DISCUSSION PAPER SERIES

No. 4474

THE BUFFER-STOCK MODEL AND THE AGGREGATE PROPENSITY TO CONSUME: A PANEL-DATA STUDY OF THE US STATES

Maria José Luengo-Prado
and Bent E Sørensen

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP4474.asp**

ISSN 0265-8003

THE BUFFER-STOCK MODEL AND THE AGGREGATE PROPENSITY TO CONSUME: A PANEL-DATA STUDY OF THE US STATES

Maria José Luengo-Prado, Northeastern University
Bent E Sørensen, University of Houston and CEPR

Discussion Paper No. 4474
July 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4474

July 2004

ABSTRACT

The Buffer-Stock Model and the Aggregate Propensity to Consume: A Panel-Data Study of the US States

We simulate a buffer-stock model of consumption, explicitly aggregate over consumers, and estimate aggregate marginal propensities to consume out of current and lagged income using simulated data generated by the model. We calculate the predicted marginal effects of changing persistence of income shocks, aggregate-level uncertainty, and individual-level risk. Next, we estimate marginal propensities for US states using panel-data methods. We find effects of persistence that clearly correspond to the predictions of the model and while the effect of aggregate uncertainty cannot be determined precisely, indicators of individual level uncertainty have strong effects consistent with the model. Overall, the buffer-stock model clearly helps explain differences in consumer behaviour across states.

JEL Classification: E21
Keywords: buffer-stock, consumption and precautionary saving

Maria José Luengo-Prado
Departments of Economics
Northeastern University
301 LA
Boston MA 02115
USA
Tel: (1 617) 373 4520
Email: m.luengo@neu.edu

Bent E Sørensen
Department of Economics
University of Houston
204 McElhinney Hall
Houston, TX 77004
USA
Tel: (1 713) 743 3841
Fax: (1 713) 743 3798
Email: bent.sorensen@mail.uh.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=160003

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=126829

Submitted 20 May 2004

1 Introduction

The buffer-stock model of consumption, pioneered by Deaton (1991) and Carroll (1997), is a promising candidate for replacing the Friedman (1957)-Hall (1978) Permanent Income Hypothesis (PIH) as the benchmark model of consumer behavior. In this paper, we examine if the buffer-stock model predicts the cross-state variation in the marginal propensities to consume (MPC) out of current and past income observed in U.S. state-level aggregate data. Rather than testing if a particular implementation of the model is literally true, we examine if the model predicts the directions in which the MPCs vary across states.

Hall's (1978) paper shocked the profession by demonstrating that under simple assumptions 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 smaller reaction of consumption to income shocks—a phenomenon known as "excess smoothness".[1]

The buffer-stock model was designed to improve on the PIH-model while still preserving many of its insights stemming from rational forward-looking behavior. Roughly speaking, by assuming consumers cannot (or endogenously will not) borrow and allowing consumers to be more prudent and less patient than in Hall (1978), Deaton (1991) and Carroll (1997) demonstrate the buffer-stock model has the potential to fit microeconomic and macroeconomic data better. In particular, the buffer-stock model predicts a high correlation of consumption with income regardless of expected future income, and that consumers will save more when they are subject

[1] For 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).

to more uncertainty ("precautionary saving").[2]

Several papers have focused on identifying the presence of precautionary saving using micro data. The major challenge faced by this line of work is finding variables that reflect uncertainty but do not correlate with other characteristics that may increase saving. Commonly used uncertainty measures are household income variance, occupation, education, and unemployment risk. The literature offers diverse estimates ranging from no evidence of precautionary saving (e.g., Dynan 1993) to quite large amounts of precautionary saving (e.g., Carroll and Samwick 1997). Browning and Lusardi (1996) provide a comprehensive survey of the empirical literature. Gourinchas and Parker (2001) estimate a structural model of optimal life-cycle consumption in the presence of realistic labor income uncertainty and find that households behave like buffer-stock consumers early in their working lives and like PIH households when they approach retirement. More recently, Carroll, Dynan, and Krane (2003) identify a precautionary effect on saving induced by unemployment risk for households with moderate levels of permanent income but not for households with low levels of income. Engen and Gruber (2001) utilize the variation in unemployment insurance programs across states and find lower precautionary saving for households with high unemployment risk in states with more generous unemployment insurance. Consistent with the buffer-stock model, McCarthy (1995) finds a higher marginal propensity to consume for households with relatively low-wealth.

The amount of empirical evidence supporting the buffer-stock model with aggregate data is limited. Hahm (1999) finds higher saving rates and steeper consumption growth paths for countries with more earnings uncertainty using OECD data for 22 countries. Hahm and Steigerwald (1999)—using U.S. aggregate data and a survey measure of GDP uncertainty as an approximation of aggregate income uncertainty—find more precautionary saving in periods of high

[2]Numerous authors have studied the effects of uncertainty with non-quadratic preferences. Leland (1968) names the difference between saving under uncertainty and saving under certainty "precautionary saving". Kimball (1990) discusses the conditions under which one can expect to observe precautionary saving. Carroll and Kimball (1996) prove the concavity of the consumption function under precautionary saving. Caballero (1991) uses the exponential utility function and evaluates the effect precautionary saving may have empirically. Under the assumption that consumers have constant relative risk aversion utility functions, it is not possible to obtain closed-form solutions and, therefore, numerical methods must be used to study the reaction of consumption and saving to uncertainty. Early numerical studies of precautionary saving include Zeldes (1989b), Hubbard, Skinner, and Zeldes (1994), Skinner (1988) and Aiyagari (1994). Evidence of credit rationing has been provided by, e.g., Zeldes (1989a) using a sample splitting method, and by Jappelli (1990) and Perraudin and Sørensen (1992) using survey data.

uncertainty of GDP. Ludvigson and Michaelides (2001) demonstrate that the buffer-stock model under certain conditions of incomplete information—*à la* Pischke (1995)—can explain at least part of the deviations from the predictions of the PIH-model at the aggregate level.

Our approach in this paper is as follows. First, we estimate a process for aggregate income for each U.S. state.[3] State-level income growth is well described by an autoregressive (AR) model of order one for most states. However, since the buffer-stock model predicts different consumption behavior for agents facing more transitory uncertainty, we generalize the process for income as the sum of an AR-model and a temporary shock.

Second, we simulate the buffer-stock model for aggregate consumption by simulating individual-level consumption and explicitly aggregating across consumers. The logarithm of individual-level income is modelled as the sum of aggregate state-level income and individual-level idiosyncratic income, where the latter is the sum of a random walk and a temporary shock. The standard deviations of the idiosyncratic components are calibrated to have the values chosen in most of the literature. We simulate the income and consumption processes for 3,000 consumers and regress the simulated aggregate series for consumption growth on the simulated aggregate current or lagged income growth. We repeat the simulations changing the benchmark parameters one at a time in order to examine the predicted effects in empirically relevant directions. Specifically, we add unemployment to the benchmark model; i.e., we add the possibility of "disastrous," large, independently identically distributed (i.i.d.) temporary shocks which happen with low probability. Alternatively, we increase aggregate uncertainty or decrease individual-specific risk. We calculate the marginal effect of each of these changes to compare with the consumption regressions described next.

Third, using the panel of U.S. states, we estimate the MPCs out of current and lagged income by regressing consumption growth on current and lagged income growth, respectively. We allow the estimated MPCs to vary with persistence, with the unemployment rate, with aggregate uncertainty, and with indicators for high (low) individual-level uncertainty; viz., the share of

[3]The advantage of using state-level data are the following: states display significant variation in the behavior of aggregate income, the data are collected in a consistent manner and most institutional features do not vary across states, making our results less likely to suffer from left-out variable bias compared to international data. Further, as argued by Ostergaard, Sørensen, and Yosha (2002) and Sørensen and Yosha (2000), the use of panel-data regressions with time-fixed effects will make the results more robust to potential biases that might obtain because the U.S. as whole is unable to borrow internationally at fixed interest rates.

agriculture (government) employment in each state.

We find strong effects of income persistence on the marginal MPC out of current income in the direction predicted by the buffer-stock model. Persistence does not have a significant effect on the MPC out of lagged income but this result is quite consistent with our simulation results if we allow for a non-negligible proportion of the consumers to behave according to the PIH-model as suggested by Gourinchas and Parker (2001). We find statistically insignificant evidence of unemployment affecting the current MPC but significant effects of unemployment on the lagged MPC; the buffer-stock model would have predicted clear effects on both. The signs for both current and lagged MPCs are consistent with the predictions of our model. We do not find significant effects of aggregate uncertainty—maybe because the estimates of aggregate-level uncertainty are too imprecise. However, we find large effects of individual-level uncertainty on the aggregate MPC out of lagged income, consistent with the predictions of the model. Overall, the buffer-stock model clearly helps explain differences in consumer behavior across states.

The remainder of the paper is organized as follows: Section 2 estimates a state-level aggregate labor income. Section 3 describes the buffer-stock model and the simulation results. Section 4 estimates panel-data models for consumption and compares the estimated MPCs to the theoretical results found in Section 3. Section 5 summarizes.

2 Estimating the Process for Aggregate State-Level Disposable Labor Income

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 4 use a shorter sample, 1976–1998, due to lack of availability of other variables prior to 1976. However, in order to obtain more precise parameter estimates we use the larger sample here.

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.[4] 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.[5] We, therefore, treat labor income growth as a stationary series.

State-level income data (see Ostergaard, Sørensen, and Yosha 2002) are typically well-approximated by first order autoregressive (AR) models of order 1. Nonetheless, some states are highly dependent on agriculture, an industry that is subject to large temporary shocks (weather, epidemics of livestock diseases, etc.). Since temporary uncertainty is crucial for the consumption behavior of buffer-stock savers, it is potentially important to capture such temporary components of the income process and, accordingly, we model state-level labor income as the sum of a "permanent component" that follows an AR(1) model after differencing and a white noise "temporary component."

We assume the growth rate of real per capita disposable labor income, $\Delta \log Y_{it}$, in state i follows the model:

$$\Delta \log Y_{it} \ = \mu_i + \log G_{it} + \sigma_{Wi} \, (W_{it} - W_{it-1}), \tag{1}$$

where W_{it} is a temporary shock (an i.i.d. variable) with variance 1. The parameter σ_{Wi} is the standard deviation of the temporary shock. We refer to states with a higher value of σ_{Wi} as states with higher "temporary uncertainty." μ_i is a state-specific constant. The permanent component of the growth rate follows the AR(1) model:

$$\log G_{it} = a_i \log G_{it-1} + \sigma_{Gi} \, \epsilon_{it} \ , \tag{2}$$

where ϵ_{it} are i.i.d., mean zero innovations with variance 1 and the parameter σ_{Gi} is the standard

[4]For the 1976–1998 sample, we can reject a unit root in labor income for only 7 states at the 5 percent level of significance and for no states at the 1 percent level. Using the longer 1964–1998 sample, we can reject the unit root for only 2 states at the 5 percent level and for none at the 1 percent level.

[5]Panel unit root tests are not attractive for these series since they are highly correlated across states. Ostergaard, Sørensen, and Yosha (2002) show that panel unit root tests for disposable income—when aggregate income is subtracted, making the data less correlated—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 would, therefore, be similar.

deviation of the innovation to the permanent income component. The larger the parameter a_i, the stronger the impact of past shocks on current income or, equivalently, the stronger the impact of a current innovation to income on permanent income (the interest rate times the present value of the income shock plus the change in the expected value of future income). We say shocks are "more persistent" the larger a_i.

The model is deceptively simple and although it is known to be equivalent to an ARMA(2,1) with restriction across the parameters, it is complicated to utilize this knowledge for estimation. Instead, we estimate the model by Maximum Likelihood.

Maximum Likelihood estimation of the income processes.

Let $y_{it} \equiv \Delta \log Y_{it}$ denote the growth rate of state-level labor income. We apply a Kalman filter technique to evaluate the likelihood function recursively, assuming the initial observation is generated by the long-run stationary distribution implied by the model. The general workings of Kalman filters are described in econometric textbooks. Concerning the technical details, we parameterize the filter in terms of a transition equation for state i of the form $x_{it} = Ax_{i,t-1} + w_{it}$ where

$$A = \begin{pmatrix} a_i & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 0 \end{pmatrix},$$

and $w_t' = (u_{it}, v_{it}, v_{it})$. The measurement equation is $y_{it} = (1, 1, 0)\, x_{it}$. a_i is the persistence parameter, $u_{it} = \sigma_{Gi}\epsilon_{it}$, and $v_{it} = \sigma_{Wi}W_{it}$. The remaining parts of the Kalman filter implementation are standard and we leave out the details.

Although it is fairly straightforward to estimate the model using the Kalman filter, the estimations resulted in local minima for the likelihood function for a low number of states. These local minima were always found for $\hat{\sigma}_{Wi} = 0$ and because the standard deviation cannot be negative, this value is at the boundary of the parameter space—a situation that sometimes complicates inference. Consequently, we hedged against local minima by performing grid-searches. Basically, we chose a grid for $\sigma_{Wi} = 0$ and optimized over the remaining two parameters; this grid search revealed the likelihood function to have one global maximum and no other local

maxima for all states.[6]

The results from the estimation are presented in Table 1. The estimates of the standard deviation of the aggregate transitory shock, $\hat{\sigma}_{Wi}$, is non-zero for 17 states and significant at conventional levels for only a few. However, the states with relatively large transitory shocks are typically agricultural (Iowa, Nebraska, North Dakota, etc.) and since our prior is that agricultural states will be subject to temporary shocks, we find it important to examine if the states with a large estimated component of temporary uncertainty have different MPCs.

We find the lowest values of a_i to be around 0 for Iowa, Montana, Nebraska, North Dakota, and South Dakota, while the largest value of 0.81 is found for Hawaii followed by 0.78 for Oklahoma. In general, it looks like income shocks tend to be more persistent in oil-states such as Alaska and Wyoming but we do not know the economic reasons behind these large differences across states.

3 Uncertainty and Aggregate Consumption: Theory

This section presents the buffer-stock model used to predict the effects of persistence and uncertainty on aggregate consumption. The specification is standard with the exception of the addition of transitory shocks to the aggregate income process. In order to explore the aggregate implications of the buffer-stock model, we simulate income and consumption paths for individuals and find aggregate income and consumption paths by averaging across consumers.

3.1 The model

Consumer j's maximizes the present discounted value of expected utility from consumption of a nondurable good, C. Let $\beta < 1$ be the discount factor and R be the interest factor. A_{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 RA_{jt-1} plus Y_{jt} units of labor income. The agent

[6] We experimented with a few grid searches for the other parameters and none of those indicated further problems with local minima.

chooses optimal consumption C_{jt} according to the maximization problem:

$$\max_{C_{jt}} \ \mathrm{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \, U(C_{jt}) \right\}$$
$$\text{s.t. } A_{jt} = RA_{jt-1} + Y_{jt} - C_{jt}.$$

Utility is assumed to be constant relative risk aversion (CRRA): $U(C_{jt}) = \frac{C_{jt}^{1-\rho}}{1-\rho}$. With $\rho > 0$ the agent is risk-averse and has a precautionary motive for saving.

In the literature, buffer-stock saving behavior has been derived from two different assumptions. Deaton (1991) explicitly imposes a no borrowing constraint ($A_{jt} > 0$) but assumes agents always receive positive income. Carroll (1997), 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, p. In this case the agent will optimally never want to borrow to avoid $U(0) = -\infty$. We use Deaton's specification as our benchmark with p, the probability of the disastrous state, set to zero. Following Michaelides (2003), we also consider the case where p is not zero and impose different lower bounds for the transitory shock. A positive lower bound may be interpreted as an income replacement program (unemployment benefits, welfare, disability payments, etc.). We refer to the disastrous state as unemployment.

Income is stochastic and the only source of uncertainty in the model. It is assumed to be exogenous to the agent. We assume:

$$Y_{jt} = P_{jt} V_{jt} W_t,$$

$$P_{jt} = G_t P_{jt-1} N_{jt}.$$

Labor income, Y_{jt}, is the product of permanent income, P_{jt}, an idiosyncratic transitory shock, V_{jt}, and an aggregate transitory shock, W_t. G_t can be thought of as the growth in permanent income attributable to aggregate productivity growth in the economy. N_{jt} is a permanent idiosyncratic shock. $\log N_{jt}$, $\log V_{jt}$ and $\log W_t$ are independent and identically normally distributed with means $-\sigma_N^2/2$, $-\sigma_V^2/2$ $-\sigma_W^2/2$, and variances σ_N^2, σ_V^2 and σ_W^2, respectively. $\log G_t$ is assumed to be an $AR(1)$ process with persistence a, unconditional mean μ_G, and variance σ_G^2.

This income specification is particularly useful since it allows for consumers to share in general 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, $\Delta \log Y_{jt} = \log G_t + \log N_{jt} + \log V_{jt} - \log V_{jt-1} + \log W_t - \log W_{t-1}$, consistent with microeconomic evidence; see for example MaCurdy (1982) and Abowd and Card (1989). By the law of large numbers, aggregate income growth can be written as $\Delta \log Y_t = \log G_t + \log W_t - \log W_{t-1}$.

3.2 Solution Method and Calibration

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, $X_{jt} \equiv RA_{jt-1} + Y_{jt}$.[7] 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:

$$U'(c_{jt}) = \max\{U'(x_{jt}), \beta R \mathrm{E}_t[(G_{t+1}N_{jt+1})^{-\rho} U'(c_{jt+1})]\}, \tag{3}$$

where $c_{jt} = C_{jt}/P_{jt}$ and $x_{j,t+1} = (G_{t+1}N_{j,t+1})^{-1} R(x_{jt} - c_{jt}) + V_{j,t+1}W_{t+1}$.[8]

Equation (3) can be solved numerically to obtain an optimal normalized consumption function for given values of the parameters. In other words, the numerical technique delivers a consumption function $c(x)$: normalized consumption as a function of normalized cash-on-hand.[9]

In this specification of the model, uncertainty arises from stochastic income alone and a good calibration of the income process is essential to obtain qualitative and quantitative predictions. We use our estimated values from the state-level time-series estimations to calibrate the aggregate shocks, G and W. Idiosyncratic income shocks are taken from previous studies—see, for example, Carroll and Samwick (1997) and Gourinchas and Parker (2001). For the bench-

[7] The budget constraint becomes $A_{jt} = X_{jt} - C_{jt}$ 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: $X_{jt+1} = R(X_{jt} - C_{jt}) + Y_{jt+1}$.

[8] A necessary condition for the individual Euler equation to define a contraction mapping is $\beta R E_t[(G_{t+1}N_{j,t+1})^{-\rho}] < 1$. This is the "impatience" condition common to buffer-stock models which guarantees that borrowing is part of the unconstrained plan.

[9] We use Euler equation iteration to solve equation (3). x is discretized and the income shocks are approximated by 10-point discrete Markov processes following Tauchen (1986). Interpolation is used between points in the x grid. More details on how to solve this equation can be found, for example, in the appendix of Ludvigson and Michaelides (2001).

mark calibration, we set risk aversion $\rho = 2$ and $\beta = 0.9524$ (which implies a discount rate of 5 percent). The standard deviations of the idiosyncratic shocks are $\sigma_V = 0.1$ and $\sigma_N = 0.1$, typical values in the literature. For the aggregate shocks, we start with $\sigma_W = 0$—no aggregate transitory shocks—since this was the estimate for most states. For $\log G_t$, we choose persistence $a = 0.39$, unconditional mean $\mu_G = 0.0165$, and standard deviation $\sigma_G = 0.0265$ (the averages of the parameter-values estimated for our state-level income data). When unemployment is allowed for, we set $p = 0.03$ and consider two different replacement rates: 30 percent of income and 0.

3.3 Consumption Functions and Aggregation Procedure

Figure 1 depicts the optimal consumption functions for the case where G is assumed to be an i.i.d. process for certain variations of the parameters.[10] We present Deaton's (1991) case with no disastrous state for two different values of the standard deviation of the idiosyncratic transitory shock, Carroll's (1997) case with unemployment and no replacement, and Michaelides's (2003) case with unemployment and a 30 percent replacement rate. In all cases, the MPC out of cash-on-hand is higher for the cash-on-hand poor. In the "Deaton case," the MPC out of cash-on-hand is equal to 1 when the liquidity constraint is binding. In the "Carroll case," agents optimally choose to never borrow and the MPC is always below 1. Figure 1 illustrates that when uncertainty increases—either by increasing the probability of the disastrous state, decreasing unemployment insurance, or increasing the variance of the transitory shock—the consumption functions shift down because of precautionary saving. Figure 2 shows the changes in the current MPC out of normalized cash-on-hand for the Deaton case with no unemployment risk, and for the Carroll case with positive unemployment risk and no replacement. The MPC is clearly higher in the Deaton case for the cash-on-hand poor. However, cash-on-hand is endogenous, and due to precautionary saving agents typically have more cash-on-hand in the Carroll case. The question then is whether the result suggested by Figure 2—the more uncertainty, the lower the MPC—survives aggregation. This needs to be verified through simulations.

Given the nonlinearity of the consumption functions, one must aggregate explicitly in order

[10] We show the i.i.d. case instead of our baseline case with persistence because with persistence the optimal policy function has 10-branches, one for each shock in our 10-point Markov discrete approximation of G. This would make the figure too busy for illustration.

to obtain aggregate implications. Moreover, since normalized consumption and cash-on-hand are not typically observed in aggregate data, we focus on the relationship between per capita aggregate consumption and income. Using the optimal consumption functions, we calculate aggregate consumption, C_t, and aggregate income, Y_t, as the averages over all consumers. We report results from two regressions, consumption growth on current income growth and consumption growth on lagged income growth:

$$\Delta \log C_t = \mu + \alpha \, \Delta \log Y_t + \varepsilon_t,$$

$$\Delta \log C_t = \upsilon + \beta \, \Delta \log Y_{t-1} + \epsilon_t.$$

$\hat{\alpha}$ is the estimated aggregate MPC out of current income and $\hat{\beta}$ is the MPC out of lagged income. For these regressions, aggregate consumption and income are derived by aggregating 3,000 consumers using 200 periods.[11] We repeat this process 100 times and report average current and lagged aggregate MPCs across the 100 simulations.

We study the effects of persistence, unemployment, and changes in transitory and permanent uncertainty by changing the parameters of our baseline simulation one at a time. This way we can calculate marginal effects on the aggregate MPCs of changes in persistence and uncertainty. Our ultimate goal is to analyze how the predictions of the aggregated buffer-stock model match the estimated effects of persistence and uncertainty from U.S. state-level data.

Our simulation exercise is similar in spirit to that of Ludvigson and Michaelides (2001). These authors calibrate their income process to match U.S. aggregate income rather than state-level income. The focus of their paper is on explaining excess sensitivity and excess smoothness of aggregate consumption and they do not study the marginal impact of higher persistence or different types of uncertainty on the aggregate MPCs.

3.4 Simulation Results

Table 2 presents results comparing an explicitly aggregated buffer-stock model, as just described, to the closed-form predictions from a log-approximation to a representative-agent PIH-model

[11] Three thousand consumers are enough to maintain the aggregate results. We simulate 215 periods but drop the first 15 periods to guarantee that our results do not depend on the initial conditions. Using a different number of periods would change the standard error of the regressions but not the point estimates, at least not considerably.

(where the representative agent receives the aggregate income process).[12] The table presents the aggregate MPCs out of current and lagged income for both models as well as the average saving rate for the buffer-stock model. The average saving rate is calculated as the average across consumers of the ratio of cash-on-hand minus consumption to cash-on-hand. Standard errors are given in parentheses.[13] An asterisk ($*$) in the table indicates that the corresponding (current or lagged) MPC is significantly different from the corresponding (current or lagged) MPC for the baseline case. A complementary table, Table 3, reports the marginal effects of changing the parameters of our simulations on the aggregate MPCs. The marginal effects are computed as the change in the corresponding MPC relative to the baseline case divided by the change in the parameter being altered. Table 3 includes marginal effects for the aggregated buffer-stock model, for the log-linear approximation of the PIH-model, and for a combination of the two (a half-half combination of buffer-stock and PIH consumers). In Section 4, we compare the sign and size of these marginal effects to our empirical findings.

The baseline simulations have no aggregate transitory shocks and allow for aggregate permanent shocks to be persistent (i.e., $a > 0$). In this case, the PIH predicts an MPC out of current income higher than 1. The predicted MPC out of lagged income is 0 because the agent is not subject to borrowing constraints and is able to adjust consumption immediately following income shocks. 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: the average saving rate is only 6.93 percent. Hence, individuals cannot increase consumption as much as PIH consumers when facing a persistent positive permanent shock and as a result sensitivity to lagged income appears. For our choice of parameters, the MPC out of current income is 1.08 (compared to 1.37 for the PIH-model) and the MPC out of lagged income is 0.106 (compared to 0 in the PIH).

Decreasing persistence to 0—the case where permanent aggregate shocks are i.i.d.—lowers the MPCs out of current and lagged income in the buffer-stock model (from 1.08 to 0.96 and from 0.106 to 0.054 respectively) and the MPC out of current income in the PIH-model (from

[12] The PIH-model is formulated in levels rather than log levels and we make the approximation that the process for the first difference of labor income can be described by the process described for the log difference. We then use the closed-form solution for the response of consumption to income shocks. See Appendix A for details. A similar approximation is made by Ludvigson and Michaelides (2001).

[13] We present the standard error of the *average* MPC reported in the table rather than the standard error of the MPC from a single simulation. Ludvigson and Michaelides (2001) present standard errors for a single simulation.

1.37 to 1). Table 3 shows the marginal effect of increasing persistence on the aggregate MPC out of current income is 0.95 in the PIH-model and 0.33 in the buffer-stock model. The marginal effect of persistence on lagged income is 0 in the PIH-model and 0.27 in the buffer-stock model.

We add unemployment to our simulations using a probability of unemployment of 3 percent. We consider both the case with an unemployment benefit that replaces 30 percent of average income and the case with no unemployment protection. Adding unemployment decreases the MPCs in the buffer-stock model greatly due to higher precautionary saving; not surprisingly, the decrease in the current MPC is more substantial if no income replacement program is present. The average saving rate goes up dramatically from 6.93 percent to 26.74 percent in the case with unemployment benefits and to 45.41 percent in the case with no income replacement. The marginal effect—shown in Table 3—is quite large: a 1 percent change in the probability of job loss will lower the MPC out of current income by 0.06 in the case with no replacement. Changes of this magnitude are probably common over the business cycle. In the PIH-model, the MPCs are not affected.

Next, we examine the marginal effects of changing 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 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 will save less, which might be expected to lead to higher current MPCs.[14] The average saving rate goes down from 6.93 percent to 4.5 percent when reducing the standard deviation of the idiosyncratic permanent shock, σ_N, from 0.10 to 0.05. It declines to just 2.25 percent if the standard deviation of the idiosyncratic transitory shock, σ_V, is reduced from 0.10 to 0.05. However, we do not observe significant differences in the current MPC from the baseline case. This is because with liquidity constraints, less saving implies agents cannot increase consumption more than income in response to a persistent positive permanent shock. This effect would tend to lower the MPCs out of current income and offset the former effect. Moreover, because of lower saving agents are liquidity constrained more often, resulting in a clearly higher MPC out of lagged income. The marginal effects in Table 3 look quite large, but a unit increase in the standard deviation of any of the income shocks corresponds to quite a massive increase in

[14] This is indeed what happens when aggregate permanent shocks are i.i.d. See Table 6.

uncertainty.

Finally, more aggregate uncertainty is introduced by changing the standard deviation of the aggregate shocks one at a time. We *increase* the standard deviation in this case.[15] More aggregate permanent uncertainty results in a higher MPC out of current income in these simulations. Because agents hold more assets due to higher uncertainty (the saving rates goes up to 7.42 percent), they can adjust consumption more promptly in response to persistent positive permanent income shocks, which increases the MPC out of current income. Also, consumers are constrained less often and the MPC out of lagged income decreases slightly. Introducing aggregate transitory uncertainty lowers both MPCs due to higher precautionary saving. In this case, the MPC out of current income decreases in the PIH-model as well. Intuitively, the MPC depends only on the persistence of shocks in the PIH-model and the effect of higher transitory uncertainty comes from the temporary shocks getting larger relative to the persistent shocks, thereby lowering the average persistence of shocks.

We explore the robustness of the results to variations in the parameters. We replicate Table 2 for different values of the risk aversion parameter—$\rho = 1$ and $\rho = 3$—and for the case with no persistence of aggregate shocks. Results are presented in Tables 4, 5, and 6, respectively. Decreasing risk aversion decreases saving, which results in a lower MPC out of current income and a higher MPC out of lagged income. In our simulations, the effect of persistence in the buffer-stock model proves very robust: the higher persistence, the higher the MPCs out of current and lagged income. Regarding our measures of uncertainty, we can summarize our findings as follows. First, with no unemployment benefit, unemployment clearly reduces the current MPC. With unemployment benefits, this is generally the case as well, except when risk aversion is low (the $\rho = 1$ case). The effect of unemployment on the lagged MPC tends to be negative. Second, increases in idiosyncratic uncertainty (transitory or permanent) clearly lower the lagged MPC. The effects on the current MPC are neither large nor robust as they change sign across simulations. Third, higher aggregate transitory uncertainty decreases the MPC out of current income, while higher aggregate *permanent* uncertainty increases it. The effect on the lagged MPC of increasing aggregate transitory uncertainty is negative but not always significant. There is not much of an effect on the lagged MPC when increasing aggregate

[15] The direction of the changes are chosen such that the model satisfies the convergence condition of footnote 8.

permanent uncertainty. Finally, when permanent aggregate shocks are i.i.d., the lagged MPC is very close to zero in all specifications but the effects on the current MPC change little from the persistent case.

In summary, we systematically analyze the effects of persistence and different sources of uncertainty in an aggregated buffer-stock model. For our parameter choices, with income calibrated to match state-level income, the model predicts a clear positive effect of persistence on the current MPC and a positive but smaller effect on the lagged MPC. Unemployment has a negative effect on the MPCs. More idiosyncratic uncertainty, either permanent or transitory, produces a lower lagged MPC, while the effect on the current MPC can go in any direction. More aggregate permanent uncertainty results in a higher current MPC, while there is no clear direction on the effect on the lagged MPC. Finally, more aggregate transitory uncertainty results in a lower current MPC and generally also in a lower lagged MPC. Our findings are potentially important for macroeconomic forecasting because they demonstrate that small changes in uncertainty result in substantial changes in the behavior of aggregate consumption. In particular, a small change in the probability of job loss can significantly change the way consumption reacts to income growth.

The purpose of these simulations is not to replicate the exact size of MPCs out of current and lagged income observed in the data. The MPC out of current income is much larger in our simulated data than its empirical counterpart, while the MPC out of lagged income is generally smaller than what is observed in the data. Ludvigson and Michaelides (2001), Michaelides (2001), and Luengo-Prado (2001) show how incomplete information, habits, and durable goods, respectively, can bring the MPCs for the buffer-stock model closer to their empirical counterparts.

4 Panel-data Estimation of the MPCs

4.1 Econometric Implementation

Let $c_{it} \equiv \Delta \log C_{it}$ denote the growth rate of state-level consumption.[16] In our implementation, we regress c_{it} on income growth, y_{it}, and lagged income growth, y_{it-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.[17] In symbols, we regress $c_{it} - \bar{c}_{.t}$ on $y_{it} - \bar{y}_{.t}$ and $y_{it-1} - \bar{y}_{.t-1}$, respectively, where $\bar{c}_{.t} = \frac{1}{N}\Sigma_{i=1}^{N} c_{it}$ 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. 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): $z_{it} = x_{it} - \bar{x}_{.t} - \bar{x}_{i.} + \bar{x}_{..}$, where $\bar{x}_{i.} = \frac{1}{T}\Sigma_{t=1}^{T} x_{it}$ 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_{it} 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

$$c_{it} = \mu_i + v_t + \alpha y_{it} + \varepsilon_{it} \ ,$$

where the μ_i 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 is measuring the marginal

[16] We tested the level of consumption (non-durable retail sales) for unit roots and could only reject the unit root for 4 and 0 states at the 5 and 1 percent level, respectively, for the 1976–1998 sample (the numbers are 2 and 1, respectively, for the 1964–1998 sample). Therefore, the growth rate of consumption is reasonably modelled as a stationary variable.

[17] Empirically, it matters little if the data are 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.

propensity to consume (MPC).[18] The main focus of our empirical work is to examine if the MPC changes in the way predicted by the theoretical model. If X is a variable that might affect the MPC, we examine if the MPC changes with X by estimating the regression

$$c_{it} = \mu_i + v_t + \alpha y_{it} + \zeta(X_{it} - \bar{X}_{.t})(y_{it} - \bar{y}_{.t} - \bar{y}_{i.} + \bar{y}_{..}) + \varepsilon_{it}.$$

In this regression, the MPC is $\alpha + \zeta(X_{it} - \bar{X}_{.t})$ where the time-specific average of X_{it} is subtracted in order to remove U.S.-wide aggregate effects. The term $\zeta(X_{it} - \bar{X}_{.t})$ is multiplied by $(y_{it} - \bar{y}_{.t} - \bar{y}_{i.} + \bar{y}_{..})$ rather than just y_{it} because the inclusion of the fixed effects in the regression implies the term y_{it} multiplying α *de facto* has the form $(y_{it} - \bar{y}_{.t} - \bar{y}_{i.} + \bar{y}_{..})$ and since we are interested in how α changes as a function of X, the income-term multiplying X has to have the same form. 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 whole point of the exercise is to gauge 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 average was subtracted.

In our implementation, we will often include more than one interaction variable and each of them will be treated as explained here. 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 1976–1998. State-level disposable labor income is constructed from BEA data as described in Section 2. 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*).[19] We transform retail sales and labor income to per capita terms using population data from the BEA and deflate them using the Consumer Price Index from the Bureau of Labor Statistics (BLS).

[18] Including time-fixed effects implies that we are measuring the effect on state-specific consumption of state-specific changes in income. If there are systematic differences between the time-series processes for aggregate and state-specific income, we implicitly assume this difference is not observed by consumers.

[19] Retail sales are a somewhat noisy proxy for state-level private consumption but no better data seem to exist. The retail sales data are available from 1963–1998. The correlation between annual growth rates of aggregate U.S. total (nondurable) retail sales and aggregate U.S. total (nondurable) private consumption from the National Income and Product Accounts, both measured in real terms and per capita, is 0.84 (0.65).

4.2 Selection of Regressors

We turn to the empirical estimation of the MPCs as functions of state-level variables. We use as interaction terms in the MPC-estimations variables that approximate the parameters of the theoretical model.

Persistence of aggregate shocks. Our measure of the persistence of aggregate shocks in state i is the estimated parameter $\hat{a}_i$.

Aggregate uncertainty. We have two estimated parameters of aggregate uncertainty: the standard error of the innovation to the persistent component of aggregate income $\hat{\sigma}_{Gi}$ and the standard error of the innovation to the temporary component of aggregate income $\hat{\sigma}_{Wi}$.

Individual-level income volatility. We use the share of farmers in a state. Farmers are subject to substantially higher temporary income uncertainty than other income groups as documented in table 4 in Carroll and Samwick (1997). In other words, farmers may be particularly subject to the type of uncertainty that is captured by the parameter σ_V in the model. In our simulations, all agents at the disaggregated level are subject to the same stochastic process for uncertainty and the simulations revealed that the aggregate MPC out of lagged income will be lower when σ_V is higher. Based on this result, we will examine if the MPC out of lagged income is lower in states where a relatively large number of consumers can be expected to have high variance of temporary 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). As discussed earlier, agricultural states may also have a high level of aggregate uncertainty and one might question if the share of agricultural employees can then be interpreted as a measure of individual-level uncertainty. The interpretation is, however, valid in a multiple regression that also includes a separate measure of aggregate uncertainty—in our case, the measure discussed above.

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 temporary uncertainty—see table 4 in Carroll and Samwick (1997). One may also hypothesize that government employees are less likely to be subject to

permanent idiosyncratic shocks (captured by σ_N in the model) but since the predicted impact of temporary and permanent idiosyncratic shocks on the aggregate MPCs are similar, it is not important for our current purpose to sort this out.

We obtain the total number of employees (including proprietors and partners) by state, as well as the number of employees in farming and government, from the BEA.

Individual-level probability of job loss. Our final indicator of individual-level uncertainty is the unemployment rate. When the unemployment rate is high, the risk of job loss is higher—this is almost true by definition—and we assume that this is associated with a higher risk of "catastrophic" income loss. In practice, the risk of job loss is quite unevenly distributed across the population and not necessarily well captured by the unemployment rate. On the other hand, casual empiricism suggests that periods of high unemployment are periods of high income uncertainty, so it is well-motivated to examine if high unemployment affects the MPCs in the direction suggested by our model. Because we focus on the potential income loss from job separation, we multiply the state-level unemployment rate by one minus the state's average unemployment insurance replacement rate; for brevity, we use the term unemployment.

State-level unemployment rates—available since 1976—are from the BLS. State-level income replacement rates are from the Unemployment Insurance Financial Data Handbook published by the U.S. Department of Labor.

Correlation matrix for regressors. Table 7 presents the correlations of our regressors: unemployment, the share of farmers in total employment, the share of government employment, and the estimated persistence and standard deviations of aggregate permanent and transitory shocks to income found in Section 2. Unemployment has low correlation with the other regressors, while the share of agriculture is strongly correlated with persistence of income as well as with both parameters for aggregate uncertainty. The share of government is not highly correlated with other regressors, while persistence is highly negatively correlated with the variance of the permanent shocks. Finally, we observe a high positive correlation between the standard deviations of permanent and temporary shocks.

4.3 Empirical results of the panel-data estimations

Current MPC. From Table 8, the MPC from a simple regression of consumption growth on current income growth and persistence is 0.37. This is clearly lower than the coefficients near 1 found in Table 2; our benchmark version of the buffer-stock model is, therefore, not the full truth. This is, of course, the well known "excess smoothness" result.

Our main focus is on the interaction effects. We experimented with the specification, but for brevity we do not tabulate the results of all permutations of the regressors. The main conclusions are quite clear though. The effect of persistence is estimated robustly with very large t-statistics and this variable is, therefore, included in all of columns (1)–(6). The estimated value of the coefficient to persistence is larger than 1 in all regressions. This corresponds to a large economic impact. For example, using the estimated values in column (1), the MPC in Iowa (with persistence 0) is predicted to be about 0.37 while the MPC in Oklahoma (with persistence 0.78) is predicted to be 1.20. No other regressors are robustly or significantly estimated at the conventional 5 percent level. We illustrate this by including each of them one-by-one.

The coefficient to unemployment, in column (2), is numerically quite large and negative as expected, but it is not statistically significant. In column (3), the MPC depends positively on the share of agriculture in the state. The t-statistic is 1.58 which is not significant at the conventional 5 percent level, and maybe more telling is that the coefficient is not robust to (untabulated) permutations in the specification. From column (4), the MPC out of current income is not dependent on the share of government employees in the state—the coefficient estimate is small and clearly insignificant. In columns (5) and (6), we include the estimated standard deviations for the components of aggregate income. The point estimates are both positive but not significant. While the permanent shock has the right sign, the transitory component does not. The t-statistics are not very far from the 5 percent significance level but the estimates are very fragile—for example, including both estimated standard deviations together makes the estimates very large with opposite signs. It may be that aggregate uncertainty is not important for the MPC out of current income but a more reasonable interpretation may be that our econometric tools are not sharp enough to settle this issue.[20]

[20] Since these regressors are based on interaction variables from an initial regression, they are obviously measured with error, which results in bias of the estimates. It is fairly straightforward to adjust for these problems. This,

Lagged MPC. Table 9 examines the same specification in terms of the MPC out of lagged income ("excess sensitivity"). The sensitivity of consumption to lagged income is clearly and robustly lower in agricultural states. The estimated value is around –5, implying that a 10 percent increase in agricultural employment can be expected to lower the MPC out of lagged income by 0.5. The order of magnitude and the level of significance is very robustly estimated and we include the share of agriculture in all specifications. In column (2), we add the share of government employees which is also robustly significant with the expected positive sign. The effect is very similar in absolute magnitude to that found for agriculture. Column (3) includes unemployment but not the share of government employees, and column (4) includes interaction terms for both unemployment and government employees. Unemployment is strongly significant with the expected negative sign. It is also robustly estimated with a coefficient of around –10 which implies that an increase in the unemployment rate of 1 percentage point will lower the MPC by around 0.1. Because agricultural employment share, government employment share, and unemployment are all robustly and significantly estimated, we keep these three variables in columns (5)–(7) where we examine the impact of the remaining regressors. Column (5) reports the effect of persistence. Surprisingly, given the strong effect of persistence in the current MPC, the MPC out of lagged income is not different in states with high vs. low persistence of aggregate shocks. The last two columns examine the effect of aggregate uncertainty. We find clearly insignificant effects of both aggregate standard deviations. These regressors both have the expected negative sign but they are not robustly estimated.

4.4 Match of Theory and Data.

Match for the MPC out of current income. Comparing the empirical findings of the panel-data estimations in Table 8 with the predictions of the buffer-stock (and PIH) model in Table 3, the buffer-stock model and the data strongly agree on a clear effect of persistency with the correct sign: the more persistence, the higher the MPC. The coefficient found in the data is larger than predicted and it is also slightly larger than predicted by the PIH-model. Nonetheless, the results are fairly consistent with a model in which some consumers are PIH consumers and some consumers follow the buffer-stock model. We are not able to significantly match the

however, will not change our conclusions since the coefficients to these variables are not robustly estimated.

predicted effect of unemployment although the sign of the estimated coefficient is as anticipated. Further, the estimations clearly miss the effects of transitory aggregate uncertainty, while the insignificant indicators of individual-level uncertainty (government and agricultural employment, respectively) are consistent with the near-zero effects predicted by the model.

Match for the MPC out of lagged income. A comparison of the empirical results in Table 9 with the theoretical predictions for the lagged MPC reveals a very good fit. The buffer stock model predicts a small effect of persistence—even smaller if a significant fraction of consumers are well-described by the PIH-model as Gourinchas and Parker (2001) suggest—and we empirically find no effect. Aggregate uncertainty is predicted to have a small effect and we found none. Empirically, we found a significant negative effect of unemployment in line with the model. The estimated effect is much larger than predicted, but measured unemployment is not directly a measure of the probability of job loss so it is not really meaningful to compare the orders of magnitude. This is also the case for the other indicators of individual-level uncertainty. Agricultural and government shares in employment are strongly significant with the predicted signs. Overall, the buffer-stock model is very successful in predicting the differences in the lagged MPCs across states. The PIH-model, on the contrary, has the implication that these are all equal to zero!

5 Conclusion

The contributions of our paper are both theoretical and empirical. Based on simulating suitably calibrated versions of the buffer-stock model, we find that the model predicts large effects on the marginal propensities to consume out of current and lagged shocks to labor income from persistence of aggregate shocks, aggregate uncertainty, and individual-level uncertainty. To the best of our knowledge, a systematic quantification of the marginal impact of these economic variables has not been done for this model.

Using panel-data regressions, we document varying propensities to consume across U.S. states. More importantly, the propensities to consume vary in ways explained by a suitably calibrated buffer-stock model. The good match of the positive predictions of the model with aggregate data suggests that it is important to use this model to supplement the standard PIH-

model in macroeconomic forecasting. Our results are derived from cross-state differences in consumer behavior but similar results will likely hold in aggregate data as uncertainty changes over the business cycle. A challenging topic for future research is to document this conjecture—this will be challenging because the limited number of aggregate business cycles provides few degrees of freedom for econometricians.

References

Abowd, J., and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–445.

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

Attanasio, O., and G. Weber (1993): "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies*, 60, 631–649.

Blinder, A., and A. Deaton (1985): "The Time Series Consumption Function Revisited," *Brookings Papers on Economic Activity*, 1985:2, 465–521.

Browning, M., and A. Lusardi (1996): "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 34, 1797–1855.

Caballero, R. (1991): "Earnings Uncertainty and Aggregate Wealth Accumulation," *American Economic Review*, 81, 859–871.

Campbell, J. Y., and A. Deaton (1989): "Why is Consumption so Smooth?," *Review of Economic Studies*, 56, 357–374.

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

Carroll, C. D., K. Dynan, and S. C. Krane (2003): "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets," *Review of Economic and Statistics*, 85, 586–604.

Carroll, C. D., and M. S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64, 981–992.

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.

Dynan, K. (1993): "How Prudent are Consumers?," *Journal of Political Economy*, 101, 1104–1113.

Engen, E. M., and J. Gruber (2001): "Unemployment Insurance and Precuationary Saving," *Journal of Monetary Economics*, 47, 545–579.

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 (2001): "Consumption over the Life Cyle," *Econometrica*, 70, 47–91.

Hahm, J.-H. (1999): "Consumption Growth, Income Growth and Earnings Uncertainty: Simple Cross-Country Evidence," *International Economic Journal*, 13, 39–58.

Hahm, J.-H., and D. Steigerwald (1999): "Consumption Adjustment under Time-Varying Uncertainty," *Review of Economic and Statistics*, 81, 32–40.

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., and T. Sargent (1981): "A Note on Wiener-Kolmogorov Predition Formulas for Rational Expectations Models," *Economics Letters*, 8, 255–260.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy*, 40, 59–125.

Jappelli, T. (1990): "Who is Credit Constrained in the U.S. Economy?," *Quarterly Journal of Economics*, 105, 219–235.

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

Leland, H. (1968): "Saving and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics*, 82, 465–473.

Ludvigson, S., and A. Michaelides (2001): "Can Buffer Stock Saving Explain the Consumption Excesses?," *American Economic Review*, 91, 631–647.

Luengo-Prado, M. J. (2001): "Durables, Nondurables, Down Payments and Consumption Excesses," Mimeo.

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.

McCarthy, J. (1995): "Imperfect Insurance and Differing Propensities to Consume Across Households," *Journal of Monetary Economics*, 36, 301–327.

Michaelides, A. (2001): "Buffer Stock Saving and Habit Formation," Mimeo.

——— (2003): "A Reconciliation of Two Alternative Approaches towards Buffer Stock Saving," *Economics Letters*, 79, 137–143.

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.

Perraudin, W., and B. E. Sørensen (1992): "The Credit-Constrained Consumer: An Empirical Study of Demand and Supply in the Loan Market," *Journal of Business & Economic Statistics*, 10, 179–193.

Pischke, J.-S. (1995): "Individual Income, Incomplete Information and Aggregate Consumption," *Econometrica*, 63, 805–840.

Skinner, J. (1988): "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics*, 22, 237–255.

Sørensen, B. E., and O. Yosha (2000): "Intranational and International Credit Market Integration: Evidence from Regional Income and Consumption Patterns," in *Intranational Macroeconomics*, ed. by Gregory Hess and Eric van Wincoop. Cambridge University Press, New York, NY.

Tauchen, G. (1986): "Finite State Markov Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

Zeldes, S. (1989a): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 142, 305–346.

——— (1989b): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104, 275–298.

A Appendix. The MPC in the PIH case

We show how to calculate the approximate MPC out of current income for the PIH-model in Table 2. The MPC is defined as $\frac{cov(\Delta C_t, \Delta Y_t)}{var(\Delta Y_t)}$. Let $g_t = \log G_t$ and $w_t = \log W_t$. Given the assumptions about the income process for labor income, $\Delta \log Y_t = g_t + w_t - w_{t-1}$ and $g_t = a g_{t-1} + \varepsilon_t$. Since the PIH is formulated in levels rather than log levels, we assume the process for the first difference of labor income can be approximately described by the process for the log difference $\Delta Y_t = g_t + w_t - w_{t-1} = a\Delta Y_{t-1} + \varepsilon_t + w_t - (1+a)w_{t-1} + a w_{t-2}$, which implies $var(\Delta Y_t) = \frac{\sigma_\varepsilon^2 + 2(1+a+a^2)\sigma_W^2}{1-a^2}$. Let $u_t = \varepsilon_t + w_t$. Hansen and Sargent (1981) and Deaton (1992) show that if income can be represented by the ARMA process $a(L)Y_t = b(L)u_t$, the PIH predicts that: $\Delta C_t = \frac{r}{1+r} u_t \times b\left(\frac{1}{1+r}\right) / a\left(\frac{1}{1+r}\right)$.

The ARMA representation for our income process: $(1-aL)(1-L)Y_t = u_t - (1+a)w_{t-1} + a w_{t-2}$ is going to be of the form: $(1-aL)(1-L)Y_t = (1 + b_1 L + b_2 L^2)u_t$. Solve for b_1 and b_2. Define $Z_t \equiv (1-aL)(1-L)Y_t$. $E(Z_t Z_{t-2}) = a\sigma_W^2 = b_2 \sigma_U^2$, which implies $b_2 = \frac{a\sigma_W^2}{\sigma_U^2}$. $E(Z_t Z_{t-1}) = [-(1+a) - (1+a)a]\sigma_W^2 = (b_1 + b_1 b_2)\sigma_U^2$. Plugging in for b_2, $b_1 = \frac{-(1+a)^2 \sigma_W^2}{\sigma_U^2 + a\sigma_W^2}$, where $\sigma_U^2 = \sigma_\varepsilon^2 + \sigma_W^2$. Finally, $\Delta C_t = \frac{1+r}{1+r-a}\left[1 + \frac{b_1}{1+r} + \frac{b_2}{(1+r)^2}\right] u_t$ and $cov(\Delta C_t, \Delta Y_t) = \frac{1+r}{1+r-a}\left[1 + \frac{b_1}{1+r} + \frac{b_2}{(1+r)^2}\right] \sigma_U^2$.

Table 1: Parameters of Time Series Process for State-Level Disposable Labor Income.

	(1) Persistence	(2) σ_G	(3) σ_W
Alabama	0.48 (0.15)	2.21 (0.26)	0.00 (0.68)
Alaska	0.51 (0.22)	4.05 (0.92)	0.84 (1.15)
Arizona	0.46 (0.15)	2.37 (0.28)	0.00 (0.82)
Arkansas	0.25 (0.16)	3.44 (0.40)	0.00 (1.91)
California	0.54 (0.20)	1.70 (0.35)	0.51 (0.30)
Colorado	0.46 (0.15)	1.77 (0.21)	0.00 (0.59)
Connecticut	0.43 (0.15)	2.06 (0.24)	0.00 (0.56)
Delaware	0.46 (0.15)	2.19 (0.25)	0.00 (0.52)
Florida	0.45 (0.15)	2.30 (0.27)	0.00 (0.56)
Georgia	0.45 (0.15)	2.37 (0.28)	0.00 (0.91)
Hawaii	0.81 (0.12)	1.57 (0.35)	0.85 (0.21)
Idaho	0.32 (0.39)	2.99 (1.15)	1.46 (0.87)
Illinois	0.38 (0.25)	2.28 (0.51)	0.40 (0.82)
Indiana	0.22 (0.16)	3.27 (0.38)	0.00 (1.46)
Iowa	0.00 (0.32)	4.11 (1.72)	2.10 (1.75)
Kansas	0.27 (0.32)	2.83 (0.79)	0.27 (2.80)
Kentucky	0.30 (0.15)	2.51 (0.29)	0.00 (1.43)
Louisiana	0.47 (0.26)	1.92 (0.50)	0.42 (0.64)
Maine	0.36 (0.29)	2.53 (0.66)	0.18 (2.82)
Maryland	0.57 (0.18)	1.96 (0.38)	0.19 (0.84)
Massachusetts	0.47 (0.14)	2.15 (0.25)	0.00 (0.62)
Michigan	0.40 (0.15)	3.10 (0.36)	0.00 (0.62)
Minnesota	0.10 (0.16)	3.50 (0.41)	0.00 (4.19)
Mississippi	0.39 (0.15)	2.70 (0.31)	0.00 (0.78)
Missouri	0.26 (0.16)	2.77 (0.32)	0.00 (1.97)
Montana	–0.00 (0.22)	3.72 (0.41)	0.48 (2.24)
Nebraska	–0.17 (0.40)	4.43 (2.14)	0.72 (7.41)
Nevada	0.43 (0.24)	2.14 (0.52)	0.75 (0.45)
New Hampshire	0.32 (0.16)	2.63 (0.31)	0.00 (0.83)
New Jersey	0.30 (0.29)	2.13 (0.56)	0.38 (1.02)
New Mexico	0.42 (0.15)	1.93 (0.22)	0.00 (0.81)
New York	0.48 (0.23)	1.79 (0.45)	0.77 (0.33)
North Carolina	0.39 (0.15)	2.53 (0.29)	0.00 (1.14)
North Dakota	0.00 (0.28)	8.51 (3.32)	4.56 (3.31)
Ohio	0.42 (0.15)	2.32 (0.27)	0.00 (0.65)
Oklahoma	0.78 (0.24)	1.11 (0.58)	1.11 (0.31)
Oregon	0.42 (0.22)	2.61 (0.53)	0.08 (4.72)
Pennsylvania	0.59 (0.14)	1.72 (0.20)	0.00 (0.31)
Rhode Island	0.46 (0.15)	2.05 (0.24)	0.00 (0.74)
South Carolina	0.61 (0.19)	2.01 (0.43)	0.45 (0.45)
South Dakota	0.00 (0.16)	6.08 (1.71)	2.46 (2.11)
Tennessee	0.43 (0.15)	2.55 (0.30)	0.00 (0.59)
Texas	0.41 (0.39)	1.94 (0.77)	0.51 (0.93)
Utah	0.53 (0.14)	1.83 (0.21)	0.00 (0.30)
Vermont	0.27 (0.16)	2.73 (0.32)	0.00 (5.20)
Virginia	0.53 (0.16)	2.07 (0.24)	0.00 (0.50)
Washington	0.49 (0.15)	2.03 (0.24)	0.00 (0.56)
West Virginia	0.56 (0.18)	2.13 (0.39)	0.26 (0.66)
Wisconsin	0.46 (0.15)	2.25 (0.26)	0.00 (0.50)
Wyoming	0.54 (0.14)	2.66 (0.31)	0.00 (0.55)

Notes: The table displays the estimated parameters of a time series model for real disposable labor income. Let y_{it} be the log of per capita disposable labor income (deflated by the CPI) in state i. The model is: $y_{it} = \mu_i + \log G_{it} + \sigma_{Wi}\,(W_{it} - W_{it-1})$, where μ_i is a constant for each state, W_{it} (an i.i.d. innovation to the level of y_{it}) is normally distributed and independent of G_{it}, and $\log G_{it} = a_i \log G_{it-1} + \sigma_{Gi}\,\epsilon_{it}$ where ϵ_{it} are i.i.d. normal innovations. a_i is persistence, σ_{Gi} is the standard deviation of the permanent component and σ_{Wi} is the standard deviation of the transitory component. The table reports the estimates of a_i in column (1), 100 times σ_{Gi} in column (2), and 100 times σ_{Wi} in column (3). The model is estimated by Maximum Likelihood using a Kalman Filter. Standard errors in parentheses. Sample 1964–1998.

Table 2: Sensitivity to Current and Lagged Income in Simulated Data

	Buffer-Stock			PIH	
	MPC		Average Saving	MPC	
	Current	Lagged	Rate	Current	Lagged
Baseline	1.076	0.106	6.93%	1.37	0
	(0.003)	(0.008)	(0.016)		
No Persistence	0.946*	–0.001*	6.63%	1	0
	(0.001)	(0.007)	(0.002)		
Unemployment					
Replacement	1.015*	0.054*	26.74%	1.37	0
	(0.003)	(0.008)	(0.028)		
No replacement	0.893*	0.063*	45.41%	1.37	0
	(0.003)	(0.007)	(0.028)		
Less idiosyncratic uncertainty					
$\sigma_N = 0.05$	1.077	0.157*	4.50%	1.37	0
	(0.002)	(0.008)	(0.009)		
$\sigma_V = 0.05$	1.072	0.160*	2.25%	1.37	0
	(0.002)	(0.008)	(0.008)		
More aggregate uncertainty					
$\sigma_G = 0.035$	1.097*	0.103	7.42%	1.37	0
	(0.002)	(0.008)	(0.023)		
$\sigma_W = 0.004$	1.046*	0.099	6.94%	1.23	0
	(0.003)	(0.008)	(0.016)		

Notes: For the columns "Buffer-Stock", the column labelled "MPC" and "Current" reports the estimated value of the parameter α from the regression: $\Delta \log C_t = \mu + \alpha \Delta \log Y_t + \varepsilon_t$. The column labelled "MPC" and "Lagged" reports the estimated value of the parameter β for lagged income from the regression $\Delta \log C_t = \upsilon + \beta \Delta \log Y_{t-1} + \epsilon_t$. Consumption is simulated as described in the text. The simulated data are based on the individual-level income process $\Delta \log Y_{jt} = \log G_t + \log W_t - \log W_{t-1} + \log N_{jt} + \log V_{jt} - \log V_{j,t-1}$, where $\log G_t$ is an AR(1) process with persistence a, unconditional mean μ_G and standard deviation σ_G. $\log N_{jt}$, $\log V_{jt}$ and $\log W_t$ are independent and identically normally distributed with standard deviations σ_N, σ_V and σ_W, and means $-\sigma_N^2/2$, $-\sigma_V^2/2$ and $-\sigma_W^2/2$, respectively. Baseline parameters: $\mu_G = 0.0165$, $\sigma_G = 0.0265$, $a = 0.39$. $\sigma_W = 0$ and $\sigma_N = \sigma_V = 0.1$. The interest rate is 2 percent and the discount rate is 5 percent. The coefficient of risk aversion, ρ, is 2 and the unemployment probability is 0. In the case with unemployment, $p = 0.03$ and the replacement rate is 30 percent when present. Aggregate consumption and income are averages over 3,000 simulated individuals. The length of the sample for the regressions is 200 periods. We report MPCs averaged over 100 simulations. Standard errors are in parentheses. "*" is added when the current or lagged MPC is significantly different from the corresponding MPC for the baseline case. The column "Average Saving Rate" presents the average saving rate across consumers— the saving rate is the ratio of cash-on-hand minus consumption to cash-on-hand. For the columns labelled "PIH" the numbers are calculated using 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.

Table 3: Sensitivity to Current and Lagged Income in Simulated Data. Marginal Effects

	Buffer-Stock		PIH		Combined	
	Current	Lagged	Current	Lagged	Current	Lagged
Persistence	0.33	0.27	0.95	0	0.64	0.14
Unemployment						
Replacement	–2.05	–1.71	0	0	–1.03	–0.86
No replacement	–6.12	–1.43	0	0	–3.06	–0.71
Idiosyncratic uncertainty						
permanent	–0.01	–1.02	0	0	0.00	–0.51
transtiory	0.09	–1.09	0	0	0.05	–0.55
Aggregate uncertainty						
permanent	2.39	–0.29	0	0	1.19	–0.14
transitory	–7.52	–1.60	–35.0	0	–21.26	–0.80

Notes: The columns present marginal effects calculated using the MPCs in Table 2. The column headings for the first four columns correspond to those of Table 2. Each marginal effect is calculated as the change in the corresponding MPC relative to the baseline case divided by the change in the parameter being altered. The table includes marginal effects for the simulated buffer-stock model, the log-linear approximation of a representative-agent PIH-model, and a combination of the two. The columns "Combined" present the case where half the agents are PIH consumers and half are buffer-stock consumers.

Table 4: Sensitivity to Current and Lagged Income in Simulated Data. $\rho = 1$.

	Buffer-Stock					PIH	
	MPC		Marginal Effect		Average Saving	MPC	Marginal Effect
	Current	Lagged	Current	Lagged	Rate	Current	Current
Baseline	1.037	0.176	–	–	3.14%	1.37	–
	(0.002)	(0.007)			(0.006)		
No persistence	0.961*	–0.003*	0.19	0.45	3.05%	1	0.95
	(0.001)	(0.007)			(0.001)		
Unemployment							
replacement	1.048*	0.056*	0.36	–3.99	13.61%	1.37	0.00
	(0.003)	(0.008)			(0.019)		
no replacement	0.925*	0.077*	–3.72	–3.29	30.80%	1.37	0.00
	(0.003)	(0.007)			(0.018)		
Less idiosyncratic uncertainty							
$\sigma_N = 0.05$	1.044*	0.216*	–0.13	–0.80	2.52%	1.37	0.00
	(0.001)	(0.007)			(0.005)		
$\sigma_V = 0.05$	1.025*	0.266*	0.24	–1.81	0.72%	1.37	0.00
	(0.001)	(0.007)			(0.002)		
More aggregate uncertainty							
$\sigma_G = 0.035$	1.048*	0.184	1.25	0.93	3.25%	1.37	0.00
	(0.002)	(0.008)			(0.008)		
$\sigma_W = 0.004$	1.017*	0.163	–5.12	–3.23	3.14%	1.23	–35.00
	(0.002)	(0.007)			(0.006)		

Notes: For the columns "Buffer-Stock", the column labelled "MPC" and "Current" reports the estimated value of the parameter α from the regression: $\Delta \log C_t = \mu + \alpha \Delta \log Y_t + \varepsilon_t$. The column labelled "MPC" and "Lagged" reports the estimated value of the parameter β for lagged income from the regression $\Delta \log C_t = \upsilon + \beta \Delta \log Y_{t-1} + \epsilon_t$. Consumption is simulated as described in the text. The simulated data are based on the individual-level income process $\Delta \log Y_{jt} = \log G_t + \log W_t - \log W_{t-1} + \log N_{jt} + \log V_{jt} - \log V_{j,t-1}$, where $\log G_t$ is an AR(1) process with persistence a, unconditional mean μ_G and standard deviation σ_G. $\log N_{jt}$, $\log V_{jt}$ and $\log W_t$ are independent and identically normally distributed with standard deviations σ_N, σ_V and σ_W, and means $-\sigma_N^2/2$, $-\sigma_V^2/2$ and $-\sigma_W^2/2$, respectively. Baseline parameters: $\mu_G = 0.0165$, $\sigma_G = 0.0265$, $a = 0.39$. $\sigma_W = 0$ and $\sigma_N = \sigma_V = 0.1$. The interest rate is 2 percent and the discount rate is 5 percent. The coefficient of risk aversion, ρ, is 1 and the unemployment probability is 0. In the case with unemployment, $p = 0.03$ and the replacement rate is 30 percent when present. Aggregate consumption and income are averages over 3,000 simulated individuals. The length of the sample for the regressions is 200 periods. We report MPCs averaged over 100 simulations. Standard errors are in parentheses. The columns labelled "Marginal Effect" present marginal effects calculated using the MPCs in the previous two columns. Each marginal effect is calculated as the change in the corresponding MPC relative to the baseline case divided by the change in the parameter being altered. "*" is added when the current or lagged MPC is significantly different from the corresponding MPC for the baseline case. The column "Average Saving Rate" presents the average saving rate across consumers—the saving rate is the ratio of cash-on-hand minus consumption to cash-on-hand. For the columns labelled "PIH" the numbers are calculated using 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. The marginal effects are calculated analogously.

Table 5: Sensitivity to Current and Lagged Income in Simulated Data. $\rho = 3$.

	Buffer-Stock					PIH	
	MPC		Marginal Effect		Average Saving	MPC	Marginal Effect
	Current	Lagged	Current	Lagged	Rate	Current	Current
Baseline	1.137	0.041	–	–	13.72%	1.37	–
	(0.003)	(0.009)			(0.041)		
No persistence	0.936*	–0.002*	0.51	0.11	12.52%	1	0.95
	(0.001)	(0.007)			(0.007)		
Unemployment							
replacement	1.040*	0.027	–3.23	–0.46	37.70%	1.37	0.00
	(0.004)	(0.008)			(0.052)		
no replacement	0.892*	0.044	–8.17	0.09	55.95%	1.37	0.00
	(0.003)	(0.007)			(0.040)		
Less idiosyncratic uncertainty							
$\sigma_N = 0.05$	1.090*	0.132*	0.94	–1.82	6.00%	1.37	0.00
	(0.002)	(0.008)			(0.013)		
$\sigma_V = 0.05$	1.147*	0.074*	–0.21	–0.67	0.27%	1.37	0.00
	(0.003)	(0.009)			(0.002)		
More aggregate uncertainty							
$\sigma_G = 0.035$	1.177*	0.029	4.69	–1.35	16.73%	1.37	0.00
	(0.003)	(0.009)			(0.074)		
$\sigma_W = 0.004$	1.098*	0.039	–9.61	–0.47	13.74%	1.23	–35.00
	(0.003)	(0.008)			(0.043)		

Notes: See notes to Table 4. In this table $\rho = 3$.

Table 6: Sensitivity to Current and Lagged Income in Simulated Data. The i.i.d Case

	Buffer-Stock					PIH	
	MPC		Marginal Effect		Average Saving	MPC	Marginal Effect
	Current	Lagged	Current	Lagged	Rate	Current	Current
Baseline	0.946	–0.001	–	–	6.63%	1	–
	(0.001)	(0.007)			(0.00002)		
Unemployment							
replacement	0.842*	0.008	–3.44	0.28	26.42%	1	0.00
	(0.001)	(0.006)			(0.00008)		
no replacement	0.729*	0.019*	–7.22	0.67	45.17%	1	0.00
	(0.002)	(0.005)			(0.00011)		
Less idiosyncratic uncertainty							
$\sigma_N = 0.05$	0.970*	0.004	–0.49	–0.08	4.40%	1	0.00
	(0.001)	(0.007)			(0.00001)		
$\sigma_V = 0.05$	0.983*	–0.002	–0.75	0.02	2.09%	1	0.00
	(0.0004)	(0.007)			(0.00001)		
More aggregate uncertainty							
$\sigma_G = 0.035$	0.957*	0.001	1.31	0.17	6.85%	1	0.00
	(0.001)	(0.007)			(0.00003)		
$\sigma_W = 0.004$	0.921*	–0.003	–6.16	–0.70	6.64%	0.96	–10.00
	(0.001)	(0.007)			(0.00004)		

Notes: See notes to Table 4. In this table $a = 0$ and $\rho = 2$.

Table 7: Correlation Matrix of Regressors

	(1)	(2)	(3)	(4)	(5)	(6)
(1) Unemployment Rate	1.00	–0.39	0.20	0.35	–0.29	–0.27
(2) Farm Share		1.00	0.05	–0.66	0.73	0.64
(3) Govt. Share			1.00	0.29	0.07	0.18
(4) Persistence				1.00	–0.75	–0.40
(5) $\hat{\sigma}_G$					1.00	0.76
(6) $\hat{\sigma}_W$						1.00

Notes: The table shows how the variables correlate across the 50 U.S. states. "Unemployment Rate" is the average unemployment rate for each state multiplied by one minus the state's average unemployment insurance replacement rate over the period 1976–1998. "Farm Share" is the ratio of employees (including proprietors) in farming to the total number of employees in each state over the same sample. "Govt. Share" is defined similarly. "Persistence" refers, for each state, to the number in the column "Persistence" in Table 1, while the standard deviations in rows (5) and (6) are the numbers given in columns (2) and (3) of Table 1, respectively.

Table 8: Sensitivity to Current Income: Non-Durable Retail Sales

	(1)	(2)	(3)	(4)	(5)	(6)
μ_i	yes	yes	yes	yes	yes	yes
v_t	yes	yes	yes	yes	yes	yes
y_{it}	0.37	0.37	0.36	0.37	0.35	0.34
	(7.90)	(7.95)	(7.24)	(7.44)	(7.16)	(6.87)
Interaction terms:						
Persistence	1.06	1.17	1.36	1.06	1.33	1.33
	(7.36)	(6.24)	(5.56)	(7.27)	(5.62)	(6.48)
Unemployment Rate	-	–3.27	-	-	-	-
	-	(0.90)	-	-	-	-
Farm share	-	-	1.98	-	-	-
	-	-	(1.53)	-	-	-
Govt. share	-	-	-	0.13	-	-
	-	-	-	(0.11)	-	-
$\hat{\sigma}_G$	-	-	-	-	2.71	-
	-	-	-	-	(1.42)	-
$\hat{\sigma}_W$	-	-	-	-	-	4.33
	-	-	-	-	-	(1.85)

Notes: Model: $c_{it} = \mu_i + v_t + \alpha y_{it} + \zeta(X_{it} - \bar{X}_{.t})(y_{it} - \bar{y}_{.t} - \bar{y}_{i.} + \bar{y}_{..}) + \varepsilon_{it}$. y_{it} is state i's labor income growth (real and per capita). c_{it} is state i's nondurable consumption growth (real and per capita). μ_i 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—see Table 7 for definitions. t-statistics in parentheses. Sample 1976–1998.

Table 9: Sensitivity to Lagged Income: Non-Durable Retail Sales

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
μ_i	yes	yes	yes	yes	yes	yes	yes
v_t	yes	yes	yes	yes	yes	yes	yes
y_{it-1}	0.27	0.21	0.27	0.21	0.21	0.21	0.21
	(6.42)	(4.85)	(6.59)	(4.80)	(4.40)	(4.86)	(4.81)
Interaction terms:							
Farm share	–4.06	–3.63	–5.70	–5.78	–5.78	–4.75	–5.01
	(5.55)	(4.94)	(5.78)	(5.92)	(4.37)	(3.13)	(3.29)
Govt. share	-	3.32	-	4.01	3.98	4.40	4.25
	-	(3.85)	-	(4.72)	(4.53)	(4.29)	(4.34)
Unemployment Rate	-	-	–9.16	–12.68	–12.43	–13.00	–12.77
	-	-	(2.43)	(3.34)	(3.21)	(3.36)	(3.33)
Persistence	-	-	-	-	–0.01	-	-
	-	-	-	-	(0.04)	-	-
$\hat{\sigma}_G$	-	-	-	-	-	–2.00	-
	-	-	-	-	-	(0.85)	-
$\hat{\sigma}_W$	-	-	-	-	-	-	–2.05
	-	-	-	-	-	-	(0.64)

Notes: Model: $c_{it} = \mu_i + v_t + \alpha y_{i,t-1} + \zeta(X_{it} - \bar{X}_{.t})(y_{i,t-1} - \bar{y}_{.,t-1} - \bar{y}_{i.} + \bar{y}_{..}) + \varepsilon_{it}$. $y_{i,t-1}$ is state i's lagged labor income growth (real and per capita). c_{it} is state i's nondurable consumption growth (real and per capita). μ_i 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—see Table 7 for definitions. t-statistics in parentheses. Sample 1976–1998.

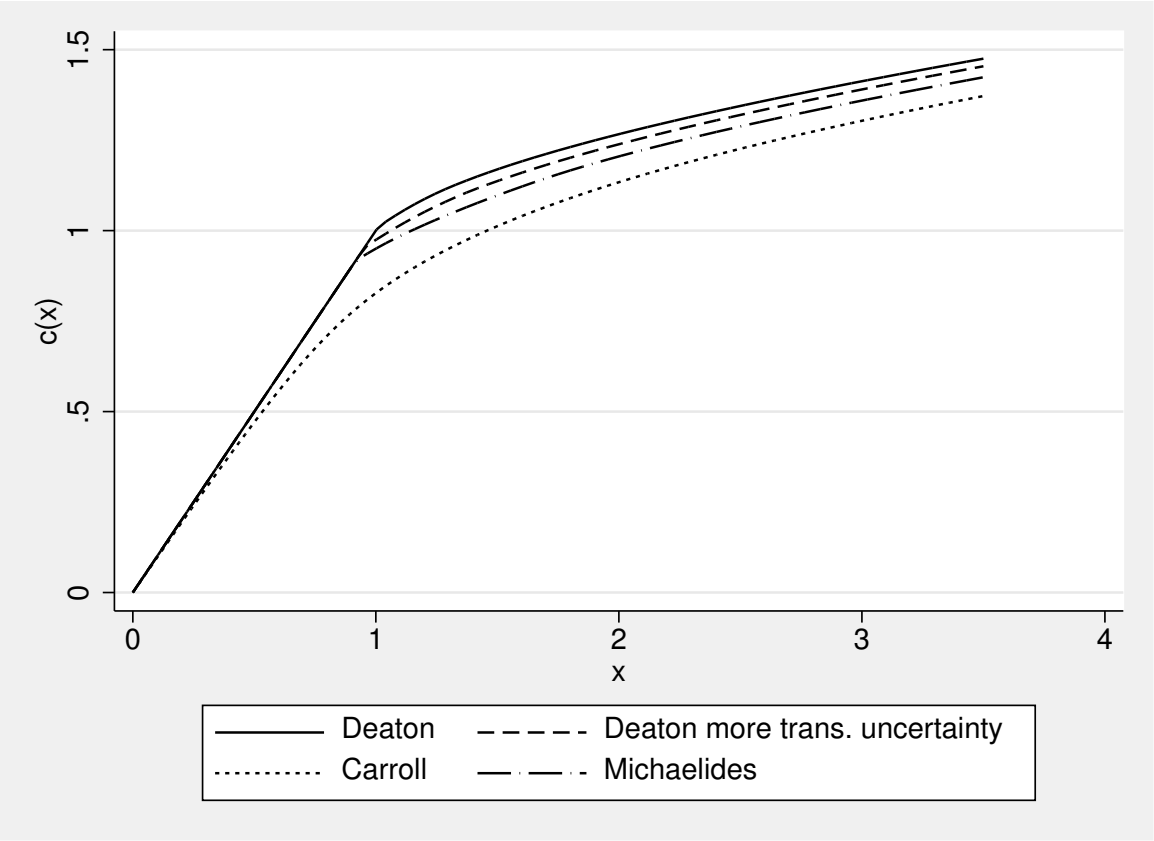

Figure 1: Consumption Functions

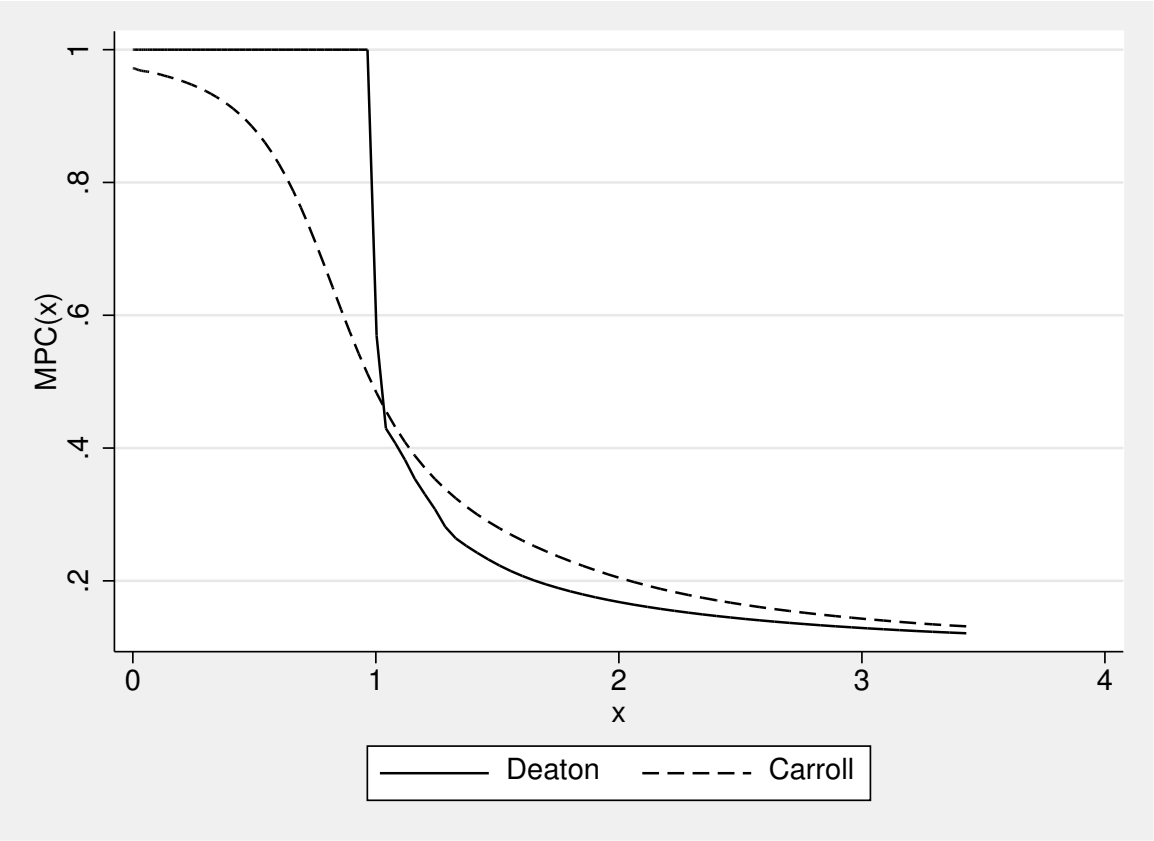

Figure 2: Marginal Propensities to Consume