

C EX
7/23 11:30

The Response of Consumption Growth to Expected Future Income Growth: Synthetic Panel Evidence*

Jeremy J. Nalewaik[†]
Graduate School of Business
University of Chicago

First Draft: January 5th, 2002

This Draft: March 19th, 2002

*The January 5th, 2002 version of this work circulated under the title "Consumer Prescience". I developed this work with the help of insights from Orazio Attanasio, John Cochrane, Steve Davis, Lars Hansen, Anil Kashyap, Lior Menzly, Paul Willen, and participants in the Chicago GSB Macro Lunch.

[†]E-mail: pjnalewa@gsbphd.uchicago.edu

Abstract

The main contribution of this paper is to show that consumers are forward-looking to a degree that some economists may find surprising. Using synthetic panel data from 1980-1999 on household consumption and income constructed from the Consumer Expenditure Survey (CEX) and the Current Population Survey (CPS), the paper shows that consumption growth forecasts income growth as much as six years in advance of its realization. Consumption growth forecasts both business cycle variation in income growth and longer term trends such as the expanding income gap between the least and most educated. The paper's use of micro data allows it to rule out competing stories for the forecasting power of consumption growth, such as Keynesian feedback, that have plagued previous macro studies in the same vein. Heterogeneity in the forecasting power of consumption growth over the life cycle and across education groups is documented; the forecasting power of consumption growth is concentrated at longer lags for households with more educated and more middle-aged male heads. In addition to the potential policy relevance of many of the results in this paper, its findings serve generally to inform our choice of models of household behavior.

1 Introduction

1.1 Some Motivation

Most existing empirical work characterizing households' intertemporal allocation of consumption expenditures focuses on estimating Euler equations (see Browning and Lusardi, 1996, for a survey). Many of these studies estimate Euler equations that hold with equality from models with quadratic utility over consumption, or log-linearized Euler equations that hold with equality from models with CRRA utility over consumption, where higher order moments such as the conditional variance of consumption growth are assumed constant. Many of these studies further assume that households' utility from consumption is separable from other determinants of household utility, especially leisure. For obvious reasons, we will refer to this model throughout the paper as the benchmark model, or the benchmark certainty equivalent model.

The focus on estimating Euler equations is easily explained: they provide information on an economically interesting parameter (the elasticity of intertemporal substitution), they require few auxiliary assumptions about some aspects of household behavior, and they provide simple and informative characterizations of the empirical failures of the benchmark model. These failures have come in the form of violations of the condition that the growth rate or change in the marginal utility of consumption should be orthogonal to lagged or predictable variables, especially lagged income growth. Characterization of these orthogonality violations has been particularly fruitful, with the empirical results spurring on the development of more complicated and realistic models incorporating precautionary savings motives, liquidity constraints, and non-separabilities between consumption and leisure or other determinants of utility. However the relative importance of each of these extensions of the benchmark model for explaining any given stylized fact (the correlation between lagged income and consumption, for example) remains controversial.

There are different ways to characterize the degree to which households intertemporally optimize and are forward-looking. The implication of forward-looking behavior studied here is the flip side to the orthogonality implication that income growth should not forecast consumption growth: here we study the implication that consumption growth should forecast income growth. Admittedly, the two implications are closely intertwined. Take the hypothesis that consumption growth should forecast income growth. Break it down into its component parts and define it more completely as the following joint hypothesis: (a) households receive advance information

about their future income, and (b) they alter their consumption in response to this information in the manner that the certainty equivalent benchmark model predicts: $\Delta C_{t+1} = \frac{r}{1+r} \sum_{j=0}^{\infty} (E_{t+1} - E_t) \frac{Y_{t+j+1}}{(1+r)^j}$, where C is consumption, Y is income, E_t is the expectation operator conditional on the household's information at time t , and, for simplicity, we've taken the case with a constant interest rate r and an infinite time horizon. If condition (a) is met but condition (b) is not met, then, as long as present value budget balance holds in expectation, the euler equation orthogonality conditions will be violated. To see this, assume that households under- or over-adjust their consumption response to a current change in expected future income. Then that under- or over-adjustment must be offset by an expected consumption change at some future date; otherwise the budget constraint will be violated in expectation. This future consumption change will be related to the current signal about income; hence the orthogonality violation.

The previous paragraph illustrates that, if the household's information set were fully observable by the econometrician, the degree to which households conform to the benchmark certainty equivalent model could be perfectly characterized by orthogonality violations. However, that is the rub: the econometrician does not observe the information set of the household, and the household's information regarding its own income is certainly superior to the information available to the econometrician. Orthogonality tests in the form of regressions of consumption on lagged income, while robust to the superior information of households, do not exploit it, and are thus bound in their characterization of consumer behavior by the information set of the econometrician. However if conditions (a) and (b) in the previous paragraph are true, then, as the equation in that paragraph suggests, consumption itself should reveal to the econometrician (part of) the household's superior information about its future income. This suggests that regressions of income on lagged consumption, rather than the other way around, will increase the amount of information employed by the econometric specification, allowing econometricians to construct more complete characterizations of consumer behavior. The specification employed in this paper, originally derived by Hansen, Roberds and Sargent (HRS, 1991) to test present value budget balance, exploits the consumer's superior information:

$$\Delta y_{t+1} = \beta_0 \Delta c_{t+1} + \beta_1 \Delta c_t + \dots + \beta_q \Delta c_{t+1-q} + e_{t+1}. \quad (1)$$

The specification is quite simple, a regression of income growth on contemporaneous consumption growth and lags of consumption growth. If consumers are able to forecast their future income growth and alter their consumption in response to these forecasts, the lags of consumption growth should help determine current income growth.

1.2 Review and Critique of Previous Empirical Work

A number of studies have employed macroeconomic data to estimate forecasting regressions of income by lagged consumption. Econometricians have employed two types of transformations to deal with the non-stationarity of the forecasting variable consumption: (i) scaling, and (ii) differencing. Papers employing the first transformation scale consumption by income, examining savings, savings rates, and log consumption-income ratios; most focus on forecasting income growth one or two quarters ahead. Campbell (1987) regresses the first difference of labor income on savings lagged one quarter, and finds a coefficient of -0.18 (0.05 standard error) using non-durable goods and services expenditures. Campbell and Deaton (1989) regress income growth on the savings rate lagged one quarter and find a coefficient of -0.18 (0.06 standard error). Cochrane (1994) regresses income growth on the log consumption-income ratio lagged one quarter and finds a coefficient of 0.08 (with a t-statistic of about 3.5). Campbell and Mankiw (1989) look at longer horizons, using the log consumption-income ratio to forecast the present discounted value of income growth as far out as their time series allows. They find coefficients ranging from 1.3 to 2.0, depending on the consumption variable used, with standard errors ranging between 0.2 and 0.4, but unfortunately, they do not break down this forecasting power by horizon.

A problem with the scaled consumption specifications is that the forecasting power for future income may be coming from the scaling variable, current income, rather than from current consumption. Take the results of Campbell and Mankiw (1989), for example. If income is slowly mean reverting, as the evidence in Cochrane (1988) and Cochrane (1994) suggests, then the detrended level of current income may forecast the discounted value of future income changes far into the future in a purely mechanical fashion, with the scaling of income by the economically interesting consumption variable serving primarily as window dressing. Another problem with the scaled consumption specifications is that current income appears in both the explanatory and dependent variables; if it is measured with error the regression coefficients are biased in the direction predicted by the theory. Indeed, the one quarter ahead forecasting results in the previous paragraph reflect no true forecasting power if about 10-20 percent of the variance of income is measurement error. A final problem is that the manner in which the scaling is done (often by estimation of a cointegrating vector) can be somewhat arbitrary and can influence the results. This is especially true when the consumption measure of choice is a subset of total consumption expenditures.

Although it avoids the problems discussed in the previous paragraph, fewer macroe-

conomic studies employ the strategy of first differencing consumption. Cochrane (1994) includes consumption growth lagged one quarter in his regression of income growth on the lagged scaled consumption variable; the coefficient on lagged consumption growth is 0.52 (with a t-statistic of about 3.8). And of course there is Hansen, Roberds and Sargent (HRS), but the focus of their work is more on testing present value budget balance than on examining the forecasting power of consumption growth.¹

Deaton (1992) has pointed out that all of these macroeconomic studies are open to reinterpretation:

In the aggregate economy, it is easy to think of other mechanisms - simple Keynesian feedback mechanisms being the obvious example - that generate correlation between saving and future income change, for example if positive consumption shocks are propagated into income increases in subsequent periods. (Deaton, 1992, p. 133.)

To rule out these competing explanations for the forecasting power of consumption, we must go to the micro data, to examine whether deviations of household's consumption growth from the aggregate forecast deviations of household's income growth from the aggregate.

Few empirical studies have examined the forecasting power of consumption using micro data, and those that have done so exploit only the life cycle variation in income and consumption to make inferences. A paper by Carroll (1994) is closest in spirit to this paper; it tests whether group differences in current consumption are related to group differences in expected future life cycle income profiles, and finds that they are not. Carroll's work is hampered by the fact that he uses only one cross section of consumption data. Further, Carroll regresses the level of consumption on the level of an expected future income variable constructed by him, a specification which does not exploit consumers' superior information.

1.3 Summary of Results and Contributions

The regression specification (1) has several insightful aspects to it not present in previously used empirical specifications. The derivation of (1) in HRS shows that, under the certainty equivalent benchmark model, the β s are the coefficients on a time series moving average representation of the observable dimension of the consumer's information set.

¹There are also some results using first differences in Campbell and Mankiw (1989) but they do not report their regression coefficients.

As such, they decompose the consumer's information about income growth by how far ahead of the arrival of income growth the information is revealed. This is in contrast to previous work using scaled consumption specifications, which cannot easily implement such decompositions, and by definition cannot decompose the consumer's information about income growth into the part which was known before the arrival of income growth and the part that is only revealed by income growth's actual arrival.

A further advantage of (1) is that the test of present value budget balance in HRS provides a benchmark value for a linear combination of the regression coefficients, telling us how reasonable the coefficients may be from the standpoint of the benchmark model. This is again in contrast to much previous work, which has often been content to examine simple sign tests (savings should negatively forecast future income, for example), or has implemented tests whose content is questionable. Specifically, inspection of the test advocated by Campbell (1987) reveals that it is little more than an Euler equation orthogonality test. Campbell and Mankiw (1989) provide an exception to the criticisms in this paragraph.

Section 2.2 reports a condensed version of some of the certainty equivalent benchmark results in HRS. Section 2.3 discusses the potential impact of some departures from the benchmark model on the β s in (1), taking insight from some of the rejections of the benchmark model found in the Euler equation literature. In particular, the paper discusses the impact of precautionary savings motives and non-separability between consumption and leisure on the regression coefficients; the discussion of precautionary motives relies heavily on the insights of Carroll (1992, 1997). The conclusion of the paper argues that precautionary motives might be important in understanding some of the empirical results in section 4.

Section 3 discusses the data, aggregation and econometric specification used to estimate equation (1). The paper uses twenty years of synthetic panel data to estimate the equation, grouping households by the educational attainment and birth cohort of the male head of the household, following Attanasio and Davis (1996). The consumption data are from the Consumer Expenditure Survey (CEX), while the income data are from the Current Population Survey (CPS). This section, and the paper in general, gives particular attention to various forms of measurement error in consumption growth.

The main results in section 4 of the paper provide remarkably strong support for forward-looking behavior on the part of consumers: the bulk of the regression specifications show consumption growth forecasting income growth five and six years in advance of its realization. Multiple types of variation in income growth are captured by the

forecasting power of consumption growth. Consumption growth forecasts the continued widening of income inequality between the least and most educated in the late 1980s and 1990s; much of the forecasting power at four, five and six year horizons comes from this source of variation in the data. However consumption growth also shows some ability to forecast what appears to be standard business cycle variation five and six years ahead of time, perhaps the most striking result in the paper. These results are robust to the inclusion of a wide range of control variables in (1), and hold true in specifications that explicitly correct for the main source of measurement error in our data - sampling variability. The results actually strengthen substantially when lags of income growth are included in the regression specification (1).

The results here help to rule out the alternative explanations raised by Deaton for the correlation between consumption growth and future income growth. It strains credibility that Keynesian models based on short-run wage or price rigidity can explain the long-horizon forecasting results in this paper. More importantly, the finding that deviations of groups' consumption growth from aggregate consumption growth forecast deviations of their income growth from aggregate income growth clearly cannot be explained by aggregate feedback.

Finally, the results in this paper demonstrate interesting heterogeneity in the forecasting power of consumption growth. The forecasting power of consumption growth is concentrated at longer lags for households headed by more educated males, and for households headed by males closer to middle age compared to households with younger heads. The paper provides discussion of some of the (multiple) economic interpretations of this heterogeneity.

Section 5 concludes with a discussion of the plausibility of the some of the results in the paper. Appendix A offers some details on the log-linearization of the household budget constraint in section 2.1; Appendix B describes computation of the robust asymptotic standard errors used in the empirical work.

2 Theory

As the introduction mentions, interpreting results from our regression specification (1) does not require lots of theory: if households receive advance information about their income growth and alter their consumption in response to this information, lags of consumption growth should predict income growth. Nevertheless, a theoretical derivation of the regression specification will prove useful, as it will (a) aid in our interpretation

of the regression coefficients, and (b) provide a testable implication on the magnitude of a linear combination of the regression coefficients. The benchmark certainty equivalent model in which present value budget balance holds - the case examined by HRS - provides a point of departure for study of the effects of features like precautionary savings motives and non-separabilities on the regression coefficients in (1). Non-interested readers can skip this section, peruse section 3.1 and then go directly to section 4.

2.1 Log linear consumption functions

Consider the problem of a infinitely-lived household h who maximizes utility over consumption subject to a standard law of motion for liquid assets. Call C_t^h the level of consumption at time t , A_t^h asset holdings, r the (assumed constant) interest rate, and X_t^h the level of time t exogenous “labor” income; although it may include transfer income, from now on we will simply refer to all exogeneous income as labor income. The law of motion for liquid assets is: $A_{t+1}^h = (1 + r) (A_t^h + X_t^h - C_t^h)$. The household utility function takes the following standard form:

$$U_t^h = E_t \sum_{j=0}^{\infty} (\beta^h)^j N_{t+j}^h \frac{(C_{t+j}^h)^{1-\gamma^h}}{1-\gamma^h},$$

where E_t is an expectation taken with respect to the time t information set of the consumer, and N_{t+j}^h represents factors that impact the household's utility from consumption, such as household size. Then optimization on the part of the household and lognormality of consumption growth yield the much researched and discussed log-linear euler equation:

$$E_t (\Delta c_{t+j}^h) = \frac{r - \delta^h}{\gamma^h} + \frac{E_t (\Delta n_{t+j}^h)}{\gamma^h} + \frac{\gamma^h}{2} \text{var}_t (\Delta c_{t+j}^h), \quad (2)$$

where lower case variables denote variables for which logs have been taken, $-\delta^h = \ln(\beta^h)$, and $\text{var}_t(\cdot)$ is the variance operator conditional on time t information.

Iterating forward the law of motion for liquid assets and imposing the no Ponzi condition yields the present value version of the budget constraint:

$$A_t^h + \sum_{j=0}^{\infty} \frac{X_{t+j}^h}{(1+r)^j} = W_t^h = \sum_{j=0}^{\infty} \frac{C_{t+j}^h}{(1+r)^j}, \quad (3)$$

where we've defined wealth W_t^h in the standard way. Appendix A discusses log-linearization techniques that convert the expected value of this budget constraint into the following

decomposition of consumption growth:

$$\Delta c_{t+1}^h = E_t \Delta c_{t+1}^h - \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) \Delta c_{t+j+1}^h + \sum_{j=1}^{\infty} \lambda^{j-1} (E_{t+1} - E_t) \Delta y_{t+j}^h, \quad (4)$$

where $\lambda = \frac{1}{1+r}$, and y_{t+j}^h is the log of total income, including income from capital. Realized consumption growth from t to $t+1$ is its time t expectation, minus the revision from t to $t+1$ in the consumer's expectations about future values of consumption growth, plus the revision from t to $t+1$ in the consumer's expectations about current and future values of income growth. Campbell (1993) does similar decompositions to innovations in consumption; his third term involves innovations to the return on the wealth portfolio, though, instead of income growth.

Substituting from (2) into (4) yields:

$$\begin{aligned} \Delta c_{t+1}^h = & \frac{r-\delta^h}{\gamma^h} + \frac{1}{\gamma^h} E_t (\Delta n_{t+1}^h) + \frac{\gamma^h}{2} \text{var}_t (\Delta c_{t+1}^h) \\ & + \frac{1}{\gamma^h} \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) (\Delta n_{t+j+1}^h) - \frac{\gamma^h}{2} \sum_{j=1}^{\infty} \lambda^j (\text{var}_{t+1} - \text{var}_t) (\Delta c_{t+j+1}^h) \\ & + (E_{t+1} - E_t) \Delta y_{t+1}^h + \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) (\Delta y_{t+j+1}^h). \end{aligned} \quad (5)$$

The first line of (5) constitutes the predictable variation in consumption growth, due to predictable variation in taste shifters and the conditional variance of consumption growth. The second line of (5) shows the variation in consumption growth due to innovations to the taste shifters and the conditional variance of consumption growth. The third line of (5) shows the variation in consumption growth due to innovations to current and expected future income growth.

In the next sub-section we will need some notation for the time series process for household income. Assume the log of the household's income is stationary in first differences, and write income growth in its $MA(q)$ representation:

$$\Delta y_{t+1}^h = \rho^h(L) \varepsilon_{t+1}, \quad (6)$$

where ε is the n -dimensional column vector of (white noise) innovations that generate the information in the economy, and $\rho^h(L)$ is an n -dimensional row vector of polynomials of order q in the lag operator. The fact that the vector of innovations to the information set of agents is of arbitrarily large dimension reflects the realistic assumption that agents have more information about their income growth than the econometrician, who we may assume observes only income growth and consumption growth.

2.2 Magnitudes for the regression coefficients: The HRS benchmark

The benchmark model discussed in this subsection assumes that the first two lines of (5) are constant, so that all of the variation in consumption growth is due to innovations to current and expected future income growth. This benchmark case is the one discussed in Hansen, Roberds and Sargent (1991), whose results I condense and repeat here for the convenience of the reader. Given (6), it is straightforward to show:

$$\Delta c_{t+1}^h = \rho^h(\lambda) \varepsilon_{t+1}, \quad (7)$$

so consumption growth equals the present discounted value of the vector of innovations to current and expected future income growth.

While the n -dimensional vector of innovations to the information set of the households is unobserved by the econometrician, consumption growth reveals one dimension of that information set. To prove this, HRS construct the matrix Q , whose first row is $\frac{\rho^h(\lambda)}{|\rho^h(\lambda)|}$, and whose other $n - 1$ rows are orthonormal vectors orthogonal to $\rho^h(\lambda)$. Then $\rho^h(\lambda) Q'$ is a row vector of zeros, except for its first element which is $|\rho^h(\lambda)|$. Define $\varepsilon_{t+1}^+ \equiv Q\varepsilon_{t+1}$, and break up ε_{t+1}^+ into its first element, $\varepsilon_{1,t+1}^+$, and the following $n - 1$ elements, $\varepsilon_{2,t+1}^+$, so $Q\varepsilon_{t+1} = [\varepsilon_{1,t+1}^+ \quad \varepsilon_{2,t+1}^+]$. Then since $Q'Q = I_n$ write:

$$\begin{aligned} \Delta c_{t+1}^h &= \rho^h(\lambda) \varepsilon_{t+1} \\ &= \rho^h(\lambda) Q'Q\varepsilon_{t+1} \\ &= |\rho^h(\lambda)| \varepsilon_{1,t+1}^+. \end{aligned} \quad (8)$$

$\varepsilon_{1,t+1}^+$ is the one dimension of the consumer's information set revealed to the econometrician by consumption growth. Since $\varepsilon_{1,t+1}^+ = \frac{\rho^h(\lambda)}{|\rho^h(\lambda)|} \varepsilon_{t+1}$, elements of ε_{t+1} that have larger discounted sums of polynomial coefficients, and hence comprise a larger part of the variance of Δy_{t+1}^h , receive more weight in $\varepsilon_{1,t+1}^+$. In this sense $\varepsilon_{1,t+1}^+$ is the maximal linear combination of elements of ε_{t+1} .

Define $\rho^h(L) \equiv \rho^h(L) Q'$, and again break up $\rho^h(L)$ into its first element, $\rho_1^h(L)$, and the following $n - 1$ elements, $\rho_2^h(L)$, so $\rho^h(L) Q' = [\rho_1^h(L) \quad \rho_2^h(L)]$. Then:

$$\begin{aligned} \Delta y_{t+1}^h &= \rho^h(L) \varepsilon_{t+1} \\ &= \rho^h(L) Q'Q\varepsilon_{t+1} \\ &= \rho_1^h(L) \varepsilon_{1,t+1}^+ + \rho_2^h(L) \varepsilon_{2,t+1}^+. \end{aligned} \quad (9)$$

The polynomial coefficients $\rho_1^h(L) = \frac{\rho^h(\lambda)}{|\rho^h(\lambda)|} \rho^h(L)'$ are the polynomial coefficients corresponding to $\varepsilon_{1,t+1}^+$, and are hence maximal in the same sense that $\varepsilon_{1,t+1}^+$ is maximal. Evidently $\rho_1^h(\lambda) = |\rho^h(\lambda)|$.

We can use the expression for consumption growth (8) to substitute $\varepsilon_{1,t+1}^+$ out of (9), yielding:

$$\begin{aligned} \Delta y_{t+1}^h &= \frac{\rho_1^h(L)}{|\rho^h(\lambda)|} \Delta c_{t+1}^h + \rho_2^h(L) \varepsilon_{2,t+1}^+ \\ &= \beta_0 \Delta c_{t+1}^h + \beta_1 \Delta c_t^h + \dots + \beta_q \Delta c_{t+1-q}^h + e_{t+1}^h, \end{aligned}$$

where in the second line we have rewritten the regression coefficients in the notation of equation (1), where $\beta(L) = \frac{\rho_1^h(L)}{|\rho^h(\lambda)|} = \frac{\rho^h(\lambda)}{|\rho^h(\lambda)|^2} \rho^h(L)'$. Then under this model the $\beta(L)$ in (1) are the moving average coefficients corresponding to a linear combination of elements of ε_{t+1} , with elements accounting for a larger fraction of the variance of income growth receiving more weight. As moving average coefficients, the β s essentially partition by horizon the variance of income growth falling within the observable dimension of the consumer's information set: a comparison of $(\beta_k)^2$ with $(\beta_{k-1})^2$ is a comparison of the incremental increase in the variance of income growth Δy_t^h captured by the (observable dimension of) the consumer's information set as we move from $t-k-1$ to $t-k$ with the incremental increase as we move from $t-k$ to $t-k+1$. For example, if $(\beta_k)^2$ is twice $(\beta_{k-1})^2$, twice as much of the consumer's (observable) uncertainty about income growth is resolved moving from $k+1$ to k periods ahead as is resolved moving from k to $k-1$ periods ahead of income growth's arrival.

Present value budget balance in this model provides us with a benchmark testable implication for the discounted sum of the regression coefficients. The restriction $\rho_1^h(\lambda) = |\rho^h(\lambda)|$ implies that $\beta(\lambda) = 1$, or:

$$\beta_0 + \lambda\beta_1 + \lambda^2\beta_2 + \dots + \lambda^q\beta_q = 1. \quad (10)$$

The discounted sum of the regression coefficients must sum to one. To understand essence of this test, assume that current income growth declines by one percent. To offset this decline in income and maintain present value budget balance, it is evident that the present discounted value of either past, current, or future consumption growth must decline by one percent. If consumption growth is unpredictable, as maintained in (7), then current income growth cannot have an impact on future consumption growth, and the consumer must either decrease current consumption growth or have reduced past consumption growth by one percent. This is the essence of the test (10).

2.3 Magnitudes for the regression coefficients: some departures from the benchmark

This section considers some selected departures from the benchmark model, largely motivated by previous findings in the consumption literature.

- **Saving for retirement:**

Consider a household who behaves as in the previous subsection, except that it scales down its consumption responses to income innovations by some fixed fraction $(1 - \chi)$, so:

$$\Delta c_{t+1}^h = (1 - \chi) \rho^h(\lambda) \varepsilon_{t+1}. \quad (11)$$

It is straightforward to repeat the results of the previous subsection, yielding $\beta(\lambda) = \frac{1}{(1-\chi)}$. The discounted sum of the regression coefficients should sum to more than one; since the consumption process is scaled down, a consumption growth rate of 1% signals a larger than 1% innovation to income growth, a $\frac{1}{(1-\chi)}\%$ innovation, in fact. All of the regression coefficients will be scaled up by the inverse of the fraction by which the consumption process is scaled down.

Household saving for retirement seems likely to produce consumption behavior of approximately this kind. On average a household will be saving some fraction of its income for retirement, and adjusting its retirement savings to income innovations, accumulating or decumulating assets depending on whether it is receiving positive or negative income shocks. These asset adjustments will effectively scale down the contemporaneous consumption response to income innovations, producing a model similar to (11).

For the sample of households studied in this paper, whose male heads are in their prime earning years, the effect of retirement savings on the regression coefficients could be significant.

- **Precautionary savings effects and liquidity constraints:**

The presence of the conditional variance terms in (5) provide the key departure from certainty equivalent behavior on the part of the consumer. Carroll (1997) has demonstrated a negative relationship between liquid asset holdings and expected consumption growth, essentially because larger stocks of liquid assets make it less costly for the agent to smooth consumption in the face of future income shocks, decreasing the conditional variance of consumption growth and, by the

euler equation (2), consumption growth itself. Consider the consumption-savings response Δc_{t+1-k}^h of a household to news $\varepsilon_{1,t+1-k}^+$ about expected future income change Δy_{t+1}^h , in the face of Carroll's precautionary savings effects. The household must carry forward the change in savings k periods ahead to smooth out the expected future income shock, changing the agent's expected consumption growth for those k intervening periods in the same direction as the initial consumption response Δc_{t+1-k}^h . For example, if $\varepsilon_{1,t+1-k}^{i,+}$ is negative, then Δc_{t+1-k}^h declines and savings increases, which decreases the variance of subsequent consumption growth rates and the consumption growth rates themselves for k periods. This effect on subsequent consumption growth rates must dampen the initial response Δc_{t+1-k}^h to $\varepsilon_{1,t+1-k}^{i,+}$ compared to the certainty equivalent case; if Δc_{t+1-k}^h changed by the discounted value of $\varepsilon_{1,t+1-k}^{i,+}$ with subsequent consumption growth rates responding in the same direction, we would have an overall over-response of consumption to the innovation and a violation of present value budget balance.

As in the retirement savings example, the dampened consumption response to the expected future income innovation implies that small consumption changes may signal larger, more than proportionate future income changes, increasing the β_k in (10) so that the discounted sum of these coefficients may sum to more than one for reasonable discount rates. It seems the effect of precautionary savings on β_k should increase in k : a larger k means more intervening consumption growth rates change in the same direction as the initial response of consumption growth, on average implying a smaller initial consumption response if present value budget balance holds.

The precautionary effects shift the response of consumption growth to news about future income forward in time, closer to the actual arrival time of the future income and further away from the arrival time of the news. Tight liquidity constraints will likely have similar effects; it is well known that models with liquidity constraints have features very similar to models with substantial precautionary savings motives (see Deaton (1992) and Carroll and XXXX (200X)). The fact that consumption growth on average responds to lagged income growth in these models implies that consumption growth is not responding enough to news about current and future income growth (from the standpoint of the certainty equivalent model).

- **Non-separable leisure:**

If the taste shifters Δn_{t+j+1}^h in (5) include work hours, their innovations will cer-

tainly be positively correlated with innovations to the income process - an upward sloping labor supply function dictates that some positive fraction of all expected future earned income changes will consist of expected future work hours changes. Assuming that work hours raise the marginal utility of consumption (as is posited in Ghez and Becker (1975), for example), then the response of consumption growth to an increase in expected future income will be less than it would if preferences over consumption and leisure were separable. The increase in expected future income will consist partially of higher work hours, raising future consumption requirements and increasing current savings so the agent may maintain constant the marginal utility of consumption in those future periods of increased work; the increase in savings will reduce the immediate increase in consumption growth in response to the expected future income increase. By the same logic as the previous two examples, this should serve to increase β_k , compared to the benchmark case of separability.

3 Empirical Implementation

3.1 Data Description

Twenty years of Consumer Expenditure Surveys - from 1980 to 1999 - provide the information on household consumption used in this paper. The CEX is a rotating panel; new households have been continuously entering the survey since 1980, and are generally interviewed quarterly for four consecutive quarters. About 5,000 to 6,000 households are interviewed in any given year. The CEX is the most comprehensive source of data on household expenditure patterns available in the US. The household consumption measures used here is that of Attanasio and Davis (1996); it is non-durable goods and services, excluding expenditures on durables, medical care, education, and housing. Each household's consumption was deflated by the annual PCE deflator from the National Income and Product Accounts. The CEX food at home data was corrected for discontinuities in 1982 and 1987 introduced by changes in survey design; see the data appendix for details.

The Current Population Survey Annual Demographic Files provide the paper's primary source for income data over the twenty year period for which the CEX consumption data is available. The CPS interviews 50,000 to 60,000 households each March, asking about their income in the previous calendar year. The CPS, like the CEX, has a rotating

panel aspect to it. Household income from the CPS is computed as the sum of earned income, transfer income, and asset income, minus taxes. Some top-coding adjustments are made to the earned income variables. Although information on taxes paid is not available from the CPS, enough information is available to estimate taxes paid for each household using the NBER's TAXSIM program. Each household's income was deflated by the annual PCE deflator from the National Income and Product Accounts. Once again, the data appendix provides details.

The CEX queries households about their income as well as their consumption; some regression specifications below employ this income information. The construction of household income from the CEX follows the CPS construction as closely as possible; in particular, I used TAXSIM to estimate taxes paid by CEX households, rather than using the tax information provided by the CEX.

The paper uses grouping techniques to exploit the full twenty year time series dimension of the CEX and CPS data. While grouping averages out much of the household-level idiosyncratic variation in the data, we can form the groups in such a way that we capture some important distributional variation in the data, as well as the aggregate variation which will naturally remain after grouping. It is well known that the US economy has experienced large relative income variation across groups differentiated by experience and by educational attainment, and that consumption is imperfectly insured against this relative income variation (see Cutler and Katz (1991), Katz and Murphy (1992), and Attanasio and Davis (1996)). Forming groups on the basis of the birth cohort and educational attainment of household heads will naturally isolate this relative income and consumption variation.

The sample of households from which we form groups is again that of Attanasio and Davis (1996), households with a male head aged 23 to 59. These restrictions exclude about a half of all households from the sample. In addition, the CEX sample excludes rural households and households classified as incomplete income reporters, reducing the CEX sample to about a third of total households. The sample is then partitioned by the 5-year birth cohort and the educational attainment of the male household head. The four categories for educational attainment are: **1.** high school drop-outs (< 12 years of schooling), **2.** high school graduates with no other schooling (12 years), **3.** some post-secondary education less than a 4-year college degree (13-15 years), and **4.** 4-year college graduates and higher (≥ 16 years). For each educational group, we have a set of five-year birth cohorts. Index the five year birth cohort, birth_5 , by the year of birth of its youngest member, so $\text{birth}_5 = 1925$ members were born from 1921 to 1925. After

imposing the sample selection restrictions on age, 9 birth cohort groups remain in the sample: birth5 = 1930, 1935, . . . , 1970.

Synthetic panel on consumption is computed as the annual geometric means of the CEX consumption variable for each group. The household consumption data is monthly, usually spanning two years. When aggregating, data from one household usually ends up in two adjacent years, except when noted below. Synthetic panel on income is the annual geometric means of the CPS income variable. There are potentially 20 annual time series observations for each group; for the 36 total groups, we have 524 annual observations before differencing. The panel is unbalanced; the youngest and oldest cohorts have short panel dimensions. After taking differences and computing several lags of consumption growth, some cohorts drop from the sample. For example, the panel dimension of birth5=1930 and birth5=1970 is too short to compute six annual lags of consumption growth, and these cohorts drop from the sample in specifications employing this number of lags.

Table 1 reports some cell count summary statistics for number of households. In computing the cell counts, each CEX household is assigned to a unique year to avoid double counting. As can be readily seen, the much larger sample of households in the CPS is an attractive feature of that data source.

**Table 1: Cell Count Summary Statistics
Number of Households per Cell**

Variable	Data Source	Min	Means by Education				Max
			< 12	12	13 - 15	> 16	
Income	CPS	289	638	1401	995	1183	2120
Consumption	CEX Interview	26	56	121	106	130	269

3.2 Aggregation

Since the number and identity of households in the CEX and CPS samples varies from year to year, the number and identity of households in any group i will vary from year to year, and we write $i(t + 1)$. Define the sample average of any variable x as $\overline{x^{i(t+1)}} = \sum_{h=1}^{i(t+1)} x^h / i(t + 1)$. Since we form groups on the basis of largely time invariant characteristics, and since the effect of mortality and family dissolution/recombination will be minimized by the sample selection restrictions on age, treat the set of households in the population belonging to group i as fixed, and define the average of x over all

members of group i in the population as \bar{x}^i .

Define the difference between the group i , time $t + 1$ measured sample average of household income and the corresponding true population average as $\eta_{t+1}^{y,i(t+1)}$; this error term contains both residual response error and true deviations of the sample from the population. Then measured group i income growth can be written as:

$$\overline{y_{t+1}^{i(t+1)}} - \overline{y_t^{i(t)}} = \overline{\rho^i}(\text{L}) \varepsilon_{t+1} + \eta_{t+1}^{y,i(t+1)} - \eta_t^{y,i(t)}. \quad (12)$$

The variation in group i membership from t to $t + 1$ introduces into our time series of growth rates the cross sectional variation in mean income between the group i sample at time t and the group i sample at time $t + 1$. With a non-stationary income process or large within-group initial differences in earnings capacity, this cross sectional variation can be a substantial source of measurement error. Of course, the panel dimension to the CPS cuts down on this source of measurement error.

Next consider household consumption. It is convenient to introduce some new notation; write the log of household consumption as: $c_{t+1}^h = \mu_{t+1}^{c,h} + \rho^h(\lambda)(\varepsilon_{t+1} + \varepsilon_t + \dots) + u_{t+1}^h$. Here the household specific mean of log consumption $\mu_{t+1}^{c,h}$ reflects initial earnings capacity, as well as the cumulated impact of household-specific taste shifters, impatience levels, and exposure to income and other types of risk - see (2) and the first two lines of (5). The term u_{t+1}^h reflects the time $t + 1$ innovations to the taste shifters and conditional variance of consumption growth - see the second line of (5). Then defining the sampling deviations of the group means of log consumption from their true population values as $\eta_{t+1}^{c,i(t+1)}$, we have:

$$\overline{c_{t+1}^{i(t+1)}} - \overline{c_t^{i(t)}} = \overline{\Delta \mu_{t+1}^{c,i}} + \overline{\rho^i}(\lambda) \varepsilon_{t+1} + \overline{u_{t+1}^{i(t+1)}} - \overline{u_t^{i(t)}} + \eta_{t+1}^{c,i(t+1)} - \eta_t^{c,i(t)}. \quad (13)$$

As with the CPS, the panel dimension to the CEX reduces the variance of the sampling deviations in this equation.

From here we can proceed with theoretical results of section 2.2, substituting the polynomial coefficients of the group i population means $\overline{\rho^i}$ for those of an individual household. Group i consumption growth computed from the sample averages does not cleanly identify a linear combination of innovations to group i income growth; if we define $\beta(\text{L}) = \frac{\overline{\rho_1^i}(\text{L})}{|\overline{\rho^i}(\lambda)|}$ we have:

$$\overline{y_{t+1}^{i(t+1)}} - \overline{y_t^{i(t)}} = \beta(\text{L}) \left(\overline{c_{t+1}^{i(t+1)}} - \overline{c_t^{i(t)}} \right) + e_{t+1}^i, \quad (14)$$

where the composite error term contains various forms of measurement error in consumption growth and income growth, and $n - 1$ dimensions of time series variation in

household's information about the income process:

$$e_{t+1}^i = -\beta(L) \left(\overline{u_{t+1}^{i(t+1)}} - \overline{u_t^{i(t)}} + \overline{\Delta\mu_{t+1}^{c,i}} + \overline{\eta_{t+1}^{c,i(t+1)}} - \overline{\eta_t^{c,i(t)}} \right) + \overline{\eta_{t+1}^{y,i(t+1)}} - \overline{\eta_t^{y,i(t)}} + \overline{\rho_2^i(L)} \varepsilon_{2,t+1}^+ \quad (15)$$

The benchmark values for the coefficients remain $\beta(\lambda) = 1$, although three sources of errors in consumption growth as a measure of income innovations may bias the regression coefficients: (i) contemporaneous innovations to taste shifters and conditional variances, (ii) predictable variation in consumption growth, and (iii) sampling variability.

3.3 Econometric Considerations

This section considers the econometric specification of regression equation (14), paying particular attention to the sources of measurement error in consumption growth identified in the previous subsection. Section 2.3 provides some analysis of the effect of the first source of measurement error, the contemporaneous innovations to taste shifters and conditional variances. Not much can or perhaps should be done to minimize the effect of these terms, as they provide interesting alternative hypotheses to the benchmark certainty equivalent model. We would like to minimize the effect of the second source of measurement error, the predictable variation in consumption growth, as this simply adds noise to the regression. Most of the regression specifications purge consumption growth of some of its predictable variation with control variables; our maintained assumption is that ratio of news-induced variation to predictable variation in consumption growth is lower across-cohorts and over the life cycle than it is over the business cycle.

The third source of measurement error in consumption growth is sampling variability. Table 2 presents some summary statistics for CPS income growth and CEX consumption growth. Standard deviations are computed weighting synthetic cohort observations by the average cell count of the two cells used to compute each growth rate.² The autocorrelations are computed using OLS, weighting the data by the average cell counts of the cells producing the explanatory variable, and including as controls in the regressions cohort dummies and quartic age polynomials interacted with dummies for the four education groups. The large negative first order autocorrelation for consumption growth is consistent with substantial sampling variability, which follows an $MA(1)$ with parameter minus one if the samples are drawn independently from year to year.

²The standard deviations of the raw data are 0.040 for income growth and 0.060 for consumption growth.

Table 2: Summary Statistics

Variable	Data Source	Standard Deviation	Autocorrelations					
			1	2	3	4	5	6
Δy	CPS	0.035	-0.11 (0.09)	-0.02 (0.07)	-0.05 (0.08)	-0.09 (0.06)	-0.14 (0.08)	-0.21 (0.08)
Δc	CEX Interview	0.054	-0.34 (0.04)	-0.11 (0.06)	-0.05 (0.05)	0.06 (0.05)	-0.03 (0.05)	-0.09 (0.06)

The empirical section coming up reports results from two regression estimators, a weighted least squares (WLS) estimator and weighted errors-in-variables (WEIV) estimator. While the WLS estimator does not explicitly correct for biases introduced by sampling variability, it does attempt to minimize them. Table 1 illustrates substantial heterogeneity in cell sizes in our synthetic panel data; the weights in the WLS estimator are the average cell sizes of the explanatory consumption variables in (14), and effectively downweight the thinner synthetic panel observations more contaminated with sampling variability.

Inspection of (15) reveals that the error terms of the regression equation are likely to exhibit contemporaneous cross correlation (across synthetic cohorts) as well as autocorrelation. The asymptotic variance-covariance matrix of the WLS estimates is robust these phenomena, as well as heteroskedasticity; appendix B provides details of its computation. Two strategies for testing $\beta(\lambda) = 1$ are employed with the WLS estimator. The first fixes a value for r and simply computes the discounted sum of the regression coefficients and its standard error. The second estimates r from an estimate of λ - the positive real root of the polynomial in the regression coefficients defined in (10). The standard error of the estimate of r is computed using the delta method and numerical derivatives.

The WEIV estimator follows Deaton's (1985) suggestion to use the estimated sampling variances of the synthetic cohort cell means on consumption to estimate the variance-covariance matrix of the sampling errors $\eta_{t+1}^{c,i(t+1)} - \eta_t^{c,i(t)}$. Call this variance-covariance matrix $\varepsilon'_x \varepsilon_x$; a typical term on its diagonal is an estimate of the sampling error variance of $\left(\overline{c_{t+1-k}^{i(t+1-k)}} - \overline{c_{t-k}^{i(t-k)}} \right)$, and is computed as the average of the estimated sample variances of $\overline{c_{t+1-k}^{i(t+1-k)}}$ plus the average of the estimated sample variances of $\overline{c_{t-k}^{i(t-k)}}$. The terms on the first off diagonals are estimates of the covariance of the sampling errors of $\left(\overline{c_{t+1-k}^{i(t+1-k)}} - \overline{c_{t-k}^{i(t-k)}} \right)$ with $\left(\overline{c_{t-k}^{i(t-k)}} - \overline{c_{t-1-k}^{i(t-1-k)}} \right)$, and are computed as minus the average of the estimated sample variances of $\overline{c_{t-k}^{i(t-k)}}$. With $\varepsilon'_x \varepsilon_x$ in hand, a measurement-error

corrected regression estimate is $\beta = [X'X - \varepsilon'_x \varepsilon_x]^{-1} X'Y$, where X is the matrix of regressors (current and lagged CEX consumption growth) measured with error, and Y is the vector of CPS income growth.

The panel dimension to the CEX complicates the estimation of $\varepsilon'_x \varepsilon_x$ using the estimated sample variances. Deaton's straightforward suggestions assume that the group i sample at time $t + 1$ is drawn independently from the group i sample at time t , and will overestimate the variance of $\eta_{t+1}^{c,i(t+1)} - \eta_t^{c,i(t)}$ if computed from a rotating panel. To deal with this issue when implementing the WEIV estimator, all of the consumption observations for each CEX household are assigned to a unique year. In addition, we compute bootstrap samples from the CEX and CPS data, treating each year of the CEX and each year of the CPS data as a separate sample, and report the β estimates as the mean values of the 200 estimates computed from bootstrapped synthetic panels. The randomization introduced by bootstrapping should help to break up the panel dimension to the CEX and CPS data, making our estimates of $\varepsilon'_x \varepsilon_x$ and $\varepsilon'_y \varepsilon_y$ (the variance of sampling errors in CPS income growth) closer to consistent. The standard errors of the β estimates are naturally the standard deviation of the 200 bootstrapped estimates.

Given that $\varepsilon'_x \varepsilon_x$ is estimated, the estimator described above can occasionally produce some nonsensical estimates, for example when a strange bootstrap sample produces a negative definite $[X'X - \varepsilon'_x \varepsilon_x]$. Deaton's suggestions are derived from the work of Fuller (1987), who discusses the small sample properties of the EIV estimator in depth and proposes some corrections to improve the within sample performance of the estimator. Let k be the number of explanatory variables and n is the number of observations in the regression. Define $Z = [Y \quad X]$ and $\varepsilon'_z \varepsilon_z$ as a $k + 1$ by $k + 1$ matrix whose first row and column are zeros and whose lower right k by k block is $\varepsilon'_x \varepsilon_x$. Define $\hat{\rho}_k$ as the smallest root of $|Z'Z - \rho \varepsilon'_z \varepsilon_z| = 0$. If $\hat{\rho}_k > 1$, then let $\gamma = [1 - \frac{k}{n}]$, otherwise $\gamma = [\hat{\rho}_k - \frac{k}{n}]$. Fuller's estimator is then:

$$\beta = [X'X - \gamma(\varepsilon'_x \varepsilon_x)]^{-1} X'Y. \quad (16)$$

This estimator is the first stage of a two stage procedure, again recommended by Fuller, that produces the WEIV estimator. Using these first stage β s and (15), define the average non-sampling error variance of the residuals of the first stage regression as: $\hat{\sigma}_e = \frac{1}{n} (Y - X\beta)'(Y - X\beta) - \frac{1}{n} \beta' (\varepsilon'_x \varepsilon_x) \beta - \frac{1}{n} \varepsilon'_y \varepsilon_y$. We bound the estimated $\hat{\sigma}_e$ below at zero and above at the variance of CPS income growth measured with error. Call $(\varepsilon_x^{it})' (\varepsilon_x^{it})$ the estimated variance of the sampling error for observation it - i.e. for group i 's current and lagged CEX consumption growth at time t , and use similar notation to

define the estimated sampling error variance of CPS income growth for that particular observation. Then the estimated variance of the regression residual for observation it is: $\hat{\sigma}_e^{it} = \hat{\sigma}_e + \beta' (\varepsilon_x^{it})' (\varepsilon_x^{it}) \beta + (\varepsilon_y^{it})' (\varepsilon_y^{it})$. The inverse of these residual variances are used as weights in the second stage of Fuller estimation, producing β^{WEIV} .

4 Results

4.1 Table 3:

Panel A of table 3 shows WLS estimates of (1) with various cutoffs q , with no control variables except for a constant. It shows the patterns that will appear over and over again in this data - consumption growth forecasts income growth at long lags, as far as five years ahead, and the bulk of the forecasting power comes at lags two to five. Panel B of table 3 reverses the role of Δc_t and Δy_t , regressing consumption on lags of income. While panel B shows a statistically significant relation between Δc_t and Δy_{t-1} , a relation that has showed up as a violation of the orthogonality conditions in many Euler equation estimates, there is very little evidence of a relation between consumption growth and further lags of income growth. We get similar results if we exclude contemporaneous income growth Δy_t from these regressions, or if we regress income growth on leads of consumption growth - there is a significant relation between income growth and consumption growth one year later, but no significant relation at longer horizons. We never see anything like the patterns we observe when regressing income growth on lags of consumption growth. A comparison of the R^2 patterns between the two panels in table 3 is informative as well, although the R^2 levels themselves are not so informative; it should be kept in mind that the sampling variability in income and consumption growth is reducing them considerably.

As an aside, note that the correlation between Δc_t and Δy_t fades as we add more lags to the regressions. This is due to the tightening of the sample restrictions as we add more lags: the sample of contemporaneous growth rates in the zero lag specification runs from 1980-1 to 1998-9, and runs from 1986-7 to 1998-9 in the six lag specification (so that we may compute the sixth lags from 1980-1 to 1992-3). If the zero lag panel A specification is estimated on the same 268 observations used to estimate the six lag specification, the coefficient is 0.09, with a standard error of 0.04; a similar drop-off is obtained from the panel B specification. Evidently the contemporaneous correlation between income and consumption growth has fallen over time, perhaps due to expanded

opportunities for consumption smoothing or consumption insurance.

4.2 Table 4:

Table 4 presents many of the main results of the paper, showing the WLS regression results from six and seven lag specifications with a broad range of control variables. To get an idea of whether the regression coefficients are reasonable from the standpoint of the basic certainty equivalent model with separable preferences, it reports the discounted sum of the regression coefficients under the assumption of a 2.5% real interest rate (see equation (10)), as well as the estimated value of the interest rate for which (10) holds. The last column reports $R_{\Delta c}^2$, the regression R^2 after orthogonalizing the data with respect to the control variables.

The first specification of table 4 repeats the last line of table 3, panel A. Compare this to the second specification, which includes as controls quartic age polynomials and birth cohort fixed effects, both interacted with fixed effects for the four education groups. The purpose of the controls is to purge consumption growth of some predictable variation, arising from things like expected changes in family size and income risk over the life cycle, and mean differences in income risk across cohorts and education groups. The controls should reduce a form of measurement error in the explanatory consumption variables, and apparently they do, as the regression coefficients on the lags of consumption growth increase substantially.

The next four specifications in table 4 include year fixed effects as control variables, purging the data of its aggregate effects and isolating its cross sectional dimension. We see that relative movements in consumption growth forecast relative movements in income growth, a significant and perhaps decisive defeat for Keynesian aggregate feedback interpretations of the forecasting power of consumption. The horizon at which consumption growth forecasts income growth is actually lengthened when we include the year effects. The forecasting power of consumption growth for income growth six years ahead is now statistically reliable, and the bulk of the forecasting power now occurs at four, five, and six year horizons. This shows that a large amount of the forecasting power of consumption growth at longer horizons comes from relative movements across the education and birth cohort groups studied here.

Figure 1 plots the relative variation in income growth and expected income growth across our four education groups. The expected income growth measure is the predicted value from the fourth regression specification in table 4 - the seven lag specification with data orthogonalized with respect to year dummies. Actual and expected income

growth are averaged over the birth cohorts within each education category to produce the data plotted in figure 1. While the predicted values do not pick up much of the year-to-year fluctuations in relative income growth, they do reflect the widening income gap between households headed by the least and most educated, and (perhaps) the deceleration of this widening. The regression coefficients from this specification show that the predicted values reflect most strongly relative consumption growth four, five and six years before the arrival of relative income growth (the coefficients are 0.15, 0.20, and 0.15), and reflect hardly at all relative consumption growth contemporaneous to relative income growth (the coefficient is 0.02). The data indicate that households were able to forecast the increasing returns to education that took place in the late 1980s and 1990s, and adjusted their consumption in anticipation of these expected relative income movements.

The remainder of the specifications in table 4 include either of two sets of income growth variables as controls. One set, labelled Δy_{t-j}^{CEX} , is contemporaneous income growth and q lags of CEX income growth: $\Delta y_t^{CEX}, \Delta y_{t-1}^{CEX}, \dots, \Delta y_{t-q}^{CEX}$. This set of controls is interesting because it will likely serve to purge the explanatory consumption variables of much of their sampling variability: the $\eta_{t-k}^{c,i(t-k)} - \eta_{t-k-1}^{c,i(t-k-1)}$ are surely highly correlated with the $\eta_{t-k}^{y,i(t-k)} - \eta_{t-k-1}^{y,i(t-k-1)}$ from the CEX sample. Being computed from the same samples, both will be driven by within-group cross-sectional differences in earnings capacity introduced by changes in sample composition over time. The second set of controls, labelled Δy_{t-j}^{CPS} , is q lags of CPS income growth: $\Delta y_{t-1}^{CPS}, \Delta y_{t-2}^{CPS}, \dots, \Delta y_{t-q}^{CPS}$. Due to the larger CPS sample sizes, these are likely less contaminated with measurement error than their CEX counterparts; these specifications provide a clear answer to the question of whether consumption growth has additional explanatory power for income growth above and beyond lags of income growth.

Controlling for income growth increases the magnitude of the coefficients on the longer lags of consumption growth - lags four, five, and six, especially compared to the coefficients on contemporaneous consumption growth. Interpreted through the lense of the benchmark model in section 2.2, the specifications with Δy_{t-j}^{CEX} controls are generally consistent with the hypothesis that consumers learn nothing new about their income growth one year in advance and when it actually arrives - i.e. all information about income growth has been revealed two years in advance of its realization, and households have altered their consumption growth in response to this information. The $R_{\Delta c}^2$ generally increase with the income variables as controls, as we would expect if the controls reduce the amount of measurement error in the data. For most of these specifications

we cannot reject a reasonable value for r ; the largest departures from a reasonable value for r are with the Δy_{t-j}^{CPS} controls, when the estimated interest rate appears too high to be consistent with the benchmark certainty equivalent model. This can be interpreted as evidence for some of the departures from the benchmark model discussed in section 2.3.

Figure 2 plots the education-specific variation in income growth and expected income growth, computed as the predicted values from the seven lag specification with data orthogonalized with respect to Δy_{t-j}^{CPS} (the tenth regression specification in table 4). The figure shows that the expected income growth measure captures a non-trivial amount of business cycle-type variation in income growth for all education groups except high school dropouts. The squared correlations between the expected and actual income growth are, in increasing order of the household head's education: 0.16, 0.63, 0.70, and 0.49. The expected income growth measure is largely a function of consumption growth two to six years before income growth arrives - the data support the view that households are able to anticipate and respond to future business cycle shocks.

In table 4, the regression coefficients on the income growth controls are not reported, but are available from the author on request. Most were negative and many statistically significant. If the income process is truly multi-dimensional, then the $\bar{\rho}_2^i(L) \varepsilon_{2,t+1}^+$ term in the error (15) will be non-zero and generally forecastable, so the theory does not say that lags of income or other control variables should have no marginal forecasting power beyond that in consumption growth. What the theory does say is that the control variables should not strip consumption growth of its forecasting power, and they do not. One last interesting fact to note: compared to regressions of income growth on lags of income growth alone, regressions of income growth on lags of income growth and lags of consumption growth have larger negative coefficients on the lags of income growth. This is consistent with the theory: consumption growth should respond more strongly to permanent income changes than temporary ones, so after conditioning on consumption growth, the remaining variation in income growth is likely to be more transitory. The greater the transitory variation in the income growth process, the more negative autocorrelation we should and do see. Put in the language of section 2.2, $\bar{\rho}_2^i(L) \varepsilon_{2,t+1}^+$ should be less persistent than $\bar{\rho}^i(L) \varepsilon_{t+1}$, as much of the persistence of $\bar{\rho}^i(L) \varepsilon_{t+1}$ has been drawn out of the data by $\bar{\rho}_1^i(L) \varepsilon_{1,t+1}^+$. For evidence on this point using macro data and a vector autoregressive specification, see Cochrane (1994).

4.3 Table 5:

Table 5 shows WEIV estimates of the same set the regression specifications as in table 4. The amount of noise added to the CPS and CEX data by breaking up their panel dimensions is considerable: the standard deviation of CPS income growth from the typical bootstrapped sample is about 25% higher than in table 2, while the standard deviation of CEX consumption growth roughly doubles in these bootstrapped samples. Occasionally a bootstrap sample produced regression estimates for which there was no sensible estimate of λ (i.e. no positive real root of the polynomial in the regression coefficients defined in (10)). For this reason no estimates of r are reported, only the discounted sums of regression coefficients with $r = 2.5\%$. R^2 s are not reported either since they are not well defined for these estimates.

Specification one shows the mean WLS estimates from the 200 bootstrapped synthetic panels, with a constant as the only control. It shows the effects of breaking up the panel dimension to the data, and provides a comparison with line two, which shows the corresponding WEIV estimate. The corrections for sampling variability generally, although not always, increase the regression coefficients.

The specifications with control variables Δy_{t-j}^{CEX} and Δy_{t-j}^{CPS} require some straightforward extensions to the computations described in section 3.3. With the Δy_{t-j}^{CEX} controls, the estimates of $\varepsilon'_x \varepsilon_x$ must take into account the sampling covariances between the consumption growth and income growth terms. With the Δy_{t-j}^{CPS} controls, the estimates must take into account the sampling covariance between the dependent variable Δy_t^{CPS} and the explanatory variable Δy_{t-1}^{CPS} . Defining the vector of estimated sampling covariances between X and Y as $(\varepsilon'_x \varepsilon_y)$, the estimates with the Δy_{t-j}^{CPS} controls are computed as $\beta = [X'X - \gamma(\varepsilon'_x \varepsilon_x)]^{-1} [X'Y - \gamma(\varepsilon'_x \varepsilon_y)]$. See Fuller (1987).

Compared to the results in table 4, some of the regression coefficients decrease, and some increase substantially - for reasonable discount rates the discounted sum of the coefficients is greater than two for a couple of specifications. These large values for the coefficients again raise the possibility that the stories discussed in section 2.3 - saving for retirement, precautionary savings and liquidity constraints, and non-separable leisure - are important features of the data. For the specifications in which the discounted sum of the coefficients are smaller than their counterparts in table 4, the increase in the size of the coefficients from implementing the EIV corrections is not large enough to offset the decline in the regression coefficients brought about by the loss of the panel dimension to the data. Generally we can say that the WEIV estimates are more volatile than their WLS counterparts.

The estimates in table 5 have the advantage that they make explicit corrections for sampling variability, but these corrections are far from perfect. We must throw an important amount of signal out of the data to correct for the noise, and it is not clear that the payoff to implementing the correction outweighs the cost. The most important lesson from table 5 is that corrections for sampling variability are unlikely to appreciably alter the patterns in the regression coefficients observed in table 4.

4.4 Heterogeneity - Tables 6, 7:

A number of papers have argued that a household's propensity to be liquidity constrained varies in a systematic way over the life cycle; see Gourinchas and Parker (1999) for a recent example. To explore this possibility, the paper allows the coefficients on consumption growth in (1) to differ by the age of the household head.

Table 6 reports six sets of regression coefficients from three regression specifications with varying age cutoffs. A five year birth cohort is assigned the age of its middle-aged members, and for a given age cutoff (either 33, 38, or 43), two sets of consumption growth terms are included in the regression: one interacted with dummies for whether the cohort was younger than the cutoff (at the time of the consumption growth), and another interacted with dummies for whether the cohort was older (so the first and fourth set of reported estimates are from the same regression). Controls are quartic age polynomials and birth cohort fixed effects, interacted with education fixed effects.

The differences in the patterns and magnitudes of the regression coefficients across age groups are informative. For the older cohorts, the distribution of the regression coefficients is more heavily weighted towards the longer lags of consumption growth; this is especially true for cohorts aged 43 to 59, where the coefficients on the fourth and fifth lags are quite large. For the 23 to 33 age group, and 23 to 38 group, the largest coefficient is on contemporaneous consumption growth, and generally the distribution of the coefficients falls more heavily on the shorter lags. These differences have multiple interpretations: they could be due to differences in the arrival time of information about future income, or differences in the ability of different age groups to change their consumption in response to this information. The magnitude of the regression coefficients is greatest for the 23 to 33 age group, which is suggestive of tighter liquidity constraints or a stronger precautionary motive for the young (see section 2.3, although there we argued that precautionary motives should skew upwards the coefficients on the longer lags of consumption growth compared to the shorter lags, which does not fit with the heterogeneity in the patterns of the coefficients). The magnitude of the regression coef-

ficients is second highest for the 43 to 59 age group, possibly due to an increased rate of saving for retirement (again see section 2.3).

The other dimension over which it is fruitful to explore heterogeneity in the nature of consumption smoothing is the educational attainment of household heads. Table 7 reports regression estimates for each of the four educational classifications of the head. Estimation for each group is done separately; the control variables are cohort specific intercepts and a quartic polynomial in age. The panel for each education group contains 67 observations after differencing and computing six lags.

The differences in the patterns in the regression coefficients across education groups show some consistency; the less the education of the household head, the more concentrated the regression coefficients at shorter lags. The only statistically significant regression coefficient for high school dropouts is at lag two; for those who graduated from high school but received no more schooling, we see many more statistically significant coefficients, but again the coefficients peak at lag two. In contrast the coefficients of heads with some college peak at lag four, and the coefficients of those with at least a four year college degree show the largest coefficients at lags four and five. The data indicate that college graduates obtain information about their future income growth at least four and five years in advance, and alter their consumption growth in response to this information. Either high school dropouts are not learning about their income growth as far in advance, or, if they do, something is preventing them from acting to change their consumption growth in as timely a fashion as the college graduates.

5 Conclusion

That households are able to forecast variation in their income growth years in advance, variation largely orthogonal to hump-shaped life cycle income profiles, may at first seem implausible to some economists. Consideration of some of the known facts about the time series process governing household's income growth argues against this skepticism.

Take first the gap in income between the most and least educated, whose fairly continuous expansion in the late 1980s and 1990s is forecast four, five and six years in advance by the consumption growth data in our sample. The gap in wages between college graduates and high school graduates began its long expansion around 1979; in general, this gap experiences long, persistent swings, increasing or decreasing consistently for a half decade, decade or more - see Katz and Murphy (1992). Given that relative wage and income growth across these groups evidently contains an very persistent component,

is not surprising that households can forecast relative income growth years in advance. Moreover, the rising returns to education had already started prior to the first year of consumption data used in the paper - 1980; it should not come as a complete shock that our consumption data can forecast an already pre-existing phenomenon. It is remarkable, though, that households seemed to understand fairly early on that the increasing returns to education were going to continue throughout the 1980s and 1990s, and that they altered their consumption growth in the manner predicted by the basic theory of consumption smoothing.

Next, consider the traditional business cycle variation in income growth forecast as much as six years in advance by consumption growth. In assessing the plausibility of this result, table 2, which shows autocorrelations of the business cycle component of income growth, warrants another look. It shows that, to forecast business cycle variation in income growth six years in advance and at a reasonable level of statistical reliability, one need look no further than at its own lags. There is a slow mean reversion to the business cycle variation in income growth that is not exactly hidden; the fact that households can forecast *some* business cycle variation in income growth years in advance does not, then, seem fantastic.

What is more suspect, in the view of this author, is the interpretation under the certainty equivalent benchmark model of the relative magnitudes of the regression coefficients. Under this interpretation, where the β s are moving average coefficients on the time series process governing the observable dimension to consumers' information, the second specification of table 4 says that more business cycle information about income growth is revealed to consumers five years in advance of income growth's arrival than is revealed by income growth's actual arrival. If true, this is really something, but I am sympathetic to the view that this is implausible on its face, and argue that it is fruitful to search for other interpretations of the regression coefficients.

One potentially fruitful area of inquiry is into the effect of precautionary motives on the regression coefficients in (1), a preliminary discussion of which was provided in section 2.3. There we argued, using the logic of Carroll (1992, 1997), that, compared to the certainty equivalent benchmark, precautionary motives should increase β_k relative to β_j for $k > j$; in other words, precautionary motives should skew upwards the coefficients on the longer lags of consumption growth more than the coefficients on the shorter lags. The reason is essentially that a household's decision to smooth its consumption in response to news about expected future income will have consequences for the household's liquid asset holdings for the length of time between the arrival of the news and the expected

arrival of the future income, altering the household's ability to smooth consumption over that entire time period. Section 2.3 argued that this effect was likely to dampen the initial response of consumption growth to the news; the greater the length of time between the arrival of the news and the arrival of the future income, the larger this dampening effect and the larger the upward skewing of the β .

Note that the skewing upwards of the coefficients on the longer lags of consumption growth is not a feature of all of the popular extensions to the benchmark certainty equivalent model - it does not appear to be a consequence of non-separability between consumption and leisure, for example. Nor will it appear in some other ad-hoc models considered in the literature, for example where a fixed fraction of consumers eat their income each period. This raises the possibility that some of the more nuanced implications of the different extensions to the benchmark model may help distinguish their relative importance in explaining the important stylized facts about consumer behavior. Since there is still a large degree of uncertainty in the profession regarding the relative importance of the popular extensions, progress along these lines could provide a significant contribution to our understanding of household behavior.

Appendix A: Derivation of the Log-Linear Budget Constraint

This section begins by log-linearizing each side of (3) in turn. Consider first the log-linearization of $W_t^h = \sum_{j=0}^{\infty} \frac{C_{t+j}^h}{(1+r)^j}$. Divide each side by C_t^h and take logs, letting lower case variables denote variables for which logs have been taken:

$$\begin{aligned} c_t^h - w_t^h &= -\ln \left(1 + \frac{1}{(1+r)} \frac{C_{t+1}^h}{C_t^h} + \frac{1}{(1+r)^2} \frac{C_{t+2}^h}{C_t^h} + \dots \right) \\ &= -\ln \left(1 + \frac{1}{(1+r)} \exp(\Delta c_{t+1}^h) + \frac{1}{(1+r)^2} \exp \left(\sum_{k=1}^2 \Delta c_{t+k}^h \right) + \dots \right) \end{aligned}$$

Take a Taylor series expansion of the expression on the right with respect to the consumption growth rates, around the points of zero growth. This yields:

$$\begin{aligned} \ln \left(1 + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \exp \left(\sum_{k=1}^j \Delta c_{t+k}^h \right) \right) &\approx \ln \left(\frac{1+r}{r} \right) + \left(\frac{r}{1+r} \right) \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \sum_{k=1}^j \Delta c_{t+k}^h \\ &= \ln \left(\frac{1+r}{r} \right) + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \Delta c_{t+j}^h, \end{aligned}$$

which can be substituted into the previous equation to arrive at:

$$c_t^h \approx w_t^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \Delta c_{t+j}^h - \left(\frac{1+r}{r} \right). \quad (17)$$

Next consider the log-linearization of $W_t^h = A_t^h + \sum_{j=0}^{\infty} \frac{X_{t+j}^h}{(1+r)^j}$. In this expression for wealth, labor income is expressed as a discounted sum of future dividends, although it could also have been expressed more compactly as the cum-dividend price of human capital: $X_t^h + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} X_{t+j}^h = X_t^h + \frac{1}{1+r} H_{t+1}^h = H_t^h$. Asset income is expressed as such a cum-dividend price in the our current expression for wealth; consider breaking it up into a stream of future income flows from liquid assets - i.e. dividends. One way to do this is to write: $A_t^h = D_t^h + \frac{1}{1+r} A_{t+1}^h = D_t^h + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} D_{t+j}^h$, where the dividend at each time period is $(1 - \frac{1}{1+r}) A_t^h$. Campbell and Mankiw (1989) suggest that the level of liquid assets can be broken up in such a way, and suggest pooling together the income flows from human capital and liquid assets. Write the log of the left hand side of (3) as:

$$\begin{aligned} w_t^h &= \ln \left(X_t^h + D_t^h + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} (X_{t+j}^h + D_{t+j}^h) \right) \\ &= \ln \left(\sum_{j=0}^{\infty} \frac{1}{(1+r)^j} Y_{t+j}^h \right), \end{aligned}$$

where Y_t^h is total “dividend payments” from human capital and non-human capital. Following the same steps as the previous log-linearization yields:

$$w_t^h \approx y_t^h + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \Delta y_{t+j}^h + \ln \left(\frac{1+r}{r} \right). \quad (18)$$

Next take the time t conditional expectation of (17); this equation and its lead are:

$$\begin{aligned} c_{t+1}^h &\approx E_{t+1} w_{t+1}^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} E_{t+1} \Delta c_{t+1+j}^h - \ln \left(\frac{1+r}{r} \right) \\ c_t^h &\approx E_t w_t^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} E_t \Delta c_{t+j}^h - \ln \left(\frac{1+r}{r} \right). \end{aligned} \quad (19)$$

We’d like to write the expression for c_{t+1}^h as a function of w_t^h . We can do this using the approximate law of motion for wealth derived in Campbell and Mankiw (1989) and Campbell (1993): $w_{t+1}^h = r + k + \left(\frac{1}{\lambda}\right) (w_t^h) + \left(1 - \frac{1}{\lambda}\right) (c_t^h)$, where $\lambda = 1 - \overline{C}/\overline{W}$ is one minus the mean consumption-wealth ratio.³ Notice that if we take the unconditional mean of the right hand side of (3), we get $\frac{1}{1+r} = \lambda$. So these are interchangeable; the paper goes with the λ notation from now on. Use this law of motion to substitute out w_{t+1}^h from the first equation of (19), then multiply the second equation of (19) by $\frac{1}{\lambda}$ and add $c_t^h - \frac{c_t^h}{\lambda}$ to both sides. These manipulations yield the two equation system:

$$\begin{aligned} c_{t+1}^h &= \left(\frac{1}{\lambda}\right) E_{t+1} (w_t^h) + \left(1 - \frac{1}{\lambda}\right) (c_t^h) - \sum_{j=1}^{\infty} \lambda^j E_{t+1} (\Delta c_{t+j+1}^h) + \ln(1-\lambda) + r + k \\ c_t^h &= \left(\frac{1}{\lambda}\right) E_t (w_t^h) + \left(1 - \frac{1}{\lambda}\right) (c_t^h) - \left(\frac{1}{\lambda}\right) \sum_{j=1}^{\infty} \lambda^j E_t (\Delta c_{t+j}^h) + \left(\frac{1}{\lambda}\right) \ln(1-\lambda). \end{aligned}$$

Subtracting the first equation from the second yields, after some rearrangements:⁴

$$\Delta c_{t+1}^h \approx E_t \Delta c_{t+1}^h - \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) \Delta c_{t+j+1}^h + \left(\frac{1}{\lambda}\right) (E_{t+1} - E_t) w_t^h.$$

Finally, use (18) to substitute income growth terms for the innovation in w_t^h . This yields our decomposition of consumption growth:

$$\Delta c_{t+1}^h = E_t \Delta c_{t+1}^h - \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) \Delta c_{t+j+1}^h + \sum_{j=1}^{\infty} \lambda^{j-1} (E_{t+1} - E_t) \Delta y_{t+j}^h.$$

³The constant $k = \ln(\lambda) - \left(1 - \frac{1}{\lambda}\right) \ln(1-\lambda)$.

⁴Here we’ve used $r + \ln(\lambda) \approx 0$.

Appendix B: Computation of the Standard Errors

Let $N\bar{T} = \sum_{j=1}^N T_j$ denote the total sample size where j indexes synthetic persons and T_j denotes the number of annual observations for synthetic individual j , and let \mathbf{X} denote the $N\bar{T} \times K$ matrix of regressors, \mathbf{Z} denote the (also) $N\bar{T} \times K$ matrix of instruments, Ω represent the $N\bar{T} \times N\bar{T}$ variance-covariance matrix of regression residuals, and finally let \mathbf{W} denote the $N\bar{T} \times N\bar{T}$ diagonal matrix with the vector of weights on the diagonal.⁵ Then, following standard practice, the variance-covariance matrix of the OLS parameter estimates is computed as:

$$(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X}(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}.$$

Here Ω is likely to exhibit both both heteroskedasticity and certain forms of dependence. If taste shocks are correlated across individuals, which seems likely and certainly cannot be ruled out a priori, Ω will exhibit contemporaneous cross correlation.⁶ In addition, the use of overlapping k th differences introduces correlation between the time t residual and the time $t - k$ residual for any synthetic person.

$\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X}$ is the sum of two additive components. The first component ($\mathbf{X}'\mathbf{W}\Omega_0\mathbf{W}\mathbf{X}$, say,) is robust to both arbitrary heteroskedasticity and cross correlation. The matrix can be represented as:

$$\mathbf{X}'\mathbf{W}\Omega_0\mathbf{W}\mathbf{X} = \sum_{i=1}^N \sum_{j=1}^N \sum_{t=1}^{T_{ij}} (x'_{i,t} w_{i,t} u_{i,t} u_{j,t} w_{j,t} x_{j,t} + x'_{i,t} w_{j,t} u_{j,t} u_{i,t} w_{i,t} x_{i,t}),$$

where T_{ij} is the number of periods where there is an observation for both i and j , $w_{i,t}$ is the weight for the t th observation on the i th synthetic person, and $x_{i,t}$ corresponds to the appropriate row vector of \mathbf{X} . Rearranging summations yields the convenient expression for computing this matrix used in this paper:

$$\begin{aligned} \mathbf{X}'\mathbf{W}\Omega_0\mathbf{W}\mathbf{X} &= \sum_{t=1}^T \sum_{i=1}^{N_t} \sum_{j=1}^{N_t} (x'_{i,t} w_{i,t} u_{i,t} u_{j,t} w_{j,t} x_{j,t} + x'_{i,t} w_{j,t} u_{j,t} u_{i,t} w_{i,t} x_{i,t}) \\ &= \sum_{t=1}^T \mathbf{X}'_t \mathbf{W}_t \Omega_{0,t} \mathbf{W}_t \mathbf{X}_t. \end{aligned}$$

⁵The weights in this paper are the square root of the synthetic cohort cell counts - i.e. if $w_{i,t}$ is the weight on the i th synthetic person at time t , $w_{i,t}$ is the square root of number of households in the sample at time t who meet the qualifications to belong to the i th synthetic cohort.

⁶The standard errors computed below do rule out cross-autocorrelation.

Here N_t denotes the number of cross sectional observations in year t , while \mathbf{X}_t denotes the N_t rows of \mathbf{X} that correspond to the year t cross sectional observations. Similarly, $\Omega_{0,t}$ denotes the matrix $\mathbf{u}_t\mathbf{u}'_t$, the outer product of the vector of all time t residuals, and \mathbf{W}_t denotes the diagonal weighting matrix for time t observations. The last expression makes clear that, after sorting the data by year, the cross-correlation corrected variance-covariance matrix of residuals will be block diagonal,⁷ with each block corresponding to a year. This variance-covariance matrix has the same form as those used in clustered samples to correct for arbitrary within-cluster correlations (see Deaton (1997), p.76), the only difference being that each year plays the role of a cluster. See also Vuolteenaho (2000), who uses this procedure to construct standard errors in a study of stock market returns.

A second component of the estimated matrix $\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X}$, ($\mathbf{X}'\mathbf{W}\Omega_k\mathbf{W}\mathbf{X}$, say,) corrects for the k th order autocorrelation induced by using k th overlapping differences. This matrix is quite standard and is computed as:

$$\mathbf{X}'\mathbf{W}\Omega_k\mathbf{W}\mathbf{X} = \sum_{j=1}^N \sum_{t=1+k}^{T_j} (x'_{j,t}w_{j,t}u_{j,t}u_{j,t-k}w_{j,t-k}x_{j,t-k} + x'_{j,t-k}w_{j,t-k}u_{j,t-k}u_{j,t}w_{j,t}x_{j,t}),$$

The full $\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X}$ is then computed as

$$\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X} = \mathbf{X}'\mathbf{W}\Omega_0\mathbf{W}\mathbf{X} + \sum_{k=1}^{k'} \left(\frac{k'+1-k}{k'+1} \right) \mathbf{X}'\mathbf{W}\Omega_k\mathbf{W}\mathbf{X}.$$

The paper sets k' equal to seven, correcting for seventh order autocorrelation.


⁷Ignoring any autocorrelation for the moment

References

- [1] Attanasio, Orazio, and Davis, Steven J. "Relative Wage Movements and the Distribution of Consumption", *Journal of Political Economy* 104 (1996), pp. 1127-1262.
- [2] Browning, Martin, and Lusardi, Annamaria. "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature* 34 (1996), pp. 1797-1855.
- [3] Campbell, John. "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis," *Econometrica*, 55 (1987), pp. 1249-1273.
- [4] Campbell, John. "Intertemporal Asset Pricing without Consumption Data." *American Economic Review* 83 (1993), pp. 487-512.
- [5] Campbell, John, and Deaton, Angus. "Why is Consumption So Smooth?" *Review of Economic Studies*, 56 (1989), pp. 357-374.
- [6] Campbell, John, and Mankiw, Greg. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual* (1989), pp. 185-245.
- [7] Carroll, Chris. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* (1992), pp. 61-156.
- [8] Carroll, Chris. "How Does Future Income Affect Current Consumption?" *Quarterly Journal of Economics*, 109 (1994), pp. 111-148.
- [9] Carroll, Chris. "Buffer Stock Saving and the Life-Cycle Permanent Income Hypothesis." *Quarterly Journal of Economics*, 112 (1997), pp. 1-55.
- [10] Cochrane, John. "How Big is the Random Walk in GNP?," *Journal of Political Economy*, 96 (1988), pp. 893-920.
- [11] Cochrane, John. "Permanent and Transitory Components of GNP and Stock-Prices," *Quarterly Journal of Economics*, 109 (1994), pp. 241-265.
- [12] Cutler, David M., and Katz, Lawrence F. "Macroeconomic Performance and the Disadvantaged", *Brookings Papers Econ. Activity*, no. 2 (1991), pp. 1-61.

- [13] Deaton, Angus. "Panel Data from a Time Series of Cross Sections," *Journal of Econometrics*, 30 (1985), pp. 109-126.
- [14] Deaton, Angus. *Understanding Consumption*. Clarendon Press, 1992.
- [15] Deaton, Angus. *The Analysis of Household Surveys, A Microeconomic Approach to Development Policy*. Published by the World Bank, Johns Hopkins University Press, 1997 (third printing, 2000).
- [16] Fuller, Wayne. *Measurement Error Models*. John Wiley & Sons, 1987.
- [17] Ghez, Gilbert; and Becker, Gary. *The Allocation of Time and Goods Over the Life Cycle*. NBER, 1975.
- [18] Gourinchas, P. -O., and Parker, J. A. "Consumption Over the Lifecycle," working paper, Princeton University (1999).
- [19] Hansen, Lars Peter, Roberds, William, and Sargent, Tom. "Time Series Implications of Present Value Budget Balance and of Martingale Models of Consumption and Taxes," in *Rational Expectations Econometrics* eds.: Lars P. Hansen and Thomas Sargent, Boulder. Westview, 1991, pp. 121-161.
- [20] Katz, Lawrence, and Murphy, Kevin M. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics* (1992), pp. 35-78.
- [21] Newey, Whitney K., and West, Kenneth D. "A Simple, Positive Semi-definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55 (1987), pp. 703-708.
- [22] Sargent, Thomas. *Macroeconomic Theory*, Academic Press, 1987.
- [23] Vuolteenaho, Tuomo. "What Drives Firm Level Stock Returns?" Phd dissertation, The University of Chicago Graduate School of Business, 2000.

**Figure 1: Actual (Dashed) and Predicted (Solid) CPS Income Growth
Estimation from 7 lag specification with year dummies as controls**



**Figure 2: Actual (Dashed) and Predicted (Solid) CPS Income Growth
Estimation from 7 lag specification with lagged CPS income growth as controls**

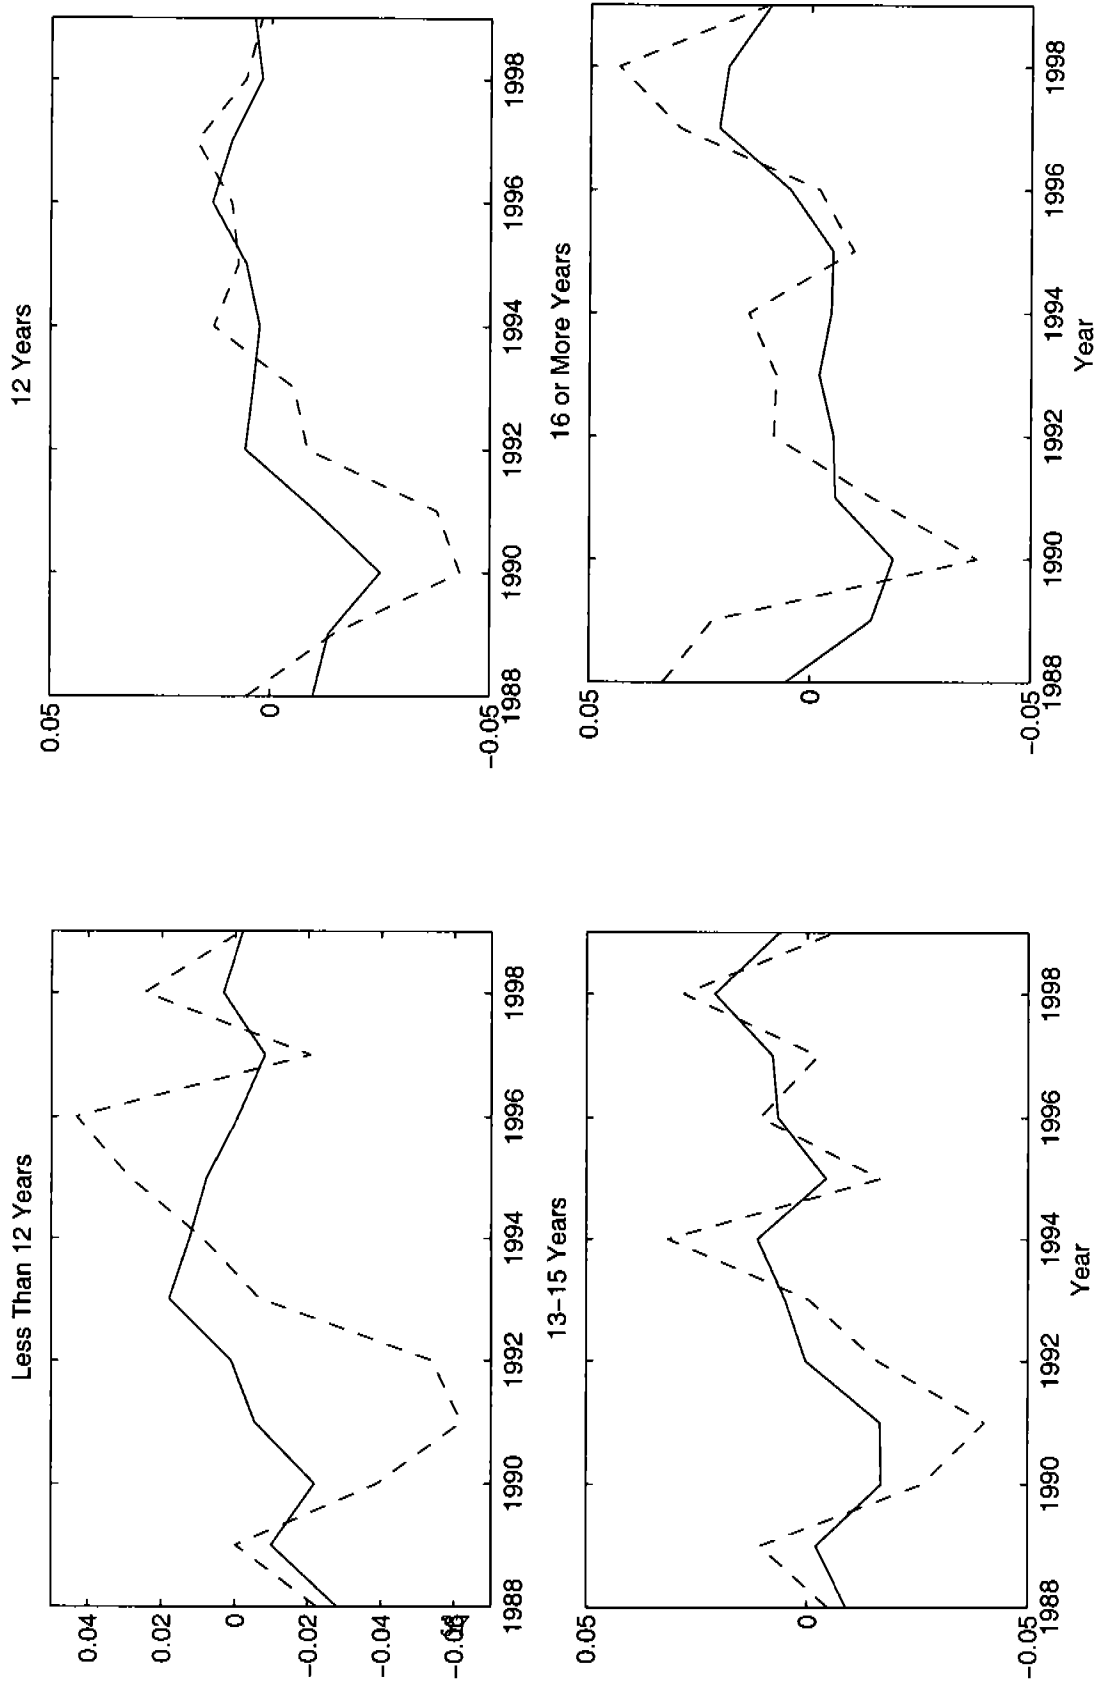


Table 3:

Panel A: Estimation of $\Delta y_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + e_t$

q	β_0	β_1	β_2	β_3	β_4	β_5	β_6	R^2	obs
0	0.17 (0.04)							0.07	484
1	0.16 (0.04)	0.12 (0.03)						0.07	444
2	0.13 (0.04)	0.12 (0.03)	0.16 (0.03)					0.11	408
3	0.13 (0.04)	0.14 (0.03)	0.18 (0.04)	0.12 (0.03)				0.15	372
4	0.10 (0.04)	0.13 (0.03)	0.19 (0.04)	0.14 (0.03)	0.09 (0.03)			0.16	336
5	0.06 (0.04)	0.09 (0.04)	0.18 (0.04)	0.16 (0.03)	0.13 (0.05)	0.12 (0.04)		0.18	300
6	0.08 (0.05)	0.07 (0.04)	0.16 (0.04)	0.15 (0.03)	0.16 (0.05)	0.14 (0.04)	0.04 (0.05)	0.17	268

Panel B: Estimation of $\Delta c_t = \beta_0 \Delta y_t + \beta_1 \Delta y_{t-1} + \dots + \beta_q \Delta y_{t-q} + e_t$

q	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	R^2	obs
0	0.40 (0.09)								0.07	484
1	0.26 (0.07)	0.35 (0.06)							0.10	444
2	0.23 (0.07)	0.36 (0.08)	0.02 (0.07)						0.09	408
3	0.24 (0.07)	0.35 (0.09)	0.01 (0.09)	0.01 (0.07)					0.08	372
4	0.20 (0.08)	0.31 (0.10)	0.03 (0.09)	0.01 (0.09)	0.11 (0.10)				0.07	336
5	0.18 (0.08)	0.28 (0.11)	0.02 (0.09)	-0.03 (0.10)	0.10 (0.11)	0.07 (0.11)			0.06	300
6	0.21 (0.09)	0.28 (0.11)	0.05 (0.09)	-0.06 (0.10)	0.01 (0.15)	-0.06 (0.11)	0.22 (0.10)		0.07	268
7	0.20 (0.11)	0.26 (0.12)	0.09 (0.09)	-0.08 (0.10)	0.03 (0.17)	-0.07 (0.15)	0.22 (0.13)	0.02 (0.17)	0.07	240

Notes to Table 3: Annual synthetic panel data grouped by 5-year birth cohort and a 4-way educational classification, 1980-1999; 36 synthetic individuals, 524 observations before differencing and lagging. The Δy_t terms are income computed from the CPS; the explanatory variables Δc_{t-k} are consumption computed from the CEX. The regressions are weighted by the average cell counts of the explanatory variables. Asymptotic standard errors are in parentheses; details of their computation are in Appendix B.

Table 4:

Weighted Least Squares Estimation of

$$\Delta y_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + \text{controls} + e_t$$

Controls for:	Estimation	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	$\sum_{j=1}^q \lambda^j \beta_j$	Implied τ	$R^2_{\Delta c}$
	WLS, $q = 6$	0.08 (0.05)	0.07 (0.04)	0.16 (0.04)	0.15 (0.03)	0.16 (0.05)	0.14 (0.04)	0.04 (0.05)		0.75 (0.05)	-0.07 (0.02)	0.17
Ag ^C , Cohorts	WLS, $q = 6$	0.09 (0.06)	0.12 (0.06)	0.22 (0.05)	0.21 (0.06)	0.22 (0.06)	0.18 (0.06)	0.08 (0.06)		1.05 (0.25)	0.04 (0.08)	0.07
Years	WLS, $q = 6$	0.05 (0.04)	0.05 (0.04)	0.10 (0.03)	0.11 (0.03)	0.12 (0.04)	0.17 (0.03)	0.15 (0.04)		0.67 (0.06)	-0.08 (0.02)	0.20
	WLS, $q = 7$	0.02 (0.05)	0.02 (0.04)	0.09 (0.03)	0.09 (0.04)	0.15 (0.04)	0.20 (0.04)	0.15 (0.04)	0.03 (0.04)	0.67 (0.07)	-0.07 (0.02)	0.18
Years, Ag ^C , Cohorts	WLS, $q = 6$	0.02 (0.05)	0.03 (0.05)	0.07 (0.05)	0.08 (0.05)	0.11 (0.05)	0.15 (0.04)	0.13 (0.04)		0.54 (0.16)	-0.12 (0.05)	0.04
	WLS, $q = 7$	-0.00 (0.06)	0.01 (0.06)	0.11 (0.05)	0.12 (0.02)	0.20 (0.06)	0.25 (0.04)	0.16 (0.03)	0.02 (0.05)	0.79 (0.18)	-0.03 (0.05)	0.07
Δy_{t-j}^{CEX}	WLS, $q = 6$	0.08 (0.03)	0.01 (0.02)	0.18 (0.04)	0.21 (0.04)	0.27 (0.05)	0.27 (0.06)	0.16 (0.03)		1.09 (0.09)	0.05 (0.02)	0.25
	WLS, $q = 7$	0.06 (0.04)	0.00 (0.02)	0.19 (0.04)	0.20 (0.04)	0.28 (0.05)	0.28 (0.06)	0.14 (0.02)	-0.06 (0.04)	1.00 (0.09)	0.02 (0.03)	0.25
Δy_{t-j}^{CPS}	WLS, $q = 6$	0.11 (0.04)	0.11 (0.04)	0.21 (0.04)	0.24 (0.04)	0.25 (0.05)	0.26 (0.04)	0.16 (0.04)		1.24 (0.13)	0.10 (0.04)	0.24
	WLS, $q = 7$	0.11 (0.04)	0.14 (0.04)	0.23 (0.04)	0.26 (0.04)	0.31 (0.05)	0.31 (0.05)	0.23 (0.05)	0.05 (0.04)	1.50 (0.15)	0.16 (0.04)	0.29
Δy_{t-j}^{CEX} , Ag ^C , Cohorts	WLS, $q = 6$	0.02 (0.04)	-0.06 (0.05)	0.11 (0.05)	0.14 (0.05)	0.23 (0.04)	0.24 (0.06)	0.15 (0.04)		0.76 (0.18)	-0.04 (0.05)	0.12
	WLS, $q = 7$	0.01 (0.05)	-0.04 (0.05)	0.15 (0.05)	0.17 (0.03)	0.28 (0.05)	0.28 (0.07)	0.12 (0.04)	-0.07 (0.05)	0.81 (0.14)	-0.03 (0.04)	0.17
Δy_{t-j}^{CPS} , Ag ^C ,	WLS, $q = 6$	0.14 (0.04)	0.17 (0.05)	0.28 (0.05)	0.33 (0.07)	0.34 (0.06)	0.33 (0.06)	0.20 (0.04)		1.65 (0.22)	0.21 (0.06)	0.16

Cohorts	WLS, $q = 7$	0.09 (0.04)	0.15 (0.05)	0.28 (0.04)	0.34 (0.05)	0.42 (0.05)	0.43 (0.07)	0.30 (0.06)	0.08 (0.05)	1.91 (0.17)	0.24 (0.04)	0.20
Δy_{t-j}^{CEX} , Years	WLS, $q = 6$	0.06 (0.04)	-0.01 (0.04)	0.10 (0.04)	0.12 (0.04)	0.13 (0.04)	0.21 (0.05)	0.20 (0.03)		0.74 (0.12)	-0.05 (0.03)	0.14
	WLS, $q = 7$	0.03 (0.05)	-0.04 (0.05)	0.11 (0.03)	0.10 (0.04)	0.14 (0.04)	0.20 (0.06)	0.16 (0.01)	-0.03 (0.04)	0.61 (0.15)	-0.09 (0.04)	0.11
Δy_{t-j}^{CPS} , Years	WLS, $q = 6$	0.09 (0.03)	0.10 (0.04)	0.14 (0.05)	0.17 (0.04)	0.17 (0.04)	0.19 (0.04)	0.19 (0.04)		0.96 (0.12)	0.01 (0.04)	0.17
	WLS, $q = 7$	0.06 (0.03)	0.07 (0.04)	0.13 (0.05)	0.15 (0.04)	0.20 (0.04)	0.26 (0.04)	0.24 (0.05)	0.12 (0.04)	1.11 (0.16)	0.05 (0.04)	0.20

Notes to Table 4: Annual synthetic panel data grouped by 5-year birth cohort and a 4-way educational classification, 1980-1999; 36 synthetic individuals, 524 observations before differencing and lagging. The Δy_t terms are income computed from the CPS; the explanatory variables Δc_{t-k} are consumption computed from the CEX. The controls for age are quartic age polynomials interacted with education fixed effects. The controls for cohorts are a set of 5-year birth cohort fixed effects interacted with education fixed effects. The controls for years are a set of year fixed effects. The controls Δy_{t-j}^{CEX} are current income growth and q lags of income growth, computed from the CEX. The controls Δy_{t-j}^{CPS} are q lags of income growth computed from the CPS. The regressions are weighted by the average cell counts of the explanatory consumption growth variables. Asymptotic standard errors are in parentheses; details of their computation are in Appendix B.

In computing the discounted sum of the regression coefficients in column 11, $\lambda = \frac{1}{1+r}$ is set to $\frac{1}{1.025}$. The implied r in column 12 is the value of r for which the discounted sum of the regression coefficients sums to one. The $R_{\Delta c}^2$ in column 13 is fraction of variance in Δy_t explained by consumption growth, after orthogonalizing the data with respect to the control variables.

Table 5:
Weighted Errors-in-Variables, Bootstrap Estimation of
 $\Delta y_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + \text{controls} + e_t$

Controls for:	Estimation	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	$\sum_{j=1}^q \lambda^j \beta_j$
	WLS, $q = 6$	0.06 (0.03)	0.09 (0.02)	0.13 (0.03)	0.13 (0.03)	0.12 (0.03)	0.11 (0.03)	0.05 (0.03)		
	WEIV, $q = 6$	0.08 (0.06)	0.09 (0.05)	0.16 (0.06)	0.15 (0.05)	0.14 (0.07)	0.14 (0.06)	0.04 (0.06)		0.74 (0.07)
Ag ^C , Cohorts	WEIV, $q = 6$	0.11 (0.07)	0.17 (0.09)	0.26 (0.09)	0.26 (0.10)	0.24 (0.09)	0.20 (0.08)	0.10 (0.07)		1.26 (0.40)
Years	WEIV, $q = 6$	0.06 (0.06)	0.07 (0.05)	0.10 (0.06)	0.11 (0.06)	0.10 (0.07)	0.16 (0.06)	0.12 (0.06)		0.66 (0.07)
	WEIV, $q = 7$	0.04 (0.08)	0.05 (0.06)	0.10 (0.06)	0.09 (0.06)	0.12 (0.08)	0.18 (0.07)	0.11 (0.06)	0.04 (0.06)	0.65 (0.07)
Years, Ag ^C , Cohorts	WEIV, $q = 6$	0.04 (0.08)	0.06 (0.10)	0.08 (0.12)	0.09 (0.13)	0.09 (0.12)	0.13 (0.09)	0.09 (0.08)		0.54 (0.56)
	WEIV, $q = 7$	0.05 (0.09)	0.11 (0.12)	0.19 (0.15)	0.21 (0.17)	0.26 (0.18)	0.28 (0.15)	0.17 (0.11)	0.07 (0.08)	1.22 (0.83)
Δy_{t-j}^{CEX}	WEIV, $q = 6$	0.08 (0.06)	0.07 (0.05)	0.16 (0.06)	0.20 (0.07)	0.21 (0.08)	0.22 (0.07)	0.11 (0.07)		0.97 (0.12)
	WEIV, $q = 7$	0.07 (0.06)	0.11 (0.06)	0.30 (0.07)	0.41 (0.08)	0.57 (0.12)	0.52 (0.11)	0.41 (0.10)	0.22 (0.08)	2.36 (0.41)
Δy_{t-j}^{CPS}	WEIV, $q = 6$	0.10 (0.06)	0.11 (0.06)	0.20 (0.06)	0.23 (0.07)	0.22 (0.08)	0.25 (0.08)	0.14 (0.07)		1.15 (0.16)
	WEIV, $q = 7$	0.12 (0.08)	0.13 (0.06)	0.21 (0.06)	0.24 (0.07)	0.27 (0.10)	0.28 (0.09)	0.19 (0.08)	0.04 (0.08)	1.37 (0.20)
Δy_{t-j}^{CEX} , Ag ^C , Cohorts	WEIV, $q = 6$	0.04 (0.08)	0.02 (0.10)	0.12 (0.11)	0.16 (0.11)	0.19 (0.11)	0.20 (0.08)	0.10 (0.07)		0.77 (0.47)
	WEIV, $q = 7$	0.04 (0.04)	0.04 (0.05)	0.11 (0.06)	0.14 (0.05)	0.17 (0.05)	0.16 (0.05)	0.07 (0.04)	-0.02 (0.04)	0.65 (0.25)

Δy_{t-j}^{CPS} , Age ^c , Cohorts	WEIV, $q = 6$	0.13 (0.07)	0.21 (0.10)	0.32 (0.11)	0.36 (0.11)	0.35 (0.11)	0.33 (0.08)	0.20 (0.07)	1.76 (0.47)
	WEIV, $q = 7$	0.12 (0.07)	0.21 (0.10)	0.34 (0.12)	0.41 (0.12)	0.48 (0.14)	0.46 (0.12)	0.32 (0.10)	2.27 (0.62)
Δy_{t-j}^{CEX} , Years	WEIV, $q = 6$	0.06 (0.07)	0.04 (0.06)	0.10 (0.06)	0.12 (0.07)	0.10 (0.08)	0.17 (0.07)	0.13 (0.07)	0.66 (0.15)
	WEIV, $q = 7$	0.03 (0.04)	0.04 (0.04)	0.10 (0.04)	0.11 (0.05)	0.11 (0.05)	0.13 (0.04)	0.09 (0.04)	0.57 (0.17)
Δy_{t-j}^{CPS} , Years	WEIV, $q = 6$	0.09 (0.06)	0.11 (0.06)	0.15 (0.06)	0.16 (0.07)	0.15 (0.07)	0.18 (0.07)	0.14 (0.06)	0.90 (0.18)
	WEIV, $q = 7$	0.06 (0.07)	0.10 (0.07)	0.15 (0.06)	0.15 (0.07)	0.19 (0.08)	0.23 (0.07)	0.18 (0.07)	1.05 (0.23)

Notes to Table 5: Annual synthetic panel data grouped by 5-year birth cohort and a 4-way educational classification, 1980-1999, 36 synthetic individuals, 524 observations before differencing and lagging. The Δy_t terms are income computed from the CPS; the explanatory variables Δc_{t-k} are consumption computed from the CEX. The controls for age are quartic age polynomials interacted with education fixed effects. The controls for cohorts are a set of 5-year birth cohort fixed effects interacted with education fixed effects. The controls for years are a set of year fixed effects. The controls Δy_{t-j}^{CEX} are current income growth and q lags of income growth, computed from the CEX. The controls Δy_{t-j}^{CPS} are q lags of income growth computed from the CPS.

The reported statistics are means and standard deviations of estimates from 200 bootstrapped synthetic panels. The weighted errors-in-variables (WEIV) estimates are of the form: $\beta = [X'X - \gamma(\epsilon_x' \epsilon_x)]^{-1} X'Y$, where where X is the matrix of regressors measured with error, Y is the vector of CPS income growth, $\epsilon_x' \epsilon_x$ is an estimate of the variance-covariance matrix of the sampling errors, and γ is a number defined in section 3.3. The estimator is computed in two stages: the first stage produces unweighted estimates used to estimate the heteroskedastic variance of the residuals, and the second stage weights the regression matrices by the inverse of the estimated variance of the residuals. See section 3.3 and section 4.3 for further details. The weighted least squares (WLS) estimates are ordinary least squares regression estimates weighted by the average cell counts of the explanatory consumption growth variables.

In computing the discounted sum of the regression coefficients in column 11, $\lambda = \frac{1}{1+r}$ is set to $\frac{1}{1.025}$.

Table 6:

**Estimation of $\Delta y_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + \text{controls} + e_t$,
with Age Heterogeneity**

Age Cutoff	β_0	β_1	β_2	β_3	β_4	β_5	β_6
Age \leq 33	0.48 (0.17)	0.29 (0.15)	0.37 (0.11)	0.33 (0.08)	0.28 (0.13)	0.18 (0.08)	0.11 (0.08)
Age \leq 38	0.22 (0.11)	0.08 (0.06)	0.19 (0.08)	0.13 (0.06)	0.10 (0.04)	0.11 (0.07)	0.09 (0.06)
Age \leq 43	0.14 (0.09)	0.17 (0.07)	0.25 (0.07)	0.22 (0.06)	0.15 (0.06)	0.13 (0.06)	0.07 (0.06)
Age $>$ 33	0.07 (0.07)	0.11 (0.07)	0.20 (0.07)	0.20 (0.07)	0.24 (0.08)	0.21 (0.07)	0.07 (0.06)
Age $>$ 38	0.03 (0.06)	0.09 (0.07)	0.20 (0.05)	0.24 (0.05)	0.31 (0.08)	0.25 (0.06)	0.05 (0.08)
Age $>$ 43	0.05 (0.08)	0.05 (0.09)	0.22 (0.09)	0.19 (0.12)	0.41 (0.09)	0.37 (0.08)	0.13 (0.11)

Table 7:

**Estimation of $\Delta y_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + \text{controls} + e_t$,
with Education Heterogeneity**

Years of Schooling	β_0	β_1	β_2	β_3	β_4	β_5	β_6	$\sum_{j=1}^q \lambda^j \beta_j$
Less than 12	0.09 (0.09)	0.09 (0.06)	0.14 (0.06)	0.08 (0.08)	-0.11 (0.18)	0.08 (0.14)	-0.02 (0.14)	0.33 (0.56)
12	0.04 (0.13)	0.08 (0.06)	0.33 (0.07)	0.25 (0.10)	0.23 (0.08)	0.14 (0.08)	0.12 (0.07)	1.10 (0.21)
13-15	0.18 (0.08)	0.18 (0.12)	0.19 (0.09)	0.21 (0.08)	0.29 (0.10)	0.18 (0.07)	0.11 (0.04)	1.25 (0.22)
16 or More	0.03 (0.09)	0.02 (0.10)	0.09 (0.11)	0.20 (0.08)	0.36 (0.09)	0.28 (0.08)	0.06 (0.10)	0.94 (0.43)

Notes to Tables 6, 7: Annual synthetic panel data grouped by 5-year birth cohort and a 4-way educational classification, 1980-1999; 36 synthetic individuals, 524 observations before differencing and lagging. The Δy_t terms are income computed from the CPS; the explanatory variables Δc_{t-k} are consumption computed from the CEX. In computing the sum of the regression coefficients, $\lambda = \frac{1}{1+r}$ is set to $\frac{1}{1.025}$. Observations are weighted by the average cell counts of the explanatory variables. Asymptotic standard errors are in parentheses; details of their computation are in Appendix B.

Table 6 reports six sets of regression coefficients from three regression specifications with varying age cutoff values. A five year birth cohort is assigned the age of its middle-aged members, and for a given age cutoff (either 33, 38, or 43), two sets of consumption growth terms are included in the regression: one interacted with dummies for whether the cohort was younger than the cutoff at the time of the consumption growth, and another interacted with dummies for whether the cohort was older. Controls are quartic age polynomials and birth cohort fixed effects interacted with education fixed effects.

In table 7, the estimation is done separately for each education group, with quartic age polynomials and birth cohort fixed effects as controls.

Data Appendix

Current Population Survey (CPS) Data

The CPS data are from the March Annual Demographic Files for survey years 1981 to 2000.¹ The BLS usually interviews around 60,000 households in any given Annual Demographic File. As stated in the text, the decision-making unit in this study is taken to be the household, and household income data is the primary CPS data employed in the study. The CPS defines a household in the following way:

A household consists of all the persons who occupy a house, an apartment, or other group of rooms, or a room, which constitutes a housing unit. A group of rooms or a single room is regarded as a housing unit when it is occupied as separate living quarters; that is, when the occupants do not live or eat with any other person in the structure, and when there is direct access from the outside or through a common hall.

Household income, broadly defined, is the sum of earned income, transfer income, asset income, and retirement income, minus state and federal taxes. Household earned income is constructed as the sum of the wage and salary income, self-employed income, and farm income of household members. Before aggregating across household members, the top-coding corrections recommend by Katz & Murphy (1992) are made for each of these three sub-components of earned income for each individual.

Prior to survey year 1990, household level transfer income is computed as the sum of several CPS income variables given at the family level,² summed over the families residing in the household. No correction for top-coding is made for any of these variables. Household level asset income and retirement income are treated in the same way. The list of family-level income variables used from survey years 1981-1989 is:

¹The income data in a survey year refers to the previous calendar year, while the demographics information refers to the current calendar year, so the paper uses twenty years of income data, from 1980 to 1999.

²The CPS defines a family in the following way:

A family is a group of two persons or more (one of whom is the householder) residing together and related by birth, marriage, or adoption. All such persons (including related subfamily members) are considered as members of one family.

**CPS transfer, asset and retirement income variables
for survey years 1981-1989**

- Transfer income
 - FINCUS: family income - money received from U.S. gov't. Includes social security and railroad retirement.
 - FINCSP: family income - supplemental security. Includes money received from U.S., state, and local gov't.
 - FINCPA: family income - public assistance and welfare. Includes aid to families with dependent children and other assistance.
 - FINCVP: family income - veterans payments etc. Includes veterans payments, unemployment compensation, and workers compensation.
 - FINCCS: family income - child support, etc. Includes alimony and child support, other regular contributions from persons not in household, and anything else.

- Asset income
 - FINCINT: family income - interest.
 - FINCDIV: family income - dividends, etc. Includes dividends, net rental income or royalties, estates or trusts.

- Retirement income
 - FINCRET: family income - retirement. Includes private pensions and annuities, military retirement, federal gov't employee pensions, and state or local gov't pensions.

Starting in the 1990 survey year, the CPS began reporting the various components of transfer, asset, and retirement income at the household level as well as the family level; the household level data is used from survey year 1990 onwards. Also in 1990, the CPS changed its classification system for transfer, asset and retirement income. The list of household-level income variables used from 1990 onwards is:

**CPS transfer, asset, and retirement income variables
for survey years 1990-2000**

- Transfer income
 - HSSVAL: HHL D income - Social Security
 - HSURVAL: HHL D income - survivor income
 - HDISVAL: HHL D income - Disability income
 - HSSIVAL: HHL D income - Supplemental Security income
 - HPAWVAL: HHL D income - Public Assistance income
 - HUCVAL: HHL D income - Unemployment compensation
 - HWCVAL: HHL D income - Worker's compensation
 - HVETVAL: HHL D income - Veteran Payments
 - HCSPVAL: HHL D income - child support
 - HALMVAL: HHL D income - alimony
 - HFINVAL: HHL D income - Financial Assistance income
 - HEDVAL: HHL D income - Education income
 - HOIVAL: HHL D income - Other income

- Asset income
 - HINTVAL: HHL D income - Interest income
 - HDIVVAL: HHL D income - dividend income
 - HRNTVAL: HHL D income - Rent income

- Retirement income
 - HRETVAL: HHL D income - Retirement income.

The paper uses the NBER's TAXSIM program to estimate taxes paid. Given demographic and income information on a tax unit in any given year, the TAXSIM program computes that tax unit's state and federal tax burden. Each household is treated as a tax unit, which may bias upwards the tax burden under the progressive U.S. income tax system, given that several returns may be filed separately within a household.

The TAXSIM program requires some data that is fairly straightforward to pull from the CPS files, such as the state of residence of the household, the marital status of the householder,³ the number of children and elderly in the household, the wage and salary income of the householder and spouse, the household's dividend income (FINCDIV and HDIVVAL), its pension income (FINCRET and HRETVAL), and its gross social security income (FINCUS and HSSVAL, HSURVAL, and HDISVAL). The construction of some other variables for TAXSIM requires more decisions. For a variable described as "Other property income, including interest, self-employment, may be negative. Also alimony, fellowships, and other taxable income," we include the household's farm and self-employment income, its interest income (FINCINT and HINTVAL), and some of its transfer income (FINCCS pre-1988 and HCSPVAL, HALMVAL, HFINVAL, and HEDVAL post-1988). In this category, we also include the wage and salary income of members of the household other than the householder and spouse of householder.⁴ For a variable described as "Other non-taxable transfer income such as welfare, municipal bond interest, child support that would affect eligibility for state property tax rebates," we include, before 1988, FINCSP, FINCPA, and FINCVP, and after 1988, HSSIVAL, HPAWVAL, HUCVAL, HWCVAL, and HVETVAL.⁵

Finally, TAXSIM includes some fields where one can input information to compute rebates and deductions, information such as medical expenses, charitable contributions, child care expenses, rent paid, and property taxes paid. Since this information is unavailable from the CPS, we set these fields to zero.⁶

The primary sample of households we consider excludes:

1. households without a male head aged 23 to 59 (without a male householder or a male spouse of householder of that age), and

³Heads of households are called "householders" by the CPS over this time period

⁴Finally, since dividend income can be negative in the CPS, and the TAXSIM program requires it to be strictly positive, negative dividend income is subtracted here and set to zero in the dividend income field.

⁵TAXSIM provides a separate field for unemployment compensation, but prior to 1988 unemployment compensation is not separated from other transfer income in the CPS variable FINCVP. For consistency, the separate field is left at a value of zero for the whole sample and unemployment compensation is always included in the "other non-taxable income" field.

⁶Such information is available for households sampled in the CEX, through. One test of the likely impact of excluding the information on deductions is to run the CEX data through TAXSIM with and without the data on deductions, and compare the different average tax payments for the synthetic cohort data. The difference between the two sets of average tax payments was minor, less than an order of magnitude of the size of the average tax payment for most synthetic cohort cells.

2. households residing in group quarters.

Synthetic cohorts are constructed by the male head of household's education and five-year birth cohort. The birth year variable is designed to run from March to February, reflecting the fact that the survey is taken in March. If any member of a five-year birth cohort violates the sample selection restriction on age in a given year, that cohort is excluded from the sample for that year. For example, a five-year birth cohort may contain heads aged 55, 56, 57, 58 and 59 at year t , in which case it would be excluded from the year t sample. The four education categories are: (i) less than 12 years of schooling, (ii) 12 years of schooling, (iii) more than 12 but less than 16 years of schooling, and (iv) 16 or more years of schooling.

Consumer Expenditure Survey (CEX) Data

The CEX data are taken from the 1980 to 1999 Interview Survey files. The unit of analysis in the CEX is the Consumer Unit (CU), defined as a group of individuals who live together and are either "related by blood, marriage, adoption, or other legal arrangement" or pool expenditures on 2 out of the following 3 expenditure categories: food, housing, or other living expenses. We equate CUs with households throughout this study, and use the terms interchangeably. Each is interviewed up to a maximum of 5 times on a quarterly basis, although no data from the first interview is published on the Interview Survey files. Households continuously rotate in and out of the survey, and about 5,000 households are in the process of being interviewed at any time.

This study primarily employs CEX data on household expenditures and the demographic characteristics of household members. The expenditures data (extracted from the interview survey MTAB files) is monthly, and normally covers each of the three months prior to the month of the interview.⁷ The demographics data (from the interview survey FMLY files) are collected at each interview, and hence are quarterly and current at the time of the interview. These data are converted to a monthly frequency by assigning the data values for a particular interview to the month of the interview and the preceding two months, or the preceding five months if a household skipped the preceding interview or if the interview is the first one for that household. These data are then merged with the data on household expenditures.

⁷In some cases, especially for the last interview, consumption data is also available for the month of the interview in addition to the three preceding months, and occasionally an interview collects information on consumption for four or five preceding months.

A household's non-durables and services consumption is the sum of its expenditures on 16 sub-categories constructed following Orazio Attanasio's classification system.⁸ These sub-categories are:

- food consumed in the home
- food consumed out of the home
- alcohol
- tobacco
- housekeeping services
- home maintenance
- fuel oil, coal, bottled gas, wood, kerosene and other fuels
- electricity and natural or utility gas
- public utilities
- telephone services
- fuel for transportation
- transportation equipment maintenance and repair
- vehicle rental and misc. transportation expenses
- public transportation
- personal care services
- non-durable entertainment expenses.

In 1982 and 1988, there were significant changes in the CEX survey questions covering food consumed at home. From 1980 to 1981, the surveyors asked the household how often it shopped for groceries, and asked what was the usual amount spent per shopping outing. From 1982 to 1987, the surveyors asked for the household's usual monthly

⁸The raw CEX data reports consumption by UCC (universal classification code); there are several hundred of these categories. A list of the UCC codes that comprise each of the sub-categories is available from the author on request.

expense on groceries, and from 1988 onwards they asked for the usual weekly expense. In addition, BLS statisticians have indicated that there were changes in how the data was processed at these break-points. The result was a large spike downwards in the food consumed at home data in 1982, and a large spike upwards in 1988 - see the aggregate data in figure B.1.

To correct these evident data problems, this paper assumes the effect of the survey changes was to scale up or down food at home expenditures of all households by the same amount. This would be the case if some survey regimes cause all households to mis-estimate their frequency of shopping by some fraction, say, 10 percent. These fractions are estimated by regressing household log real food at home expenditures on two dummy variables, one covering the 1980-81 time period, the other 1982-1987, and the log of an explanatory expenditure variable - real non-durable goods and services less food at home.⁹ Real food at home expenditures are deflated by the CPI food at home deflator. The deflator for the explanatory expenditure variable is constructed as a geometric weighted average of the CPI deflators of its 15 sub-categories, with nominal household specific consumption shares as weights (i.e. a Stone price index with household-specific weights). The list CPI deflators that are matched to each CEX expenditure sub-category is:¹⁰

CEX category	CPI categories
food (home)	food at home (SAF11)
food (away from home)	food away from home (SEFV)
alcohol	alcoholic beverages (SAF116)
tobacco	tobacco and smoking products (SEGA)
housekeeping services	household furnishings and operations (SAH3)
home maintenance	housekeeping supplies (SEHN)
fuel oil, coal, etc.	fuel oil and other fuels (SEHE)

⁹criticized as a device for estimating Engel curves, due to its inability to accommodate the budget constraint (see Deaton and Muellbauer (1980)), the specification yields reasonable corrections. A more complicated alternative, not pursued in this paper, would be to estimate corrections in the context of a fully specified demand system.

¹⁰The CPI data were downloaded using the selective access program at www.bls.gov (as of October 2001 the download program name had been changed - it was called "create customized tables (multiple screens)", at www.bls.gov/cpi/home.htm#data). The codes next to each CPI category are the post-1998 revision CPI codes. Where multiple CPI deflators were used for a single category, the price deflator was constructed as an unweighted arithmetic average of the multiple deflators.

electricity and gas	gas (piped) and electricity (SEHF)
public utilities	water and sewerage mainenance (SEHG01) and garbage and trash collection (SEHG02)
telephone services	telephone services, local charges (SEED01) and interstate toll calls (SS27051) and intrastate toll calls (SS27061)
transport fuel	motor fuel (SETB)
transp. maintenance	motor vehicle parts and equipment (SETC) and motor vehicle maintenance and repair (SETD)
misc. transp. expenses	private transportation (SAT1)
public transportation	public transportation (SETG)
personal care services	personal care products (SEGB) and personal care services (SEGC)
entertainment expenses	admissions (SERF02) and fees for lessons or instructions (SERFO3)

Unfortunately, this specification does not capture well a major feature of the data: the steady decline over time in the share of food at home in non-durable goods and services expenditures. The first panel of figure B.2 shows the NIPA log ratio of real food at home to the rest of real non-durable goods and services;¹¹ the second panel plots the ratio using the uncorrected CEX data, which is aggregated by taking the sum across households of the log expenditure variables. To address this phenomenon in the regression correction, we include as explanatory variables a time trend and a time trend interacted with our explanatory expenditure variable. We also experimented with including in the regression polynomials in the time trends and explanatory expenditure variable; this made little difference to the correction.

The dummy variable for 1980-81 took on a value of -0.044, while the dummy for 1982-1987 took on a value of -0.178. Table B.3 shows the raw and corrected real food consumed at home data, both in logs and in log ratio form. The corrections seem reasonable.

The study makes use of the income data from the CEX, in tables 4 and 5. These data come from the interview survey FMLY files and the interview survey MEMB files. The

¹¹The composition of NIPA non-durable goods and services is matched as closely as possible to the CEX composition described above; the price indices used to deflate the data NIPA data are the CPI deflators described above.

CEX asks households questions about their income from the previous twelve months at the second interview (the first for which data is available) and fifth interview; in this paper, only income data from the fifth interview is used. The income data is extrapolated to a monthly frequency by assigning the fifth interview data to the month of the interview and the previous twelve months, and is then merged with demographics data from the FMLY files extrapolated to monthly frequency as before.

Following the CPS construction, household income is constructed as the earned income of its members, plus other household income. The three components of earned income of household members are wage and salary income (SALARYX), self-employed income (NONFARMX), and farm income (FARMINCX), and are drawn from the CEX MEMB files. Before aggregating across household members, top-coding corrections were made similar to those made to the CPS data. Before tax household income is then constructed as the FMLY file variable FINCBTAX (the sum of earned income, asset income, and various forms of transfer and retirement income) minus earned income without top-coding adjustments plus earned income with top-coding adjustments.

Although the CEX asks questions on taxes paid by households, this paper follows the CPS construction and uses the NBER's TAXSIM program to estimate taxes paid for each household. The primary difference from the CPS computations is that data on medical expenses, charitable contributions, rent paid, and state, local and property taxes paid, are all available for CEX households, allowing us to account for the impact of deductions on the total tax burden. The fraction of CEX households who itemized in our computations approximately matched the fraction in the population (private correspondence with Daniel Feenberg). After-tax income for CEX households was computed as the adjusted FINCBTAX variable minus the TAXSIM estimate of taxes paid.

The primary sample of CEX households we consider in this paper excludes:

1. households without a male head aged 23 to 59 (without a male householder or a male spouse of householder of that age),
2. households living in rural areas,¹²
3. households providing an incomplete income response,¹³

¹²This selection makes the CEX samples comparable over time, since non-urban consumers were not sampled in 1982-3.

¹³The CEX considers data given by incomplete reporters to be of low quality.

4. households whose reported head either ages by more than one over the sample period of 4 quarters, changes sex, or changes education,¹⁴
5. monthly household observations with non-positive non-durables and services consumption, and
6. households residing in student housing.

Synthetic cohorts are constructed in the same manner as with the CPS data.¹⁵ The four education classifications are defined as: (1) less than high school graduate, (2) high school graduate but no subsequent schooling, (3) some college, and (4) 4-year college degree or more schooling.

¹⁴These restrictions generally are meant to eliminate households whose reference person changes due to death or other circumstances, while the last restriction serves to eliminate households engaged in schooling.

¹⁵If the age of the head changes between quarterly interviews, we can usually pin down the birth year, but if not, we can often only pin down a set of possible birth months that may span two years. A reference person is assigned to birth year t rather than birth year $t - 1$ if more than half the possible birth months of the head are in year t . For consistency with the CPS, birth years are defined to run from March to February.

Figure B.1: CEX Food at Home Expenditures, Raw Data




Figure B.2: Log (Real Food at Home/Real Non-Durable Goods and Services)






Figure B.3: Corrected and Uncorrected CEX Real Food at Home Data

Log Expenditures



Log Ratio to Real Non-Durable Goods and Services

