:

$$\max_{\{C_{jt}, K_{jt}\}} V = E_0 \left\{ \sum_{t=0}^{\infty} \beta^t U(C_{jt}, K_{jt}) \right\},$$

$$\begin{aligned} \text{s.t. } S_{jt} &= RS_{j,t-1} + Y_{jt} - C_{jt} - (K_{jt} - \psi K_{j,t-1}) - \chi(K_{jt}, K_{j,t-1}), \\ S_{jt} &\geq -(1 - \theta)K_{jt}, \end{aligned}$$

where ψ is the depreciation factor (one minus the depreciation rate) and $\chi(K_{jt}, K_{j,t-1})$ is an adjustment cost function equal to 0 if $K_{jt} = \psi K_{j,t-1}$ and $\phi\psi K_{j,t-1}$ otherwise. This adjustment cost can be interpreted as a proportional loss in the selling price of the durable stock attributed to any type of cost incurred upon sale, such as sales commissions or imperfections in the resale market for the durable. We assume that the instantaneous utility function is separable in the two goods and is of the CRRA type, $U(C_{jt}, K_{jt}) = \frac{C_{jt}^{1-\rho}}{1-\rho} + \frac{K_{jt}^{1-\rho}}{1-\rho}$.¹¹

An important aspect of the model is the collateralized constraint imposed on the agent, $S_{jt} \geq -(1 - \theta)K_{jt}$, with $\theta \in [0, 1]$. This implies that an individual's borrowing limit is a fraction $(1 - \theta)$ of the durable stock. The constraint summarizes several commonly observed aspects of collateral lending. A household can finance a fraction $(1 - \theta)$ of durable purchases but must provide a fraction θ of the durable's value as a down payment. On the other hand, a household owning a durable good can obtain a durable-equity loan with a maximum loan-to-value ratio $(1 - \theta)$. This wealth constraint alters the allocation of resources between the durable and the nondurable goods and has implications for the volatility of the consumption of the two goods.

Luengo-Prado (2006) shows that in this framework, nondurable consumption becomes smoother relative to income as down payment requirements increase for two different reasons. First, when income is transitorily low, a buffer-stock consumer on occasion liquidates the equity accumulated in the durable to prop up nondurable consumption. Since higher required down payments translate into higher levels of equity, nondurable consumption becomes smoother. Second, when an individual experiences a positive permanent income shock, he or she chooses not to fully adjust consumption due to the desire to spread out the cost of accumulating a down payment. Thus, the sluggish response of durable and nondurable consumption to changes in permanent income can generate robust excesses at the aggregate level for reasonable parameter values. In this model, *ceteris paribus*, the higher θ relative to the user cost of the durable, $R - \psi/R$, the higher the excesses.

¹¹We follow Bernanke (1984) who studies the joint consumption of durable and nondurable goods and finds that separability is a good approximation. With regards to prices, we assume that $P_t^C/P_t^K = 1, \forall t$.

As for the standard buffer-stock model, there is no closed-form solution (see Appendix D for a brief description of our computational method) and explicit aggregation is needed to obtain implications for aggregate consumption. The MPCs depend on both persistence and uncertainty levels in this model also.

Rule-of-Thumb Consumers

In this setting, agents consume their income every period ($C_{jt} = Y_{jt}$) and this model does not generate excess smoothness since $\Delta C_{jt} = \Delta Y_{jt}$. The rule-of-thumb model (ROT) generates excess sensitivity only if income growth is persistent. As in the PIH case, the aggregate implications of this model can be calculated by assuming that a representative agent receives the aggregate income process.

Time-Aggregation

We anticipate that neither the PIH model nor the standard buffer-stock model generate robust excess sensitivity at the state-level. In order to generate additional sensitivity, we take into account that consumers' decision intervals and data-sampling intervals may be different, i.e., we allow for temporal aggregation. Time-aggregation generates excess sensitivity in a representative PIH model (see Working 1960). Appendix C spells out the details. Time-aggregation also delivers additional excess sensitivity in buffer-stock models.

Risk Sharing

We find that the sensitivity of consumption to current income shocks is much lower than can be explained by the models considered thus far. A potential explanation could be that innovations to measured disposable income are not a correct measure of the changes in consumers' actual budget constraints.

The predicted change in consumption in the PIH-framework, see equation (5), is calculated using two features of the model. The first is the celebrated martingale result of Hall (1978) and

the second is simply the inter-temporal budget constraint:

$$\sum \left(\frac{1}{R}\right)^k C_{t+k} = S_t + \sum \left(\frac{1}{R}\right)^k Y_{t+k}, \quad (6)$$

where S_t measures asset holdings at the beginning of period t . An income shock ε_t gives rise to a change in the right-hand side of the budget constraint that affects consumption by the factor shown in equation (5). Consider the situation where risk sharing takes the form of a “transfer” τ_t to the representative agent which is *not part of measured disposable income* paid out of U.S. aggregate consumption. Such transfers will render equation (6) invalid for the purpose of determining consumption. We will outline the argument further below—a more detailed discussion of this issue can be found in Attanasio and Pavoni (2006).

In the real world, payments of dividends, interest, and rent are not fixed interest rate returns and the stochastic returns potentially provide substantial smoothing of the labor income of U.S. states. The variable τ will capture such deviations from the model. Stochastic returns could possibly be measured or imputed but τ also is likely to capture a host of other variables that “break” the budget constraint based on measured disposable income, such as unmeasured capital gains, financial help from family and friends, misreported income, bankruptcy, etc. While many potential unmeasured transfers come to mind, it is hard to know their quantitative importance. However, examining the degree of excess smoothness can help assessing this, as pointed out by Attanasio and Pavoni (2006). Recall that our empirical work includes time-fixed effects so any aggregate numbers will not affect our results and, consequently, τ_t , as well as ε_t , should—for the purpose of this paper—be interpreted as state specific variables that average to 0.

Assume that consumption is a martingale but the budget constraint is:

$$\sum \left(\frac{1}{R}\right)^k C_{t+k} = S_t + \sum \left(\frac{1}{R}\right)^k (Y_{t+k} + \tau_{t+k}). \quad (7)$$

If there is full risk sharing, the variable $Y_{t+k} + \tau_{t+k}$ will simply be zero (when the variables are measured as deviations from aggregate variables). However, perfect risk sharing is not the optimal outcome in a situation where moral hazard makes agents likely to supply less (or no) effort under full risk sharing. Abraham and Pavoni (2005) develop a model where agents’ effort is unobserved and agents have access to a bond market with a fixed interest rate but where agents’ asset holdings are unobserved. In such a setting, partial risk sharing will be a constrained

optimum under certain parametric assumptions. The intuition is that it has to provide positive utility for agents to provide effort even if part of the fruits from effort will be shared due to risk sharing. However, as long as agents have access to a bond market the Euler equation for consumption will hold and, in particular for the quadratic utility case, consumption will be a martingale if the interest rate equals the subjective time discount rate.¹²

If the temporal behavior of $Y + \tau$ can be described by the ARMA-process $\alpha(L)(Y_t + \tau_t) = +\beta(L)v_t$, where v_t is the innovation to after risk sharing income, then predicted consumption will satisfy:

$$\Delta C_t = \frac{R-1}{R} \times \frac{\beta(\frac{1}{R})}{\alpha(\frac{1}{R})} v_t,$$

if utility is quadratic and the interest rate equals the time discount rate. In other words, the PIH still holds but post-risk sharing income replaces income. In a regression of ΔC_t on current income of the form:

$$\Delta C_t = \mu_t + \gamma \Delta Y_t + u_t, \tag{8}$$

the right hand side variable can be interpreted as $Y_t + \tau_t$ plus a measurement error ($-\tau_t$). If the measurement error is uncorrelated with the true regressor, it is well-known downward bias will occur. If one is willing to make strong parametric and functional form restrictions, one can obtain a stronger prediction. We, again, refer the reader to Attanasio and Pavoni (2006) for details and note that they derive a model with moral hazard due to unobserved effort where $Y_t + \tau_t = \omega y_t$ for a constant $\omega < 1$. One can consider $1 - \omega$ the amount of partial risk sharing and the value of ω can be deduced from the observation that in regression (8), the estimated coefficient γ will be $\omega \times \varrho$ where ϱ is the coefficient one would obtain in the “correct” PIH-regression:

$$\Delta C_t = \mu_t + \varrho \Delta(\omega Y_t) + u_t. \tag{9}$$

In other words, if there is 50% partial risk sharing the estimated coefficient in a regression of consumption on measured disposable income will be 50% lower than predicted by a model that ignores risk sharing. The coefficient to lagged income will similarly be 50% lower. Note, that ω measures the amount of consumption insurance “starting from” measured disposable labor

¹²The assumption that asset holdings are unobserved is necessary for this outcome to be optimal. If asset holdings were observed a planner could control those and the properties of the model would change.

income. If state-level output is the relevant state-level endowment shock, a lot of risk sharing can take place between financial institutions, firms, and governments and ω would be a lower bound for “total” risk sharing. Asdrubali, Sørensen, and Yosha (1996) measure state-level overall inter-state risk sharing but do not attempt to separate the amount of consumption smoothing taking place through optimal savings behavior (“self-insurance”) from the amount of partial (unobserved) risk sharing measured by ω .

For the interpretation of our results we make approximations. We assume, as is typical in the PIH-related literature, that the relations derived for the change in variables approximately hold for the change in log-transformed variables. Also, we will assume that the considerations in this section are approximately valid in the setting where the consumer may face credit constraints.

3.3 Calibration and aggregation procedures

Calibration

A good calibration of the income process is essential to obtain qualitative and quantitative predictions. To calibrate the aggregate and state-level income shocks, we use our estimates of Section 2.1. As the baseline state-level income shock, we use the simple average across states of the estimated AR(1) parameters in Table 1. Idiosyncratic income shocks are taken from previous studies—see, for example, Carroll and Samwick (1997) and Gourinchas and Parker (2002). In particular, we set $\sigma_V = 0.07$ and $\sigma_N = 0.05$ in the benchmark calibration.

Regarding other parameters, the interest rate is set to 2%. In the buffer-stock models risk aversion is $\rho = 2$ and the discount rate is 5%, both standard in the literature. For the buffer-stock model with durables, we follow Luengo-Prado (2006) and set the adjustment cost parameter, ϕ , equal to 5%, a typical realtor’s commission. ψ , the depreciation factor, is set to 0.915, implying an annual depreciation rate of 8.5%.¹³ The down payment parameter θ is 0.3.¹⁴

¹³We obtain this number by combining data from the *National Income and Product Accounts* and the *Fixed Assets and Consumer Durable Goods Accounts* from the Bureau of Economic Analysis for the years 1959–2001. We interpret durables, K , in a comprehensive manner as the sum of residential stocks and all consumer durable goods. Accordingly, investment in durables, I , is calculated as expenditure on consumer durables plus residential private domestic investment. We assume the U.S. is in steady state and calculate the real, average ratio of investment on durables to the durable stock, which determines the depreciation rate: $1 - \psi = I/K$.

¹⁴According to the Federal Housing Finance Board, the average down payment for a house for the period 1963–2001 was 25 percent. We choose the slightly higher number because we include durables other than houses in our

Aggregation procedure

While it is possible to work with a representative-agent model for the PIH and the ROT model, it is well-known that consumption functions for a buffer-stock consumer are nonlinear, so explicit aggregation is needed to obtain implications for aggregate consumption. Our simulation exercise is similar in spirit to that of Ludvigson and Michaelides (2001), who calibrate their income process to match U.S. aggregate income and focus on explaining excess sensitivity and excess smoothness at the aggregate level. Our goal, however, is to assess how income persistence and income uncertainty impact the MPCs using state-level data, so we proceed in a different manner.

Ideally, we would like our simulation exercise to be as close as possible to the empirical strategy in Section 2.2. Briefly, we would like to simulate 50 states with different persistence and uncertainty parameters and a common aggregate productivity shock. Then, we would run panel-data regressions with both state-fixed effects and time-fixed effects. The inclusion of time-fixed effects is important because this removes the first-order impact of the more persistent aggregate (U.S.-wide) shock—due to the non-linearity of the model, the results are not, however, identical to simulating the model without any U.S. wide component. Due to computational limitations, we simulate “states” with state-specific shocks generated from a common distribution. In other words, our simulated states are ex-ante identical but ex-post different because they are subject to different shocks.¹⁵ We calculate marginal effects on the MPCs of changes in persistence, and changes in transitory and permanent uncertainty by changing the parameters of our baseline calibration (for all states) one at a time.

We simulate income paths for 30,000 individuals in 10 states—3,000 per state—for a number of periods. All individuals share a common aggregate shock each period and individuals living in the same state share state-specific shocks. Using the optimal consumption functions and the simulated income paths, we calculate state-level consumption, and state-level income— C_{st} , Y_{st} —as the average of individual consumption and income over all consumers living in “state”

interpretation of K .

¹⁵Thus, state-fixed effects are not necessary in the regressions with simulated data but are included for comparability with the regressions using actual data.

s. Then, we run the following panel regressions:

$$\Delta \log C_{st} = \mu_s + \nu_t + \alpha^c \Delta \log Y_{st} + u_{st}, \quad (10)$$

$$\Delta \log C_{st} = \mu_s + \nu_t + \alpha^l \Delta \log Y_{s,t-1} + u_{st}. \quad (11)$$

μ_s are state-fixed effects and ν_t are time-fixed effects that control for aggregate effects. Thus, $\hat{\alpha}^c$ is the estimated MPC out of current state-specific income shocks, and similarly $\hat{\alpha}^l$ is the estimated MPC out of lagged state-specific income shocks—for brevity, we refer to them as the current/lagged MPCs. We repeat this process 20 times and report, in Table 5, the average current and lagged MPCs across the 20 independent samples. We report the average of the estimated standard error (for each parameter) across the 20 samples in parentheses.¹⁶

Time-aggregation

In some simulations, we take into account that consumers’ decision intervals and data-sampling intervals may be different and allow for temporal aggregation. In particular, we assume that while agents make decisions on a bi-annual basis, we only observe annual data. In our regressions with simulated data, state-level consumption in year t is simply $C_{st} = C_{st}^1 + C_{st}^2$, where C_{st}^1 and C_{st}^2 are calculated as the average of individual consumption in state s for the first and second half of the year, respectively, and analogously for income. All relevant parameters are adjusted to the bi-annual frequency for these simulations. Table 4 compares the benchmark parameter values at both frequencies. Details on how to adjust the income parameter values are presented in Appendix E.

3.4 Models’ predictions regarding the MPCs

Table 5 presents results comparing the predicted state-level MPCs for the different models in our benchmark calibration. While for the PIH and the ROT models the MPCs are calculated analytically, for the buffer-stock model the MPCs are calculated through explicit aggregation

¹⁶Adding more samples to the simulations does not change the results significantly. For example, for the augmented buffer-stock model and 100 samples, the current and lagged MPCs are 0.758 (0.01) and 0.357 (0.033), respectively. With 20 samples, the numbers are 0.756 (0.01) and 0.358 (0.033), as shown in Table 5.

via simulations. Estimated standard errors are given in parentheses.¹⁷

We first focus on the results without time-aggregation or risk sharing. Given that the state-specific permanent shock is persistent in our benchmark calibration ($a^s = 0.07$), the PIH predicts a current (state-level) MPC higher than 1, roughly 1.07. The lagged MPC is 0. The ROT model predicts a current MPC of 1, since in this case the change in consumption is simply the change in income. In the buffer-stock model, agents cannot borrow and even though they have some assets because of prudence, asset holdings are small due to impatience. Hence, individuals cannot increase consumption as much as PIH consumers would, when facing a persistent positive permanent shock, resulting in a lower current MPC and a higher lagged MPC. For our baseline case, the current and lagged MPCs in the buffer-stock model are 1.0 and 0.04 respectively. In the buffer-stock model with durables, the current and lagged MPCs are 0.79 and 0.1, respectively. With a non-trivial down payment (30%), agents choose not to adjust consumption levels immediately when facing permanent income shocks, preferring to spread out the accumulation of required wealth holdings. As a result the current MPC is lower than in a standard buffer-stock model and the lagged MPC is higher.

All these numbers are far from their empirical counterparts: the estimated current MPC is 0.33, while the estimated lagged MPC is 0.17. Allowing for time-aggregation increases the predicted lagged MPC considerably and decreases the current MPC slightly (except for the ROT model). For the PIH, the current MPC is 1.04 and the lagged MPC is 0.15. For the buffer-stock model with durables, the current MPC goes down to 0.76, while the lagged MPC becomes 0.36. One might generate higher smoothness and more sensitivity by increasing the required down payment significantly but down payments of, say, 50% or more do not appear to be common. Adding 50% risk sharing to this model delivers a current MPC of 0.38 and a lagged MPC of 0.18, much closer to their empirical counterparts. The assumption of partial risk sharing of the simple form applied here, and by Attanasio and Pavoni (2006), allows us to fit the smoothness of consumption in any of the models considered—for example, we could assume the PIH model with time-aggregation and about 70 percent risk sharing. However, allowing for such a large amount of risk sharing would lead us to predict virtually no excess sensitivity.

¹⁷These are the average estimated standard errors of $\hat{\alpha}^c$ and $\hat{\alpha}^l$ in regressions (10) and (11), respectively, across the 20 independent samples.

Both time-aggregation and risk sharing are needed to obtain MPCs close to their empirical counterparts. Therefore, we focus the remainder of our analysis on models that include these features. In particular, we concentrate on the PIH and the augmented buffer-stock model. Table 6 presents results regarding the effects of persistence and uncertainty on the MPCs in these two models.

Decreasing persistence to -0.1 lowers the MPCs out of current and lagged income in the augmented buffer-stock model (from 0.38 to 0.35 and from 0.18 to 0.32 respectively). In the PIH, the current MPC declines from 0.52 to 0.4, while the lagged MPC increases slightly from 0.07 to 0.08. We can calculate the marginal effect of persistence as the change in the MPC relative to the benchmark case divided by the change in the persistence parameter. Table 6 shows that the marginal effect of increasing persistence on the current MPC is 0.72 in the PIH-model, quite large, and 0.15 in the buffer-stock model. The marginal effect of persistence on the lagged MPC is -0.06 in the PIH-model and 0.29 in the buffer-stock model.

Next, we examine the marginal effects of uncertainty by changing the standard deviation of the different income shocks one at a time. Contrary to the PIH case, these changes have large effects on the MPCs in the augmented buffer-stock model. We start with the idiosyncratic shocks by reducing their standard deviations by half (one at a time). Because of less uncertainty, agents save less, which might be expected to lead to higher current MPCs. However, with less savings and liquidity constraints, agents cannot increase consumption as much in response to a persistent positive permanent shock. This effect tends to lower the MPCs out of current income and offsets the former effect in the case of transitory uncertainty and dominates in the case of permanent uncertainty (in our simulations, the marginal effect of transitory uncertainty on the current MPC is 0.06, while the marginal effect of permanent uncertainty is 0.76). Note that the lagged MPCs are higher because with lower savings agents are liquidity constrained more often. The marginal effects in Table 6 look quite large: -0.5 for transitory uncertainty and -0.8 for permanent uncertainty. However, a unit increase in the standard deviation of any of the income shocks corresponds to quite a massive increase in uncertainty.

Finally, more state-level aggregate uncertainty is introduced by changing the standard deviation of the state-level shocks one at a time.¹⁸ More state-level aggregate permanent uncertainty

¹⁸We *increase* the standard deviation in this case. The direction of the changes are chosen such that the model

results in a higher current MPC in these simulations. Because agents hold more assets due to higher uncertainty, they can adjust consumption more promptly in response to persistent positive permanent income shocks, which increases the current MPC. Also, consumers are constrained less often and the lagged MPC decreases slightly.

The clear finding that states with more persistent income shocks have higher sensitivity to income provides strong support to forward-looking models such as the PIH and precautionary savings models. Instrumental in bringing the predicted smoothness to the level observed is the combination of credit constraints and durable goods together with partial risk sharing. Time-aggregation will then explain the observed sensitivity to lagged income while 50% risk sharing is not so large as to drive the predicted sensitivity to zero. The impatience that creates buffer-stock behavior is necessary to explain why states where agents are likely to face significant individual level uncertainty display less sensitivity.

A model that roughly fits the behavior of U.S. state-level consumption (when adjusted for aggregate components) is a model with time-aggregation and 50% partial risk sharing where half the population behaves like PIH-consumers while the other half faces credit constraints, is impatient, and demands lumpy durable goods. The predicted MPCs and marginal impacts of our regressors are displayed in the right-most third of Table 6. The predicted current MPC is at the high end of a 95% confidence interval for the estimated current MPC in Table 3 and the marginal effect on the current MPC of persistence is somewhat lower (by one standard deviation) than the coefficient to the persistence interaction found in Table 3. Nonetheless, the overall impression is that this model does an impressive job of matching the features of the data.

4 Conclusions

Using panel-data regressions, we document that state-level consumption displays strong excess smoothness and sensitivity. Also, the persistence of state-level income affects the current marginal propensity to consume significantly, with states with higher persistence exhibiting higher MPCs. Further, we observe that excess sensitivity is higher in states where consumers are likely to face higher income uncertainty.

satisfies the convergence condition of footnote 21.

We show that a combination of models is required to simultaneously explain these features. While the effect of income persistence on the current MPC is consistent with a standard forward-looking model such as the PIH, this setup alone cannot account for the magnitude of excess smoothness and sensitivity observed in state-level data, nor the dependency of the MPCs on uncertainty indicators. A simple explanation of excess sensitivity may be that consumers make decisions at a higher frequency than the data available to econometricians. Excess smoothness can be explained by the fact that innovations to measured disposable income may not be the right measure of changes in a consumer's actual budget constraint in the presence of unobserved risk sharing. Further, buffer-stock behavior is consistent with the observed higher excess sensitivity in agricultural states. A suitably calibrated model that combines these features can fit the data fairly well.

Our work in this article put together a model in order to fit the data. In order to actually test this model, one would need to estimate it on independent data from, say, another country. Such testing would be an interesting topic for future work.

Appendices

A The Data

We use state-level annual data from a variety of sources. We construct state-level disposable labor income for the period 1964–1998 using data from the Bureau of Economic Analysis (BEA). We define labor income as personal income minus dividends, interest, and rent, and social security contributions. We calculate after-tax labor income by multiplying labor income by one minus the tax rate, where we approximate the tax rate by total personal taxes divided by personal income for each state in each year. We refer to the resulting series as disposable labor income or—for brevity—just as labor income or income. The panel regressions in Section 2.2 use a shorter sample, 1970–1998, due to lack of availability of other variables prior to 1970. However, in order to obtain more precise parameter estimates we use the larger sample in the income estimations of Section 2.1. We also, for robustness, used the BEA disposable personal income data by state and found qualitatively similar results.

We perform state-by-state Augmented Dickey-Fuller (ADF) tests for unit roots in labor income. These tests reject the unit root null hypothesis for only a few states at conventional levels of significance. ADF tests provide somewhat weak evidence because they have low power for samples as short as ours. The overall impression is, nevertheless, that U.S. state-level labor income is well-described as an integrated process.¹⁹ We, therefore, treat labor income growth as a stationary series.

We approximate state-level consumption by state-level retail sales published in the Survey of Buying power, in *Sales Management* (after 1976, *Sales and Marketing Management*). Retail sales are a somewhat noisy proxy for state-level private consumption but no better data seems to exist. The retail sales data is available from 1963–1998. The correlation between annual growth rates of aggregate U.S. total (nondurable) retail sales and aggregate U.S. total (nondurable and services) private consumption from the National Income and Product Account, both measured in real terms and per capita, is 0.84 (0.65). We transform the retail sales and labor income series to per capita terms using population data from the BEA and deflate them using the Consumer Price Index from the Bureau of Labor and Statistics (BLS).

B Estimation of Standard Errors

Because the parameters for the state-level income processes are estimated in an initial regression they are random variables. This “generated regressors” problem leads to bias in the standard errors reported by OLS. We, therefore, use a “parametric bootstrap” procedure to calculate standard errors for all the coefficients.

Our approach is as follows. We regress consumption growth on the non-generated regressors X_{st} (including fixed effects) and the estimated regressors Y_s (which do not vary over time) using OLS (after weighting the variables with state- and time-specific estimated standard errors):

$$c_{st} = X_{st}\gamma + Y_s\delta + e_s ,$$

where $s = 1, \dots, 50$ is an index of the states, $t = 1, \dots, T$ is an index of time, and γ and δ

¹⁹Ostergaard, Sørensen, and Yosha (2002) show that panel unit root tests for disposable income—when aggregate income is subtracted—provide little evidence against the unit root hypothesis. Disposable income is highly correlated with labor income state-by-state and the results using labor income are similar.

are OLS-coefficients. From this regression, we retrieve the estimated values $\hat{\gamma}$ and $\hat{\delta}$ and the estimated standard error s_e of the residuals e_s . We proceed to estimating the standard errors of $\hat{\gamma}$ and $\hat{\delta}$ from the following Monte Carlo experiment. In each iteration l we draw from an i.i.d. $N(0, s_e)$ distribution a vector of variables, $e_{st}^{(l)}$ ($s = 1, \dots, 50; t = 1, \dots, T$). We generate the variable

$$c_{st}^{(l)} = X_{st}\hat{\gamma} + Y_s\hat{\delta} + e_{st}^{(l)} .$$

Then, for each state s and time period t , we generate $Y_s^{(l)}$ by drawing from an $N(Y_s, \Sigma_{Y_s})$ distribution where Σ_{Y_s} is the variance matrix with the estimated standard errors of Y_s reported in Table 1 in the diagonal (for example, 0.017 for persistence in Alabama).

We then perform a panel-data regression of $c_{st}^{(l)}$ on X_{st} and $Y_s^{(l)}$ and record the estimated coefficients $\hat{\gamma}^{(l)}$ and $\hat{\delta}^{(l)}$. We repeat this for $l = 1, \dots, 25000$ and then calculate the standard errors of $\hat{\gamma}^{(l)}$ and $\hat{\delta}^{(l)}$. These are the standard errors reported in the tables.

C The MPC in the PIH Case

This appendix describes how to calculate the approximate current and lagged MPCs for the PIH-model. First, we show how consumption reacts to income innovations and then we discuss time-aggregation issues.

CONSUMPTION GROWTH AND INNOVATIONS TO INCOME

Given our assumptions about the income process, state-specific income growth is $\Delta(\log Y_{st} - \log Y_t) = \log G_{st}$. For notational simplification, let $g_t = \log G_{st}$ and $\log Y_{st} - \log Y_t = y_t$. By assumption, $g_t = a^s g_{t-1} + \varepsilon_t$. We drop the superscript in s hereafter. Thus, $\Delta y_t = g_t = a\Delta y_{t-1} + \varepsilon_t$.

Because the PIH is formulated in levels rather than logs, we must assume the process for the first difference of state-specific income can be approximately described by the process for the log difference: $\Delta(\log Y_{st} - \log Y_t) \simeq \Delta(Y_{st} - Y_t)$. Hansen and Sargent (1981) and Deaton (1992) show that if income can be represented by the ARMA process $a(L)y_t = b(L)\varepsilon_t$, the PIH predicts that: $\Delta C_t = \frac{R-1}{R}\varepsilon_t \times b\left(\frac{1}{R}\right)/a\left(\frac{1}{R}\right)$.

The ARMA representation for our state-specific income process is $(1 - aL)(1 - L)y_t = \varepsilon_t$. Thus, the consumption change in period t can be written as:

$$\Delta C_t = \frac{R}{R - a} \varepsilon_t = H \varepsilon_t. \quad (12)$$

H equals 1 when $a = 0$ and increases with persistence. We can calculate the MPC out of current and lagged state-specific income as $\frac{\text{cov}(\Delta C_t, \Delta y_t)}{\text{var}(\Delta y_t)}$ and $\frac{\text{cov}(\Delta C_t, \Delta y_{t-1})}{\text{var}(\Delta y_{t-1})}$ respectively. In this case, $\text{var}(\Delta y_t) = \sigma_\varepsilon^2 / (1 - a^2)$, $\text{cov}(\Delta C_t, \Delta y_t) = H \sigma_\varepsilon^2$, $\text{cov}(\Delta C_t, \Delta y_{t-1}) = 0$. Thus, the current MPC is equal to $H \times (1 - a^2)$ and the lagged MPC is 0.

TIME-AGGREGATION

We assume that agents make consumption decisions bi-annually but we only observe data at annual frequencies (i.e., we observe $C_t = C_t^1 + C_t^2$, where C_t^i is consumption in the i^{th} half of year t , $i = 1, 2$).

Using equation (12)—valid for the relevant frequency in the agent’s optimization problem—we can derive an expression for C_t^1 and C_t^2 :

$$\begin{aligned} C_t^1 &= C_{t-1}^2 + H \varepsilon_t^1 = C_{t-1}^1 + H(\varepsilon_t^1 + \varepsilon_{t-1}^2), \\ C_t^2 &= C_t^1 + H \varepsilon_t^2 = C_{t-1}^2 + H(\varepsilon_t^2 + \varepsilon_t^1), \end{aligned} \quad (13)$$

where the last part follows from recursive backward substitution. Thus, the annual consumption change in period t is simply:

$$\Delta C_t = (C_t^1 + C_t^2 - C_{t-1}^1 - C_{t-1}^2) = H(\varepsilon_t^2 + 2\varepsilon_t^1 + \varepsilon_{t-1}^2). \quad (14)$$

Regarding income, we assume that the bi-annual income processes are AR(1), so $\Delta y_t^2 = a \Delta y_t^1 + \varepsilon_t^2$ and $\Delta y_t^1 = a \Delta y_{t-1}^2 + \varepsilon_t^1$. Annual income is then $y_t = y_t^2 + y_t^1$. Thus, $\Delta y_t = y_t^2 + y_t^1 - y_{t-1}^2 - y_{t-1}^1$ and $\Delta y_t = y_t^2 - y_{t-1}^2 + 2y_{t-1}^1 - 2y_{t-1}^2 + y_{t-1}^2 - y_{t-1}^1 = \Delta y_t^2 + 2\Delta y_t^1 + \Delta y_{t-1}^2$ and using recursive backwards substitution, we obtain:

$$\begin{aligned} \Delta y_t &= \Delta y_t^2 + 2\Delta y_t^1 + \Delta y_{t-1}^2 \\ &= (a^2 \Delta y_{t-1}^2 + \varepsilon_t^2 + a \varepsilon_t^1) + 2(a^2 \Delta y_{t-1}^1 + \varepsilon_t^1 + a \varepsilon_{t-1}^2) + (a^2 \Delta y_{t-2}^2 + \varepsilon_{t-1}^2 + a \varepsilon_{t-1}^1) \\ &= a^2 \Delta y_{t-1} + [\varepsilon_t^2 + (2 + a) \varepsilon_t^1 + (1 + 2a) \varepsilon_{t-1}^2 + a \varepsilon_{t-1}^1]. \end{aligned}$$

We can calculate the MPC out of current and lagged state-specific income as $\frac{cov(\Delta C_t, \Delta y_t)}{var(\Delta y_t)}$ and $\frac{cov(\Delta C_t, \Delta y_{t-1})}{var(\Delta y_{t-1})}$ respectively. Using equations (14) and (15), it is easy to show that:

$$var(\Delta y_t) = \frac{6+8a+6a^2}{1-a^2} \sigma_\varepsilon^2,$$

$$cov(\Delta C_t, \Delta y_t) = H(6 + 4a)\sigma_\varepsilon^2,$$

$$cov(\Delta C_t, \Delta y_{t-1}) = H\sigma_\varepsilon^2.$$

For $a = 0$, the MPC out of current income is 1 and the MPC out of lagged income is 1/6. Therefore, time-aggregation produces robust excess sensitivity in the PIH model (see Working 1960).

Furthermore, both $cov(\Delta C_t, \Delta y_t)$ and $var(\Delta y_t)$ increase with a . For small values of a , the covariance increases faster and the current MPC increases in a initially, but could eventually decrease.

D Solution Methods for the Buffer-Stock Models

The standard buffer-stock model

A closed-form solution of the model does not exist and it must be solved by computational methods. Following Deaton (1991), the model is first reformulated in terms of cash-on-hand, $X_{jt} \equiv RS_{jt-1} + Y_{jt}$.²⁰ Given the homogeneity property of the utility function, all variables can be normalized by permanent income to deal with non-stationarity, as proposed by Carroll (1997). The first order condition of the problem becomes:

$$U'(c_{jt}) = \max\{U'(x_{jt}), \beta RE_t[(A_{t+1}G_{s,t+1}N_{j,t+1})^{-\rho}U'(c_{j,t+1})]\}, \quad (15)$$

where $c_{jt} = C_{jt}/P_{jt}$ and $x_{j,t+1} = (A_{t+1}G_{s,t+1}N_{j,t+1})^{-1}R(x_{jt} - c_{jt}) + V_{j,t+1}$.²¹ Individuals distinguish aggregate from state-specific shocks and optimize accordingly. We use Euler equation

²⁰The budget constraint becomes $S_{jt} = X_{jt} - C_{jt}$ and the liquidity constraint $C_{jt} \leq X_{jt}$. Combining the definition of cash-on-hand and the budget constraint, we can write an expression for the evolution of cash-on-hand: $X_{j,t+1} = R(X_{jt} - C_{jt}) + Y_{j,t+1}$.

²¹A necessary condition for the individual Euler equation to define a contraction mapping is $\beta RE_t[(A_{t+1}G_{s,t+1}N_{j,t+1})^{-\rho}] < 1$. This is the ‘‘impatience’’ condition common to buffer-stock models which guarantees that borrowing is part of the unconstrained plan.

iteration to solve Equation (15) numerically. x is discretized and the income shocks are approximated by discrete Markov processes following Tauchen (1986). We use 5 points for N , V , A and G . Interpolation is used between points in the x grid. The numerical technique delivers a consumption function $c(x, A, G)$: normalized consumption as a function of normalized cash-on-hand and the aggregate and state-specific permanent states. In other words, our optimal policy function for consumption has 25 branches, one for each (A, G) combination of the discrete approximations of the persistent permanent income shocks.²²

The augmented buffer-stock model

In this case the problem is solved using a finite state approximation method. The technique consists of specifying a finite-state problem that approximates the continuous one we want to solve. As with the previous technique, all relevant variables are normalized by permanent income (lower case notation). Also, the problem is reformulated in terms of a variable that we call voluntary equity, $q_{jt} \equiv s_{jt} + (1 - \theta)k_{jt}$, the equity held in excess of the required down payment. Unlike s_{jt} with a lower limit which depends on the value of the durable, voluntary equity has lower limit of 0 (independent of the value of the durable), which greatly simplifies computation. From the equation for the evolution of assets, we can work out an equation for the evolution of normalized voluntary equity: $q_{jt} = (A_t G_{st} N_{jt})^{-1} \{ Rq_{j,t-1} + [\psi(1 - d\phi) - R(1 - \theta)]k_{j,t-1} \} - \theta k_{jt} + v_{jt} - c_{jt}$, where d is 1 if the household changes the durable stock and 0 otherwise. Next, using the homogeneity of degree $(1 - \rho)$ of the utility function, we can write the Bellman equation of the problem as:

$$\begin{aligned}
V(q_{j,t-1}, k_{j,t-1}, A_t, G_{st}) = & \\
& \mathbb{E}_{t-1} \left\{ (A_t G_{st} N_{jt})^{1-\rho} \max_{q_{jt}, k_{jt}: q_{jt} \geq 0} U \left[(A_t G_{st} N_{jt})^{-1} \left\{ Rq_{j,t-1} + [\psi(1 - d\phi) - R(1 - \theta)]k_{j,t-1} \right\} \right. \right. \\
& \left. \left. - \theta k_{jt} + V_{jt} - q_{jt}, k_{jt} \right] + \beta V(q_{jt}, k_{jt}, A_{t+1}, G_{s,t+1}) \right\}. \tag{16}
\end{aligned}$$

We replace the continuous state variables, k and q , with the finite sets, $\mathcal{K} = \{k_1, \dots, k_{N_k}\}$ and $\mathcal{Q} = \{q_1, \dots, q_{N_q}\}$. Note that the problem has been conveniently formulated in such a way

²²More details on how to solve this equation can be found, for example, in the appendix of Ludvigson and Michaelides (2001).

that the control variables are next period's states. The liquidity constraint is implemented by setting $q_1 = 0$ and $q_i > 0, \forall q_i \in \mathcal{Q}, i > 1$. To deal with adjustment cost, we set: $d = 0$, if $|k_{jt} - (A_t G_{st} N_{jt})^{-1} \psi k_{j,t-1}| \leq \kappa$, and $d = 1$, if $|k_{jt} - (A_t G_{st} N_{jt})^{-1} \psi k_{j,t-1}| > \kappa$, where $\kappa = (k_n - k_l)/(N_k - 1)$. The precision of our solution increases as κ falls.

As with the previous technique, all components of the income process are discretized, using 5 points for each income shock. We use 300 points for the q and k grids. While the lower bounds for the grids are set at 0, the upper bounds are determined by trial and error. We use value function iteration, which is sped up with an acceleration technique, modified policy function iteration with S states as described in Judd (1997) to solve the discretized problem. We refer the interested reader to Luengo-Prado (2006) for more details.

E Time-Aggregation and Income Growth

In this appendix, we discuss how the calibration of income needs to be adjusted when going from annual to the bi-annual frequencies. For simplicity of exposition, we first show how to deal with permanent shocks and then we consider transitory shocks.

Permanent shocks

Let us assume that there are no transitory shocks to income and there is only one permanent shock for simplicity. In particular, income in the second half of period t equals permanent income, $Y_t^2 = P_t^2$. Permanent income in turn is $P_t^2 = G_t^2 P_t^1$, with $\log G_t^1 = a \log G_t^2 + \varepsilon_t^2$. Let $\log G \equiv g$, and $\log Y \equiv y$. Thus, bi-annual income growth is an AR(1) process: $\Delta y_t^2 = a \Delta y_t^1 + \varepsilon_t^2$. Let us work out which this implies for annual income growth.

Annual income in t is defined as the sum of income in both halves of year t :

$$\begin{aligned} Y_t &= Y_t^1 + Y_t^2 = P_t^1 + e^{g_t^2} P_t^1 \\ &= P_t^1 (1 + e^{g_t^2}) \\ &= P_{t-1}^1 e^{g_t^1 + g_{t-1}^2} (1 + e^{g_t^2}), \end{aligned}$$

where the third line follows from recursive backward substitution. Annual income in $t - 1$ is:

$$Y_{t-1} = Y_{t-1}^1 + Y_{t-1}^2 = P_{t-1}^1(1 + e^{g_{t-1}^2})$$

Then, income growth in period t is:

$$\begin{aligned} \Delta \log Y_t \equiv \Delta y_t &= g_{t-1}^2 + g_t^1 + \log(1 + e^{g_t^2}) - \log(1 + e^{g_{t-1}^2}) \\ &\simeq g_{t-1}^2 + g_t^1 + \frac{1}{2}g_t^2 - \frac{1}{2}g_{t-1}^2 \\ &= \frac{1}{2}g_t^2 + g_t^1 + \frac{1}{2}g_{t-1}^2, \end{aligned}$$

where the second part follows from a first order Taylor series approximation. Using recursive backward substitution twice, we can write the expression above as:

$$\begin{aligned} \Delta y_t &= \frac{1}{2}(a^2 g_{t-1}^2 + \varepsilon_t^2 + a\varepsilon_t^1) + (a^2 g_{t-1}^1 + \varepsilon_t^1 + a\varepsilon_{t-1}^2) + \frac{1}{2}(a^2 g_{t-2}^2 + \varepsilon_{t-1}^2 + a\varepsilon_{t-1}^1) \\ &= a^2 \Delta y_{t-1} + \left[\frac{1}{2}\varepsilon_t^2 + \left(\frac{a}{2} + 1\right)\varepsilon_t^1 + \left(a + \frac{1}{2}\right)\varepsilon_{t-1}^2 + \frac{a}{2}\varepsilon_{t-1}^1 \right] \\ &= a^2 \Delta y_{t-1} + w_t. \end{aligned}$$

Because $E[w_t w_{t-k}] = 0$ for $k > 1$, Δy_t follows an ARMA(1,1) process. Note that if persistence is a at the bi-annual level, persistence becomes a^2 at the annual level. Also, $\sigma_w^2 = 0.5(3 + 4a + 3a^2)\sigma_\varepsilon^2$. In our simulations, we abstract from the MA component at the annual frequency. We assume that income growth is an AR(1) both at the bi-annual and at the annual frequencies. In other words, if bi-annual income growth is an AR(1) process with persistence a and variance σ_ε^2 , annual income growth is approximated by an AR(1) process with persistence a^2 and variance σ_w^2 . For example, an annual persistence of 0.07—the average persistence for idiosyncratic state-level income—corresponds to a bi-annual persistence of 0.265. For that level of persistence, an annual $\sigma_w = 0.018$, corresponds to a bi-annual $\sigma_\varepsilon = 0.0123$. Note that if persistence was 0, $\sigma_w = \sqrt{\frac{3}{2}}\sigma_\varepsilon$.

Transitory shocks

Let us now assume that there are only transitory shocks to income. Income in the second half of period t is: $Y_t^2 = V_t^2$. $\log V_t^2 \equiv v_t^2 \sim N(-\sigma_v^2/2 - \log(2), \sigma_v^2)$, where $-\log(2)$ guarantees that the transitory shock at the annual frequency still has mean 1. Income growth at the bi-annual frequency is simply $\Delta y_t^2 = v_t^2 - v_t^1$. At the annual frequency, $Y_t = Y_t^2 + Y_t^1 = (e^{v_t^2} + e^{v_t^1})$, and

income growth:

$$\begin{aligned}\Delta y_t &= \log(e^{v_t^2} + e^{v_t^1}) - \log(e^{v_{t-1}^2} + e^{v_{t-1}^1}) \\ &\simeq \frac{v_t^2 + v_t^1}{2} - \frac{v_{t-1}^2 + v_{t-1}^1}{2},\end{aligned}$$

where the second part follows from a first order Taylor approximation. Without time-aggregation and an annual calibration, $\Delta y_t = u_t - u_{t-1}$ (where u_t is the log of the annual transitory shock). To obtain the same income growth volatility, we need to set $\sigma_v^2 = 2\sigma_u^2$. For example, for a standard deviation of the transitory shock at the annual level of 0.07, we should use a standard deviation of approximately 0.1 at the bi-annual level.

References

- ABOWD, J., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–445.
- ABRAHAM, A., AND N. PAVONI (2005): "The Efficient Allocation of Consumption under Moral Hazard and Hidden Access to the Credit Market," *Journal of the European Economic Association*, 3, 370–381.
- ASDRUBALI, P., B. E. SØRENSEN, AND O. YOSHA (1996): "Channels of Interstate Risk Sharing: United States 1963–1990," *Quarterly Journal of Economics*, 111, 1081–1110.
- ATTANASIO, O., AND N. PAVONI (2006): "Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption," Mimeo, UCL.
- ATTANASIO, O., AND G. WEBER (1993): "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies*, 60, 631–649.
- BERNANKE, B. S. (1984): "Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence From Panel Data," *Quarterly Journal of Economics*, 99, 587–614.
- BLINDER, A., AND A. DEATON (1985): "The Time Series Consumption Function Revisited," *Brookings Papers on Economic Activity*, 2, 465–521.
- CAMPBELL, J. Y., AND A. DEATON (1989): "Why is Consumption so Smooth?" *Review of Economic Studies*, 56, 357–374.
- CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–156.
- (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55.
- (2000): "'Risky Habits' and the Marginal Propensity to Consume out of Permanent Income," *International Economic Journal*, 4, 1–40.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, 41–71.
- DEATON, A. (1987): "Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" in *Advances in Econometrics. Fifth World Congress, Vol. II*, ed. by Truman F. Bewley. Cambridge University Press, New York, NY.
- (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1121–1248.
- (1992): *Understanding Consumption*. Oxford University Press, New York, NY.
- DEJUAN, J. P., AND M. J. LUENGO-PRADO (2006): "Consumption and Aggregate Constraints: International Evidence," *Oxford Bulletin of Economics and Statistics*, 68, 81–99.
- DYNAN, K. (2000): "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 90, 391–406.

- FLAVIN, M. (1981): “The Adjustment of Consumption to Changing Expectations about Future Income,” *Journal of Political Economy*, 89, 974–1009.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Princeton University Press, Princeton, NJ.
- GALÍ, J. (1991): “Budget Constraints and Time Series Evidence on Consumption,” *American Economic Review*, 81, 1238–1253.
- GOURINCHAS, P.-O., AND J. PARKER (2002): “Consumption over the Life Cycle,” *Econometrica*, 70, 47–91.
- HALL, R. E. (1978): “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 86, 971–987.
- HANSEN, L. P., W. ROBERDS, AND T. SARGENT (1991): “Time Series Implications of Present Value Budget Balance and of Martingale Models of Consumption and Taxes,” in *Rational Expectations Econometrics*, ed. by L. P. Hansen, and T. Sargent. Underground Classics in Economics, Westview Press, Boulder, Colorado and Oxford, UK.
- HANSEN, L. P., AND T. SARGENT (1981): “A Note on Wiener-Kolmogorov Prediction Formulas for Rational Expectations Models,” *Economics Letters*, 8, 255–260.
- JUDD, K. L. (1997): *Numerical Methods in Economics*. MIT Press, Cambridge, MA.
- LUDVIGSON, S., AND A. MICHAELIDES (2001): “Can Buffer Stock Saving Explain the Consumption Excesses?” *American Economic Review*, 91, 631–647.
- LUENGO-PRADO, M. J. (2006): “Durables, Nondurables, Down Payments, and Consumption Excesses,” *Journal of Monetary Economics*, Forthcoming.
- MACURDY, T. E. (1982): “The Use of Time-Series Processes to Model the Error Structure of Earnings in Longitudinal Data analysis,” *Journal of Econometrics*, 18, 83–114.
- MICHAELIDES, A. (2001): “Buffer Stock Saving and Habit Formation,” Mimeo, LSE.
- OSTERGAARD, C., B. E. SØRENSEN, AND O. YOSHA (2002): “Consumption and Aggregate Constraints: Evidence from U.S. States and Canadian Provinces,” *Journal of Political Economy*, 110, 634–645.
- TAUCHEN, G. (1986): “Finite State Markov Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–181.
- WORKING, H. (1960): “Note on the Correlation of First Differences of Averages in a Random Chain,” *Econometrica*, 28, 916–918.

TABLE 1: PARAMETERS OF TIME SERIES PROCESS FOR STATE-LEVEL DISPOSABLE LABOR INCOME.

	(1) Persistence	(2) σ_{G_s}
Alabama	0.16 (0.17)	0.74 (0.09)
Alaska	0.37 (0.16)	4.57 (0.54)
Arizona	0.35 (0.17)	1.31 (0.15)
Arkansas	-0.14 (0.17)	1.70 (0.20)
California	0.17 (0.19)	1.17 (0.14)
Colorado	0.41 (0.15)	1.26 (0.15)
Connecticut	0.39 (0.16)	1.56 (0.18)
Delaware	-0.04 (0.17)	1.46 (0.17)
Florida	0.09 (0.17)	1.24 (0.15)
Georgia	0.33 (0.16)	0.83 (0.10)
Hawaii	0.20 (0.16)	2.65 (0.31)
Idaho	-0.53 (0.15)	2.49 (0.29)
Illinois	0.07 (0.16)	0.92 (0.11)
Indiana	-0.06 (0.17)	1.51 (0.18)
Iowa	-0.38 (0.15)	3.28 (0.39)
Kansas	-0.12 (0.16)	1.25 (0.15)
Kentucky	-0.24 (0.15)	1.21 (0.14)
Louisiana	0.50 (0.14)	1.73 (0.20)
Maine	-0.16 (0.16)	1.61 (0.19)
Maryland	0.20 (0.17)	1.38 (0.16)
Massachusetts	0.45 (0.15)	1.48 (0.17)
Michigan	0.07 (0.16)	1.78 (0.21)
Minnesota	-0.28 (0.15)	1.78 (0.21)
Mississippi	0.02 (0.17)	1.55 (0.18)
Missouri	-0.25 (0.16)	1.22 (0.14)
Montana	-0.23 (0.16)	2.78 (0.33)
Nebraska	-0.49 (0.14)	2.74 (0.32)
Nevada	0.40 (0.15)	1.58 (0.19)
New Hampshire	0.29 (0.16)	2.00 (0.24)
New Jersey	0.18 (0.18)	1.40 (0.17)
New Mexico	0.08 (0.17)	1.37 (0.16)
New York	0.10 (0.17)	1.55 (0.18)
North Carolina	-0.06 (0.17)	1.04 (0.12)
North Dakota	-0.18 (0.16)	9.28 (1.09)
Ohio	-0.28 (0.16)	1.07 (0.13)
Oklahoma	0.26 (0.16)	1.58 (0.19)
Oregon	0.05 (0.17)	1.51 (0.18)
Pennsylvania	0.07 (0.17)	0.94 (0.11)
Rhode Island	0.15 (0.16)	1.83 (0.22)
South Carolina	0.15 (0.16)	1.10 (0.13)
South Dakota	-0.30 (0.15)	5.17 (0.61)
Tennessee	-0.01 (0.16)	0.97 (0.11)
Texas	0.48 (0.15)	1.28 (0.15)
Utah	0.35 (0.16)	1.20 (0.14)
Vermont	0.03 (0.16)	1.50 (0.18)
Virginia	0.14 (0.16)	1.14 (0.13)
Washington	0.40 (0.15)	1.45 (0.17)
West Virginia	0.15 (0.16)	1.52 (0.18)
Wisconsin	-0.08 (0.16)	1.01 (0.12)
Wyoming	0.32 (0.15)	2.62 (0.31)

Notes: The table displays the estimated parameters of a time series model for real disposable labor income. Let y_{st} be the log of per capita disposable labor income (deflated by the CPI) in state i . The model is: $y_{st} = \mu_s + \log G_{st}$, where μ_s is a constant for each state, and $\log G_{st} = a^s \log G_{st-1} + \sigma_{G_s} \epsilon_{st}$ where ϵ_{st} are i.i.d. normal innovations. a^s is persistence and σ_{G_s} is the standard deviation of the permanent component. The table reports the estimates of a^s in column (1) and 100 times σ_{G_s} in column (2). Standard errors in parentheses. Sample 1964–1998.

TABLE 2: CORRELATION MATRICES OF REGRESSAND AND REGRESSORS

	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(0) c_t	1.00	0.12	-0.06	0.02	0.03	0.03	-0.01	0.07
(1) y_t		1.00	0.03	0.06	-0.46	0.72	0.60	0.41
(2) Farm Share			1.00	-0.25	-0.03	0.02	0.06	-0.06
(3) Gov't. Share				1.00	0.02	0.04	-0.01	0.15
(4) Persistence $\times y_t$					1.00	-0.50	-0.76	0.19
(5) $\hat{\sigma}_{G_s} \times y_t$						1.00	0.87	0.42
(6) (Farm Share) $\times y_t$							1.00	0.08
(7) (Gov't. Share) $\times y_t$								1.00
(0) c_t	1.00	0.05	-0.06	0.02	0.07	-0.02	-0.08	0.11
(1) y_{t-1}		1.00	0.03	0.03	-0.45	0.71	0.60	0.40
(2) Farm Share			1.00	-0.25	-0.04	0.04	0.07	-0.06
(3) Gov't. Share				1.00	-0.02	0.04	0.02	0.07
(4) Persistence $\times y_{t-1}$					1.00	-0.50	-0.76	0.19
(5) $\hat{\sigma}_{G_s} \times y_{t-1}$						1.00	0.87	0.42
(6) (Farm Share) $\times y_{t-1}$							1.00	0.08
(7) (Gov't. Share) $\times y_{t-1}$								1.00

Notes: The table shows how the variables entering the regressions correlate across the 50 U.S. states over the years 1970–1998. c_t is state s 's nondurable consumption growth (real and per capita) in period t after subtracting state- and time-specific means. y_t is real per capita disposable labor income growth minus the state- and time-specific means. “Farm Share” is the ratio of employees (including proprietors) in farming to the total number of employees in each state in period t in the top part of the panel, and for period $t - 1$ in the bottom part of the panel, with the average for each time period and for each state subtracted. “Gov't. Share” is defined similarly. “Persistence” refers, for each state, to the number in the column (1) in Table 1 minus its mean. Row (4) corresponds to that number interacted with income with time- and state-specific means subtracted. For row (5) the standard deviations from Table 1 are treated similarly. Farm-share in row (6) is the farm share as just defined minus its time-specific mean; this is then multiplied by income with time- and state-specific means subtracted. Gov't. share in row (7) is treated similarly.

TABLE 3: SENSITIVITY TO CURRENT AND LAGGED INCOME: NON-DURABLE RETAIL SALES
1970-1998

	CURRENT MPC			LAGGED MPC		
	(1)	(2)	(3)	(4)	(5)	(6)
y_{st}	0.336*** (0.049)	0.339*** (0.047)	0.333*** (0.047)	- -	- -	- -
$y_{s,t-1}$	- -	- -	- -	0.163*** (0.046)	0.163*** (0.045)	0.171*** (0.042)
Farm Share	-0.001 (0.001)	- -	- -	-0.001 (0.001)	- -	- -
Gov't. Share	0.000 (0.001)	- -	- -	0.000 (0.001)	- -	- -
Interaction terms:						
Persistence	0.608*** (0.175)	0.588*** (0.173)	0.530*** (0.136)	-0.089 (0.157)	-0.087 (0.155)	- -
$\hat{\sigma}_{G_s}$	-0.024 (0.023)	-0.024 (0.022)	-0.026* (0.015)	0.005 (0.018)	0.004 (0.018)	- -
Farm share	0.003 (0.013)	0.001 (0.013)	- -	-0.037*** (0.012)	-0.037*** (0.012)	-0.032*** (0.006)
Gov't. share	-0.008 (0.010)	-0.005 (0.009)	- -	0.029*** (0.010)	0.031*** (0.009)	0.030*** (0.008)

Notes: Model: $c_{st} = \mu_s + v_t + \alpha y_{s,t-l} + \zeta(X_{st} - \bar{X}_{.t})(y_{s,t-l} - \bar{y}_{.t-l} - \bar{y}_{s.} + \bar{y}_{..}) + u_{st}$, where l is 0 for current and 1 for lagged income. y_{st} is state s 's labor income growth (real and per capita). c_{st} is state s 's nondurable consumption growth (real and per capita). μ_s is a cross-sectional fixed effect and v_t is a time-fixed effect. X is one of the variables that may affect the MPC, listed as "interaction terms"—see Table 2 for precise definitions. Feasible GLS-estimation allowing for the innovation error variance to differ by year and by state. Standard errors in parentheses. The standard errors for columns (1)-(3), which include generated regressors, are estimated using parametric Monte Carlo simulation as described in Appendix B. Sample 1970–1998. ***, **, and * indicate significance at the 1, 5 and 10% level, respectively.

TABLE 4: BENCHMARK PARAMETERS FOR MODEL SIMULATIONS

Parameter	Annual Frequency	Bi-annual frequency
Aggregate growth, μ_A	0.016	0.008
Aggregate shock persistence, a^A	0.420	0.648
Aggregate shock volatility, σ_A	0.020	0.011
State-level shock persistence, a^s	0.070	0.265
State-level shock volatility, σ_{G_s}	0.018	0.012
Permanent idiosync. shock volatility, σ_N	0.050	0.041
Transitory idiosync. shock volatility, σ_V	0.070	0.099
Interest rate, $R - 1$	0.020	0.010
Discount rate, $1/\beta - 1$	0.050	0.025
Depreciation rate, $\psi - 1$	0.085	0.043
Adjustment cost parameter, ϕ	0.050	0.050
Down payment, θ	0.300	0.300

Notes: The aggregate and state-level income shocks are calibrated according to the income process estimation of Section 2.1. The other parameters, including the idiosyncratic income shocks, are taken from previous studies—see the main text for more particulars. Details on how to adjust the income parameter values to the bi-annual frequency are described in Appendix E.

TABLE 5: SENSITIVITY TO CURRENT AND LAGGED INCOME IN SIMULATED DATA

	NO TIME-AGGREGATION				TIME-AGGREGATION			
	PIH	ROT	Buffer	Buffer-Durable	PIH	ROT	Buffer	Buffer-Durable
NO RISK SHARING								
Current MPC	1.068	1.000	1.000 (0.003)	0.790 (0.007)	1.042	1.000	0.983 (0.006)	0.756 (0.010)
Lagged MPC	0.000	0.070	0.041 (0.048)	0.096 (0.038)	0.148	0.070	0.192 (0.046)	0.358 (0.033)
50% RISK SHARING								
Current MPC	0.534	0.500	0.500 (0.001)	0.395 (0.004)	0.521	0.500	0.491 (0.003)	0.378 (0.005)
Lagged MPC	0.000	0.035	0.028 (0.024)	0.048 (0.019)	0.074	0.035	0.096 (0.023)	0.179 (0.017)

Notes: The rows labeled “current” report the estimated value of the parameter α^c from the regression: $\Delta \log C_{st} = \mu_s + v_t + \alpha^c \Delta \log Y_{st} + u_{st}$, where C_{st} and Y_{st} are state nondurable consumption and income respectively. The columns labeled “lagged” report the estimated value of the parameter α^l for lagged income from the regression $\Delta \log C_{st} = \mu_s + v_t + \alpha^l \Delta \log Y_{s,t-1} + u_{st}$. For the columns “buffer” and “buffer-durable,” consumption is simulated as described in the text. The simulated data is based on the individual-level income process $\Delta \log Y_{jt} = \log A_t + \log G_{st} + \log N_{jt} + \log V_{jt} - \log V_{j,t-1}$, where $\log A_t$ and $\log G_{st}$ are AR(1) processes with persistence a^A and a^s respectively, unconditional means μ_A and 0, and standard deviation σ_A and σ_{G_s} . $\log N_{jt}$ and $\log V_{jt}$ are independent and identically distributed with standard deviations σ_N and σ_V , and mean values $-\sigma_N^2/2$ and $-\sigma_V^2/2$ respectively. Baseline parameters described in Table 4. “State”-level consumption and income (C_{st} , Y_{st}) are averages over 3,000 individuals in each “state.” There are 10 “states” and 50 periods. We report average MPCs for 20 independent simulations. Average estimated standard errors in parentheses. For the columns labeled “PIH” the numbers are calculated for a log-linear approximation of a representative-agent PIH-model, with the interest rate equal to the discount rate and the representative agent receiving the aggregate income process. In the case with time-aggregation, agents are assumed to make decisions bi-annually but the data is transformed to the annual frequency before running the regressions. All parameters, including income parameters, are adjusted to a bi-annual frequency in this case.

TABLE 6: SENSITIVITY TO CURRENT AND LAGGED INCOME IN SIMULATED DATA. COMPARISONS TO BENCHMARK IN SELECTED MODELS

	BUFFER-DURABLES				PIH				COMBINED			
	TIME AGG. AND 50% RISK SHARING		MARGINAL EFFECT		TIME AGG. AND 50% RISK SHARING		MARGINAL EFFECT		MPC		MARGINAL EFFECT	
	current	lagged	current	lagged	current	lagged	current	lagged	current	lagged	current	lagged
Benchmark	0.378 (0.005)	0.179 (0.017)	-	-	0.521	0.074	-	-	0.450	0.126	-	-
$a^s = -0.1$	0.352 (0.007)	0.129 (0.017)	0.152	0.291	0.399	0.084	0.719	-0.061	0.375	0.107	0.436	0.115
$\sigma_V = 0.035$	0.375 (0.006)	0.206 (0.016)	0.001	-0.005	0.521	0.074	0.000	0.000	0.448	0.140	0.000	-0.003
$\sigma_N = 0.025$	0.359 (0.006)	0.199 (0.016)	0.008	-0.008	0.521	0.074	0.000	0.000	0.440	0.137	0.004	-0.004
$\sigma_{G_s} = 0.036$	0.382 (0.005)	0.174 (0.017)	0.002	-0.002	0.521	0.074	0.000	0.000	0.451	0.124	0.001	-0.001

Notes: All definitions and parameters as in Table 5, except the parameter being altered. The marginal effects are calculated as the change in the corresponding MPC relative to the benchmark case divided by the change in the parameter being altered. The marginal effects relating to uncertainty parameters (the last three) are divided by 100.