



Current Consumption and Future Income Growth: Synthetic Panel Evidence

Jeremy J. Nalewaik
WP2004-04
December 20, 2004

The views expressed in this paper are solely those of the author and not necessarily those of the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce.

Current Consumption and Future Income Growth: Synthetic Panel Evidence

Jeremy J. Nalewaik*

December 20, 2004

Abstract

Using group means computed from twenty years of high quality survey data, I show a strong and robust relation between households' consumption growth and subsequent realizations of their income growth, including realizations as distant as six years later. The relation appears in multiple types of variation in income growth: variation across cohort-education groups, variation over the life cycle, and even some variation over the business cycle. The results may be evidence of forward-looking households altering their current consumption in response to information they receive about their future income; other interpretations are explored as well.

JEL classification: E21, D12, D91.

Keywords: Consumption (Empirical Analysis), Intertemporal Choice.

*Bureau of Economic Analysis, 1441 L Street NW, Washington, DC 20230 (e-mail: jeremy.nalewaik@bea.gov). The first version of this paper was written in January 2002. I developed this work with the help of insights from Orazio Attanasio, John Cochrane, Chris Carroll, Steve Davis, Lars Hansen, Erik Hurst, Anil Kashyap, Evan Keonig, Dean Maki, Lior Menzly, Greg Siourounis, Paul Willen, and seminar participants at: the Board of Governors of the Federal Reserve, the University of Chicago Graduate School of Business, the Federal Reserve Bank of Boston, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Dallas, Indiana University, the University of California Santa Barbara, the University of California San Diego, the Social Security Administration, the Bureau of Economic Analysis, the Second Annual Trans-Atlantic Doctoral Conference at the London Business School, and the 2002 NBER-CRIW Consumer Expenditure Survey Research Conference. The views expressed in this paper are solely those of the author and not necessarily those of the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce.

1 Introduction

To what extent do households make their consumption-savings decisions in a forward-looking manner? While this question is of fundamental importance to many aspects of economics, a wide range of viewpoints prevail. The massive empirical literature studying consumer behavior with Euler equations has produced mixed results (see Martin Browning and Annamaria Lusardi (1996) for a survey), with a rough consensus emerging that consumption growth responds to some types of predictable variation in income, in violation of the orthogonality tests characteristic of basic life-cycle/permanent-income (LC/PIH) models. Indeed, it has been shown that a number of empirical facts do not accord well with the implications of basic forward-looking models. But despite these misgivings, other evidence has been gathered indicating that the forward-looking model may have more than a grain of truth to it. For example, evidence dating back to Milton Friedman (1957) indicates that consumption responds more strongly to permanent income innovations than to transitory ones.¹

Further potential support for the forward-looking model is provided by a number of papers by John Y. Campbell and others, which document a relation in macroeconomic data between consumption, usually scaled by income, and income changes a quarter later.² More recently, Lettau and Ludvigson (2001, 2002a) have shown that a consumption measure scaled by a combination of assets and income (which they call *cay*) is related to investment returns from a quarter to several years later. Such regressions of income or asset returns on prior scaled consumption variables may be evidence that households receive information about their future income or investment opportunities

¹More recent evidence includes Christina H. Paxson (1992).

²See John Y. Campbell (1987); Campbell and Angus Deaton (1989); Campbell and N. Gregory Mankiw (1989); C.L.F. Attfield, David Demery and Nigel W. Duck (1990); and John H. Cochrane (1994). Antecedent or concurrent work along the same lines occurred in Thomas J. Sargent (1978); Marjorie Flavin (1981); and Lars P. Hansen, William Roberds, and Sargent (1991).

and alter their consumption in response, as the forward-looking model predicts they should. Under this interpretation, these regressions essentially tell the econometrician what households' know about their future income or asset returns, which is nice since households' information about these variables, especially their own future income, is generally vastly superior to what an econometrician can discern from observables other than consumption. So, compared to other strategies where the econometrician characterizes the household income process and tests whether consumption behaves as it should, these regressions bring more information to bear in characterizing households' forward-looking nature. Any characterization of the household income process by the econometrician will always be limited by the his or her potentially much smaller information set, and in fact the econometrician's characterizations must coincide with households' perceptions of the income process for these tests to be valid (see Hansen et al., 1991). Similarly, Euler equation orthogonality tests, while robust to households' potentially superior information, do not exploit it as consumption forecasting regressions do, as the universe of testing variables will always be limited by the information set of the econometrician.

While this macroeconomic evidence for forward-looking behavior is intriguing and potentially very valuable, corroborating evidence from micro data would take it to a new level of credibility. This is because in the macro data, plausible non-forward-looking interpretations are readily available; as Deaton (1992) points out:

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. (Angus Deaton, 1992b, p. 133.)

The microeconomic data employed in this paper permit exploration of these competing

explanations, as well as providing more types of usable variation than the macro data. While prior microeconomic studies of the relation between current consumption and future income have produced no convincing evidence for such a relation,³ the work here employs a long, high quality synthetic panel to arrive at a different result. In addition, I employ an econometric specification that I believe to be superior to what has been used in prior empirical work, which generally uses consumption scaled by other variables as a predictor. It is always hard to know whether the predictive power uncovered in these specifications really stems from consumption, rather than the variables used to scale consumption; see the recent exchange between Lettau and Ludvigson (2002b) and Brennan and Xia (2002). The specification I employ here does not suffer from these problems of interpretation; in addition, it allows easier decompositions of the predictive power of consumption by time horizon.

Although the main contribution of this paper is empirical, I begin in section 2 in the usual way, with a theory section designed to aid interpretation of the empirical results, through the lens of a forward-looking model of consumer behavior. Section 3 discusses the twenty years of micro data used in the paper (from the Consumer Expenditure Surveys (CEX) and Current Population Surveys (CPS)), and computation of the synthetic panel - 28 time series of group means of income and consumption, with households grouped according to the educational attainment and birth cohort of their male heads. Section 4 shows the main empirical result: a statistically significant relation between growth rates of group means of income (computed from the CPS) and growth rates of group means of consumption (computed from the CEX) from one to five years earlier. Section 5 undertakes an extensive robustness analysis, breaking down the relation between consumption and future income by different types of variation and different types of households; the main results are found to be remarkably robust. Sec-

³See Deaton (1992a), Christopher D. Carroll (1994), Deaton (1997), and Rob Alessie and Lusardi (1997).

tion 6 discusses explanations of the results other than the forward-looking explanation; section 7 concludes.

2 Theory

The baseline empirical specification in this paper is a synthetic panel estimate of:

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

where Δy_t^i denotes time t income growth for group i , estimated from CPS sample means, and where Δc_t^i denotes time t consumption growth for group i , estimated from CEX sample means. We provide a derivation of this equation in the context of the certainty equivalent LC/PIH, and briefly discuss how interpretations change when we relax some of the LC/PIH assumptions.

2.1 Hueristic Derivation

Taking the simplest possible case, consider an infinitely-lived household h whose information about its income derives entirely from past realizations of the income process - i.e. the household information set about its income is univariate. The growth rate of household income Δy_{t+1}^h ⁴ is assumed covariance stationary; no other assumptions are required to write Δy_{t+1}^h in its Wold moving average representation:

$$(2) \quad \begin{aligned} \Delta y_{t+1}^h &= \kappa_{t+1}^h + \rho_0^h \varepsilon_{t+1} + \rho_1^h \varepsilon_t + \rho_2^h \varepsilon_{t-1} + \dots \\ &= \kappa_{t+1}^h + \rho^h(L) \varepsilon_{t+1}. \end{aligned}$$

⁴Appendix B directs us to include dividend and interest income in this measure.

The ε_{t+1-k} are realizations of an independent and identically distributed random variable representing innovations to the household's information about Δy_{t+1}^h ; $\rho^h(L)$ is a polynomial in the lag operator (assumed of order q); κ_{t+1}^h is the linearly deterministic component of the income process.

Household h 's consumption growth Δc_{t+1}^h follows a version of the standard LC/PIH: Δc_{t+1}^h is unpredictable and equals the time $t+1$ innovation to the present discounted value of h 's current and expected future income growth. Taking a constant interest rate r , and letting $\lambda = \frac{1}{1+r}$, consumption growth is then:

$$\begin{aligned}
\Delta c_{t+1}^h &= \sum_{j=1}^{\infty} \lambda^{j-1} (E_{t+1} - E_t) \Delta y_{t+j}^h \\
&= (\rho_0^h + \rho_1^h \lambda + \rho_2^h \lambda^2 + \dots) \varepsilon_{t+1} \\
(3) \qquad &= \rho^h(\lambda) \varepsilon_{t+1},
\end{aligned}$$

where E_t is an expectation taken with respect to the time t information set of the household. Appendix A works through one possible derivation such a LC/PIH consumption function, clarifying its implicit assumptions. Since $\varepsilon_{t+1} = \frac{1}{\rho^h(\lambda)} \Delta c_{t+1}^h$, the consumption growth rates can be used to substitute the innovations out of (2):

$$\begin{aligned}
\Delta y_{t+1}^h &= \rho_0^h \varepsilon_{t+1} + \rho_1^h \varepsilon_t + \dots + \kappa_{t+1}^h \\
&= \frac{\rho_0^h}{\rho^h(\lambda)} \Delta c_{t+1}^h + \frac{\rho_1^h}{\rho^h(\lambda)} \Delta c_t^h + \dots + \kappa_{t+1}^h \\
&= \frac{\rho^h(L)}{\rho^h(\lambda)} \Delta c_{t+1}^h + \kappa_{t+1}^h.
\end{aligned}$$

This is equation (1) at the household level rather than the group level, although in (1) we have truncated the potentially infinite number of lags of consumption growth at k lags. The model tells us to interpret $\beta(L) = \frac{\rho^h(L)}{\rho^h(\lambda)}$; so the β s are the Wold moving average coefficients governing the household income process, normalized so their discounted sum

is unity - i.e. $\sum_{k=0}^q \lambda^k \beta_k = \beta(\lambda) = 1$. A statistically significant β_k translates to a statistically significant ρ_k^h , implying that some information is revealed to households about their income growth k periods in advance of its arrival.

2.2 Superior Household Information

The restrictive assumption in the previous subsection on the information set of the household is both unrealistic and unnecessary to derive (1), as Hansen, Roberds and Sargent (HRS, 1991) show. They consider an information set for household h that consists of n linearly independent,⁵ covariance-stationary variables, including Δy_{t+1}^h ; n may be arbitrarily large, so the information set of the household may be arbitrarily larger than the information set of the econometrician. The vector process can be written in a Wold moving average representation, with one row characterizing income growth again as in (2), but with ε representing an n -dimensional column vector of (white noise) innovations to the information set of the household,⁶ and $\rho^h(L)$ representing an n -dimensional row vector of polynomials of order q in the lag operator. Given the LC-PIH consumption function, we again have: $\Delta c_{t+1}^h = \rho^h(\lambda) \varepsilon_{t+1}$.⁷

Appendix B shows the HRS derivation of (1) in this multi-dimensional setting; in addition to showing the general derivation, it works through an example for a specific

⁵This “assumption” of n linearly independent variables in the household information set, to quote a previous referee who objected, is without loss of generality.

⁶Without loss of generality, we assume the variance-covariance matrix of ε is the identity matrix I_n .

⁷This consumption function assumes that households observe different types of shocks separately. While this assumption produces the cleanest derivation of (1), other models may be considered as well in interpreting the paper’s empirical results. As an example, take the case where $n = 2$, with one shock being household-specific, and the other economy-wide or specific to the household’s entire education-cohort group. The household observes only an average of these two shocks, however; it cannot disentangle one from the other. So the household may receive a shock indicating that its income will increase over the next few years, not knowing whether the shock will hit the economy or its cohort-education group as a whole, or just the household itself. Combined with the aggregation scheme employed in this paper, such a model may make the paper’s empirical results more palatable to some economists; see the next footnote. For analysis of a model of this kind, see also Jorn-Steffen Pischke (1992).

time series process where $n = 2$. In the derivations, we see that consumption growth reveals to the econometrician a linear combination of the n elements of ε_{t+1} , which can be substituted out of the reinterpeted (2); the remaining $n - 1$ elements of household information end up in the error term of the regression, and are necessarily orthogonal to consumption growth. The linear combination of ε_{t+1} revealed to us by consumption growth is $\frac{\rho^h(\lambda)}{|\rho^h(\lambda)|}\varepsilon_{t+1}$, where $|\rho^h(\lambda)| = \sqrt{\rho^h(\lambda)\rho^h(\lambda)'}$. $\frac{\rho^h(\lambda)}{|\rho^h(\lambda)|}\varepsilon_{t+1}$ weights more heavily those elements of ε_{t+1} with larger discounted sums of polynomial coefficients; as posited by Friedman (1957), consumption reflects more strongly those innovations that have a more persistent effect on the income process. The implication that $\beta(\lambda) = 1$ continues to hold in the multivariate setting - this is the HRS test of present value budget balance.

2.3 Aggregation

Aggregation of this model is easy. Let ε_{t+1} represent the vector of innovations to a new information set equal to the union of the information sets of all households in the economy. Then we can write equation (2) for an individual household using the expanded ε_{t+1} and simply setting to zero those rows of $\rho^h(L)$ corresponding to elements of ε_{t+1} not in the individual household's information set. Averaging into groups, the main results of sections 2.1-2.2 carry over with $\rho^h(L)$ replaced by $\bar{\rho}^i(L)$, the average moving average polynomials for households in group i . This produces equation (1) at the group level, with $\beta(L) = \frac{\bar{\rho}^i(L)}{\rho^s(\lambda)}$. More on aggregation can be found in my dissertation, Nalewaik (2003).⁸

⁸Again, this model where households observe different types of shocks separately produces the cleanest aggregation, but other models of household information may be more reasonable. Continuing with the example from the previous footnote, a household observes a shock indicating its income will increase over the next few years, but does not know whether the shock will hit the economy as a whole or just itself; in any event, it increases its consumption in response. It may know that most of the shocks it receives are idiosyncratic (so the variance of the idiosyncratic component of the shocks it observes is much larger than the variance of the aggregate component), in which case the household will place little stock in this shock as a predictor for the economy as a whole, and will have little to say about the future state of the economy if asked. Its *individual* consumption decisions are uninformative about

2.4 Extensions of the LC/PIH

Much empirical evidence rejects the strictest version of the LC/PIH where consumption growth is completely unpredictable, and there are sound theoretical reasons to expect some predictability, including precautionary savings motives, liquidity constraints, and non-separabilities between consumption and other determinants of utility such as leisure and demographics variables. This is one rationale for including control variables in (1): appropriately chosen controls may increase the ratio of news to predictable variation in consumption growth, bringing the data more in line with the LC/PIH theory in sections 2.1-2.3. In particular, some empirical specifications below employ controls to remove from the data its mean cross-sectional and life-cycle variation, likely the types of variation most predictable to households. However the appropriateness of these predictability assumptions is not always clear cut, and will be discussed more later.

Whatever control variables the econometrician includes in the regression, consumption growth is unlikely to conform exactly to the predictions of the LC/PIH. Liquidity constraints, precautionary savings motive, and complementarity between consumption and leisure may skew the β s away from their LC/PIH interpretation as scaled moving average coefficients; in my dissertation, Nalewaik (2003), I show that these factors will dampen the response of consumption growth to variation in expected future income growth, so consumption responses will be more muted than the one-for-one relation predicted by the LC/PIH. At the macro level, decreasing returns to capital may also

the aggregate economy as well, their movements being dominated by idiosyncratic variation. However when a shock really does hit the entire economy or group, all households in that group will be ratcheting up their consumption simultaneously, with each one thinking the shock is specific to itself. As econometricians, when we average household consumption into groups, this is exactly the type of variation we will isolate, as we will average away much of the idiosyncratic variation in consumption, isolating and cumulating the small pieces of information about aggregate and group-level income in individual households' consumption decisions. The aggregation is key to interpreting the paper's empirical results in this story: with its increased signal-to-noise ratio, the semi-aggregated consumption growth may be a useful predictor of the state of the business cycle several years in the future, or the fate of demographic groups, even if no individual household has much of a clue about these things.

dampen the response of consumption to expected future income; elastic labor supply and saving for retirement could dampen consumption’s response to both current and expected future income. Interestingly, this dampening of consumption responses will tend to *increase* the size of the β s in (1), as small consumption changes signal larger, more than proportionate changes in current or expected future income.⁹

Since some of these extensions to the LC/PIH surely play an important role in consumer behavior, the LC/PIH implication of $\beta(\lambda) = 1$ probably should not be taken too seriously; the extensions to the theory broadly predict $\beta(\lambda) > 1$. In addition, the interpretation of the β_k as scaled moving average coefficients will remain a matter of judgement. We could make powerful statements if such an interpretation held exactly: for example a result of $(\beta_k)^2 > (\beta_j)^2$ says households learn more about their income growth k periods ahead of its arrival than j periods ahead. While such statements should be taken with a grain of salt, my personal view is that the intuition from forward-looking models distilled into the LC-PIH, and its accompanying interpretation of the β s, remains a very useful place to start in thinking about regression results from (1).

3 Data and Econometrics

The Annual Demographic Files from the CPS provide the paper’s primary source of data on household income from 1980-1999; the CEX Interview Surveys provide data on household consumption over the same time period. Both sets of surveys are short rotating panels, with CPS sample sizes (50,000-60,000 households) roughly an order of magni-

⁹Other modifications to the LC/PIH will have different effects. Campbell and Mankiw (1989) consider models with “rule-of-thumb” consumers, who set $\Delta c_{t+1}^h = \Delta y_{t+1}^h$ each period. For these consumers, estimates of (1) would yield $\beta_0 = 1$, with all other $\beta_k = 0$. If some fraction of consumers behaved as “rule-of-thumb” consumers and the rest behaved as LC/PIH consumers, our estimated β s would be a weighted average of the moving average coefficients of the LC/PIH consumers, and the “rule-of-thumb” β s ($\beta_0 = 1$ and $\beta_k = 0, k \neq 0$). In other words, we should observe a spike at β_0 .

tude larger than CEX sample sizes. Household income includes earned (labor) income,¹⁰ transfer income, and asset income,¹¹ and subtracts an estimate of taxes paid computed using the National Bureau of Economic Research (NBER)'s TAXSIM program. The CEX is the most comprehensive source of micro data available on expenditures of US households, reporting several hundred expenditure categories; the consumption measure used here is non-durable goods and services and excludes durables, medical care, education, and housing.¹² The CEX data on food consumed at home was corrected for discontinuities introduced by changes in survey design in 1982 and 1987. Both income and consumption were deflated by the annual personal consumption expenditures (PCE) deflator from the National Income and Product Accounts (NIPA). For further details on the raw consumption and income data used in the paper, see the data appendix available from the author.

The sample of households studied here, households with a male head aged 23 to 59, includes about half of all households in the survey data.¹³ The CEX and CPS samples in each year are partitioned by the characteristics of their male household heads: four categories for their educational attainment (high school drop-outs, high school graduates with no other schooling, high school graduates with additional schooling less than a 4-year college degree, and 4-year college graduates), crossed with seven 5-year birth cohorts, ranging from the 1931-1935 birth cohort (i.e. male heads born in those years)

¹⁰Some top-coding adjustments are made to the earned income variables following Lawrence Katz and Kevin M. Murphy (1992).

¹¹Empirical results that excluded asset income from the household income measure were virtually identical to the results reported in this paper.

¹²This consumption measure matches that of Orazio P. Attanasio and Steven J. Davis (1996), whose programs I used to create it. I use non-durables consumption to avoid additional dynamics that may be introduced into the basic LC/PIH by durability, and because durable goods expenditures are partially savings, not the pure consumption that is the topic of the theory.

¹³The CEX sample also excludes rural households and households classified as incomplete income reporters; these are fairly standard sample selection restrictions - see Attanasio and Davis (1996).

to the 1961-1965 birth cohort.¹⁴ The synthetic panel on consumption and income is the set of means of log consumption and log income taken over the sample of households in each group in each year. For each group there is a time series of up to 20 annual cell means, but due to the sample selection restrictions on age the panel is unbalanced; we have 524 cell means in total for the 28 groups.

Stacking together the 28 time series, our baseline specification is a synthetic panel estimate of (1). Unfortunately, random differences from sample to sample introduce sampling errors into the panel - i.e. errors due to the measured sample means being different from their corresponding population means from year to year. The variance of the sampling error for each synthetic panel cell mean, on either income or consumption, is inversely proportional to the number of households used to compute it (its cell count for short); table 1 reports summary statistics on these 524 cell counts. They vary widely from observation to observation, leading heteroskedasticity to be a source of concern. As a remedy, the paper's least squares estimates of panel equations such as (1) weight by the average cell count of the cells used to produce the explanatory variables, effectively downweighting the thinner synthetic panel observations more contaminated with sampling error. Note that the sampling errors in Δy_t^i will be independent of the sampling errors in Δc_t^i , as the CPS and CEX are independent random samples; then the sampling errors in Δc_t^i should bias the β s in (1) generally towards zero, making it more difficult to identify a relation between consumption and income growth.

Table 2 reports summary statistics for the panel. Autocorrelations are computed using weighted least squares, controlling for a full set of fixed effects and education-specific quartic polynomials in the age of the male household heads.¹⁵ The controls

¹⁴Nine 5-year birth cohorts met the sample selection restrictions on age, but our specification also requires computation of several lags of growth rates, which caused the two cohorts with the shortest time series (the 1926-1930 cohort and the 1966-1970 cohort) to drop from the sample.

¹⁵A five year birth cohort is assigned the age of its middle-aged members.

purge the data of most of its variation over the life cycle and its mean cross-sectional variation, leaving mainly the time series variation in the panel. While we do see a large negative first order autocorrelation for consumption growth, this is the expected outcome of substantial sampling variability; otherwise there is little autocorrelation in either variable's time series variation. To compute the standard errors here and throughout the paper, I use an asymptotic variance-covariance matrix (described in Appendix C) that is robust to heteroskedasticity, contemporaneous cross correlation, and seventh order autocorrelation.

Before proceeding to the main result, we should note that the model in section 2.1-2.3 is particularly amenable to estimation on panel data. By construction, the errors are orthogonal to the regressors; see Appendix B. Cohort effects biasing the β s is a non-issue, since these effects (cohort averages of κ_{t+1}^h) are linearly deterministic and hence orthogonal to the innovations captured by consumption growth. Furthermore, consumption growth is white noise, so no complications arise regarding estimation of distributed lags on short panels (see Ariel Pakes and Zvi Griliches (1984)). The only real complication arises in the case where the β s are heterogeneous across cohorts; for the panel estimates to consistently estimate the mean β s across (appropriately weighted) cohorts, the heterogeneity in the β s must be independent of consumption growth (again see Pakes and Griliches (1984)). However the results in section 5.4 will indicate little heterogeneity in the β s: the relation between consumption and future income appears fairly uniformly distributed across households of different ages and education levels.

4 The Basic Result

Panel A in table 3 shows weighted least squares estimates of (1) with various cutoffs q , with no additional control variables beyond a constant. We see a statistically significant relation between consumption growth and income growth as far as five years into the

future, with the size of the regression coefficients peaking at the two to five year horizon. The results at longer horizons depend very little on whether or not we include contemporaneous consumption growth and the first lag, as the last line of the main panel shows. If consumers behave as in the LC/PIH, table 3 indicates that households receive more information about their income growth five years before it arrives than at the date of its actual arrival.

For comparison, panel B in table 3 shows univariate regressions, regressing income on each lag of consumption growth separately. These regressions hold the sample fixed for comparability, each utilizing only the 268 observations for which we can compute the sixth lag of consumption growth.¹⁶ The R^2 here may seem small, but it should be kept in mind that sampling variability in income and consumption growth is reducing R^2 ; its maximum attainable value is considerably less than unity.

For an additional comparison, table 4 reverses the roles of Δc_t and Δy_t in equation (1). While I give these results no structural interpretation, I simply note that these regressions are similar to what has been done in the Euler equation literature. While we find a statistically significant relation between Δc_t and Δy_{t-1} , a relation similar to what has shown up as a violation of the orthogonality conditions in many Euler equation estimates, we also find very little evidence of a relation between consumption growth and further lags of income growth. The coefficients here are sometimes larger than the coefficients found in table 3, but this is to be expected: the sampling error variance of the explanatory variables is about ten times smaller here than in table 3, due to the much larger CPS sample sizes. Comparing t -statistics is probably more informative, as may be a comparison (as we change q) of patterns in R^2 between the two tables. The incremental additional explanatory power of the lags of consumption growth for income

¹⁶One thing to notice here is that restricting the sample to these 268 observations reduces the contemporaneous correlation between consumption growth and income growth, reflecting a weakening over time of the degree to which consumption tracks income contemporaneously (the restricted sub-sample starts in 1986-7, as opposed to 1980-1 in the full sample).

is remarkable given the total lack of incremental additional explanatory power of the lags of income growth for consumption.

The bottom line from tables 3 and 4 is that consumption growth predicts income growth far into the future, consistent with households actively changing consumption in response to forecasts of their income growth several years ahead, forecasts that turn out to be accurate. And the reverse is not true: consumption seems to fully incorporate the information in realizations of income growth after one year. While the LC/PIH predicts that consumption should incorporate all available information about income growth *immediately*, if we allow a one year lag the theory does not perform so badly.

5 Robustness Analysis

5.1 Types of Variation: Cross Sectional, Life-Cycle, and Time-Series

What types of variation drive the results in table 3? Consider first the variation in mean growth rates across our 28 cohort-education groups. If we orthogonalize Δy and contemporaneous and lagged Δc with respect to a set of year effects (removing mean business cycle effects from the data) and then take group means, the group means of the sum of Δc_{t-4} to Δc_{t-7} (the longer lags of consumption growth) explain about 55 percent of the variation in group means of Δy_t , while group means of the sum of Δc_t to Δc_{t-3} (the shorter lags) explain only 49 percent. In a simple cross-sectional regression of the 28 group means of Δy on both sets of group means of Δc , only the average of the fourth to seventh lags is significant. Such evidence suggests that the relation between consumption and future income derives at least in part from cross sectional variation in mean growth rates, arising from such sources as the continuing expansion of income inequality between more and less educated earners in this sample period.

Next consider variation over the life cycle. The highly influential work of Carroll and Lawrence H. Summers (1991) uses CEX data to show what is essentially a contemporaneous “one-for-one” relationship between income and consumption, as they put it. The data used here tells a different story, primarily because it employs multiple repeated cross sections to track fixed cohorts over time, unlike Carroll and Summers who confound cohort and age effects in their life-cycle profiles by using a single cross section of data to produce each plot.¹⁷ Figure 1 shows, for each of our four education categories, predicted values from a synthetic panel regression of CEX consumption and CPS income on a fifth-order polynomial in the age of the male household head, controlling for a full set of cohort effects. For every education category, the peak in the age-consumption profile clearly precedes the peak in the age-income profile, indicating that the relation in table 3 may stem in part from life-cycle variation.¹⁸

We next use control variables to purge our income and consumption data of its mean cross sectional variation and its variation over the life cycle. While the strictest version of the LC/PIH in sections 2.1-2.3 provides no justification for control variables, there are at least three good reasons to include them. First, they help provide a general decomposition of the results in table 3 by different types of variation. Second, given the persistent rise in the returns to education over our sample period, some readers may question the stationarity of some components of the Δy (even in growth rates). Such concerns argue for the inclusion of group fixed effects in (1), as these controls will remove

¹⁷Part of the variation in Carroll and Summers’ plots stems from age effects, but part of it stems from cohort effects as well, since the groups they differentiate by 5-year age bands are also differentiated by 5-year cohorts. Being time invariant, cohort effects in income will translate one-for-one into cohort effects in consumption even in models where households are completely forward-looking; contemporaneous “tracking” of consumption by income is no evidence against forward-looking models when it appears in variation driven by cohort effects.

¹⁸Various other papers use synthetic panel data to plot income and consumption over the life cycle, including Attanasio and Weber (1995), Attanasio and Browning (1995), and Banks, Blundell, and Tanner (1998). Pierre-Olivier Gourinchas and Jonathan A. Parker (2002) use techniques similar to those employed here and find similar results, showing plots where consumption clearly leads income over the life-cycle for various cohort-occupation and cohort-education groups.

from the regression much of this potentially non-stationary across-group variation. And third, as discussed in section 2.4, controls may purge consumption growth of much of its predictable variation, leaving variation in consumption growth that is mostly news, as the LC/PIH posits. However the appropriateness of purging the regression of a particular type of variation will depend on assumptions about its predictability, which in turn may depend on assumptions about the information set of households. Consider again the differences in mean consumption growth across cohort-education groups. Under one set of assumptions, these differences represent differences in the average values of innovations to households' information about their income growth over the sample period: since mean Δc for households headed by college graduates exceeded mean Δc for households headed by high-school graduates, the college graduates on average received better news about their current and expected future income in the 1980's and 1990's. The fact that the wage gap between college and high school graduates grew more or less continually over this time period does not mean that households knew this would occur ex-ante; for example it may be that only midway through the 1980's did households conclude that this gap would likely continue growing throughout the decade and into the next. Under this interpretation, group means are simply another source of variation to exploit in estimating (1), and there is no need to control for group fixed effects. However under other sets of assumptions, across group differences in mean consumption growth are driven by predictable variation rather than news, so any correlation with group means of income growth would not be the result of forward-looking behavior on the part of households. Inclusion of fixed effects in (1) would then be appropriate.

To compute estimates of (1) purging the data of its mean cross sectional variation, I simply included a set of group fixed effects. To purge the data of its variation over the life cycle, I tried various sets of controls, including many measures of family size and composition (numbers of children of different ages, numbers of elderly, whether a wife was present, etc.), and various polynomials in the age of the male household head. However

once I included quartic polynomials in the age of the male household head interacted with education fixed effects, all the other demographics controls had virtually no effect on the explanatory power of consumption growth or the coefficients in (1); for this reason I report results that include only these education-specific age polynomials as controls for variation over the life cycle. Estimates of (1) that include these controls are reported in the second specification of table 5. This table reports the regression β s, as well as the discounted sum of the regression coefficients $\beta(\lambda) = \sum_{k=0}^q \lambda^k \beta_k$, where $\lambda = \frac{1}{1+r}$ and r is the interest rate set to 0.025; the HRS test of present value budget balance predicts that this quantity equal one. Table 5 also reports implied r , the estimated value of r that sets $\beta(\lambda) = 1$,¹⁹ and $R_{\Delta c}^2$, the regression R^2 after orthogonalizing the data with respect to the control variables.

Remarkably, the primary effect of including in (1) the fixed effects and age polynomials is to increase the size of the β s; just as remarkably, the controls do very little to change the relative magnitudes of the coefficients, as the peak in the β s remains at the two to five year horizon. The discounted sum of the regression coefficients $\sum_{k=0}^q \lambda^k \beta_k$ increases by about 25-30 percent compared to the first specification with no controls. While $R_{\Delta C}^2$ drops substantially and standard errors increase, the β s are plainly statistically significant. Given the nature of the controls, the variation that is left in the data is largely time series variation; the results indicate that consumption growth predicts some time series variation in income growth years in advance.²⁰ For some graphical evidence, see Nalewaik (2003), who plots the aggregate business cycle variation in the data and deviations of groups' business cycle experiences from the aggregate.

¹⁹The standard error of this quantity is computed using the delta method and numerical derivatives.

²⁰I reiterate that this does not *necessarily* imply that, on average, individual households are able to forecast the state of the business cycle several years into the future. See footnotes 7 and 8.

5.2 Aggregate vs. Non-Aggregate Variation

The next two estimates in table 5 address concerns that a single aggregate effect may be influential, by including as controls a set of year fixed effects (the fourth estimate includes the age polynomials and cohort effects as well). The year effects have little impact on $R_{\Delta C}^2$, and the majority of the coefficients remain statistically significant.²¹ The HRS test $\beta(\lambda) = 1$ fails on the low side in these specifications, but it should be kept in mind that sampling errors are biasing these β s towards zero, more in these specifications than in others, as the control variables remove from consumption growth much of its true variation and little or none of its noise from sampling errors.

Interestingly, the year effects increase the β s on the longer lags of consumption growth at the expense of the β s on contemporaneous consumption growth and the shorter lags. The β s now die out at the seven year horizon, as the relation between consumption growth and income growth six years ahead is now relatively large and statistically reliable. The fourth estimate in table 5 indicates that a typical group's consumption growth contains information about how its business cycle experience will differ from the aggregate business cycle, at horizons four, five and six years in the future.

5.3 Consumption as a Proxy for Income

The final two specifications of table 5 include seven lags of CPS income growth as control variables. Prior work on current consumption and future income has often worked with variables like savings or consumption-income ratios, essentially scaling the explanatory variable consumption by income; some have argued that the empirical results in those papers are driven by variation in the scaling variable rather than variation in consumption as the basic forward-looking theory predicts. While such criticisms are a non-issue

²¹Additional analysis confirmed that one influential year does not drive the results; dropping any individual year (any of 1987-1999) does not appreciably impact the β s.

for an econometric specification such as (1) where consumption is left to its own devices to explain future income, concerns may remain that consumption is simply proxying for income in (1). The final estimates in table 5 address these concerns, and given the large *increases* in β s and $R^2_{\Delta C}$ when the lags of income growth are included in the regression, the argument that consumption is simply proxying for income has little to no credibility. The argument loses even more credibility when one realizes that the largest source of measurement error in these data are sampling errors, and that the variance of the sampling errors in consumption growth is about an order of magnitude larger than the variance of the sampling errors in income growth. The facts from this table are clear: lags of consumption growth have substantial additional explanatory power for income growth above and beyond lags of income growth.

5.4 Heterogeneity by Age and Education

Table 6 explores heterogeneity in the β s by age, reporting six sets of regression coefficients from three specifications of (1). Two sets of consumption growth terms are included in each specification: one interacted with dummies for whether the cohort was younger than a specified age cutoff (at the time of the consumption growth), and another interacted with dummies for whether the cohort was older. The age cutoffs reported are 33, 38 and 43; the first and fourth set of reported estimates in table 6 consider an age cutoff of 33, for example. Controls are quartic age polynomials and birth cohort fixed effects interacted with education fixed effects.

The results in table 6 show some relatively large β s for the age 23-33 group. However the standard errors are large, and examination of other age cutoffs indicates that the drop in $\sum_{k=0}^q \lambda^k \beta_k$ occurs quite suddenly as the cutoff increases from 34 to 35. There is some evidence of an increase in $\sum_{k=0}^q \lambda^k \beta_k$ for older cohorts as the age cutoff increases, but the increase is neither large nor statistically significant. Overall the case for heterogeneity

in the β s over the life cycle is not strong.

Table 7 reports regression estimates for each of the four educational classifications of the male household head; estimation for each classification is done separately employing cohort specific intercepts and a quartic polynomial in age as controls. The poor results for high-school drop-outs stand out most in this table, but the lack of a statistically significant relation may be due to relatively small sample sizes and hence relatively large measurement error in consumption growth for this group (see table 1). The β s for college graduates appear to be more concentrated at longer lags than the β s for other groups, but again the evidence for heterogeneity is not so strong. Overall, the explanatory power of consumption growth for future income growth is fairly uniformly distributed across the groups in our synthetic panel, further evidence of the robustness of the paper's main results.

6 Non-Forward Looking Explanations

6.1 Aggregate Feedback

Table 5 shows that the relation between consumption and future income extends beyond the aggregate business cycle to relative variation, illustrating that the Keynesian feedback effects discussed by Deaton cannot be the whole story behind the empirical results. For such feedback effects to explain the relation in group-specific relative variation, the groups would need to function as at least somewhat autarkic economies, a condition which is clearly not met for cohort-education groups in an integrated modern economy. More generally, the results in table 5 cast doubt on all stories that work primarily through aggregate feedback or aggregate effects, such as correlation of consumption growth with aggregate variables that forecast output growth, including interest

rates and stockmarket returns.²²

6.2 “Gotta Pay the Bills”

One interesting explanation for the empirical results, quite different from the models discussed in sections 2.1-2.4, posits that non-forward-looking consumption changes (caused by random taste shocks, for example) force households to increase their income later to meet the budget constraint; for short, call this the “gotta pay the bills” story. The most obvious means by which households increase their income is by working longer hours, so if “gotta pay the bills” largely explains the empirical results, we may expect the relation between consumption and future income to stem largely from the work hours component of income, and not much from the wage component. To examine this issue, table 8 shows results from estimating several of the same regressions as in table 5, substituting the real after-tax wage of the male household head²³ for total household income.

The coefficients in table 8 are somewhat more more erratic than those in table 5, but broadly speaking the results are similar. For time series fluctuations, the $R_{\Delta c}^2$ indicate that lags of consumption growth pick up at least as much variation in wage growth as in income growth. For the specifications with year effects, the strong relation at longer lags clearly remains, and in fact statistically significant coefficients do not die out even at the seventh lag.²⁴ Consumption growth contains information about how groups’ male wages will differ from aggregate male wages over the business cycle, at horizons from

²²Given the highly skewed distribution of asset holdings across cohort-education groups, we cannot rule out the possibility that such stories may explain some of the across-group relation between consumption and future income. However arguing against the importance of these stories is the available evidence indicating that stock market returns and interest rates do not predict output growth more than about a year in advance, see Eugene F. Fama (1990) and Nai-Fu Chen (1991).

²³Computed as the annual labor income of the male head divided by his annual work hours, from March CPS data. Tax rates are computed from NBER’s TAXSIM program.

²⁴Additional lags of consumption growth were not added to the regression because of concerns about dropping more years from the sample.

four to seven years in the future. For more on the relation between consumption and future wage growth, as well as on the relation between wives' leisure and future wage growth, see my dissertation, Nalewaik (2003).

Table 8 indicates that the simplest version of "gotta pay the bills" can at best explain only part of the results in tables 3 and 5. Further, most economists believe that much of the variation used to estimate (1) in those tables, including the aggregate business cycle variation and variation in the returns to education, are caused by shifts in labor demand rather than shifts in labor supply, as quantities and prices covary positively. However other arguments could be made in favor of "gotta pay the bills": hours may be mis-measured, and earners may be able to increase their earnings without increasing work hours, for example by taking more difficult jobs, increasing work effort, and undertaking on-the-job training, all of which would translate into measured wages rather than work hours. It would be interesting to see if corroborating evidence could be marshalled to examine the validity of these more sophisticated "gotta pay the bills" stories. They certainly could be part of the story; it should be noted that the competing explanations of the results are not mutually exclusive. However I believe the evidence favors forward-looking behavior as the greater part of the explanation for this paper's results.

7 Conclusion

It is probably safe to say that very few economists would have predicted the main empirical results observed in this paper; that being the case, the evidence here should provide substantial stimulus to the debate about the extent to which households conform to basic forward-looking models. More importantly, the results open up potentially promising new avenues for future research; for example income could be broken down into its various components - taxes, transfers, the work hours and wages of the male and female household heads. Table 8 started this work, and Nalewaik (2003) goes further

down such a path.

More generally, the results in this paper argue for moving beyond traditional Euler equation specifications to examine other implications of basic forward-looking models, such as the implication that consumption should reflect household information about future income. While specifications designed to examine such implications generally employ more assumptions than Euler equation estimates, namely assumptions about the structure of the budget constraint, the additional uncertainty induced by these assumptions can be combat at least in part by a vigorous robustness analysis, including examination of different control variables and types of variation in income growth. Since the specification studied here was specifically designed to exploit households' superior information about their own future income, information that is left out of studies based on Euler equation orthogonality tests, we should at least consider the possibility that the added value of this additional information may swamp the added cost of incorporating some additional auxiliary assumptions into our econometric estimates. The intriguing empirical results in this paper certainly argue that this is the case.

References

- [1] Alessie, Rob and Lusardi, Annamaria, 1997. Saving and Income Smoothing: Evidence from panel data. *European Economic Review* 41, 1251-1279.
- [2] Attanasio, Orazio P. and Browning, Martin, 1995. Consumption over the Life Cycle and over the Business Cycle. *American Economic Review* 85, 1118-1137.
- [3] Attanasio, Orazio P. and Davis, Steven J., 1996. Relative Wage Movements and the Distribution of Consumption. *Journal of Political Economy* 104, 1127-1262.
- [4] Attanasio, Orazio P. and Weber, Guglielmo, 1995. Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy* 103, 1121-1157.
- [5] Attfield, CLF; Demery, David; and Duck, Nigel W., 1990. Saving and Rational Expectations: Evidence from the UK. *Economic Journal* 100, 1269-1276.
- [6] Banks, James; Blundell, Richard; and Tanner, Sarah, 1998. Is There a Retirement Savings Puzzle? *American Economic Review* 88, 769-788.
- [7] Brennan, Michael J. and Xia, Yihong, 2002. Tay's as good as cay. Working paper, the Wharton School, University of Pennsylvania.
- [8] Browning, Martin and Lusardi, Annamaria, 1996. Household Saving: Micro Theories and Micro Facts. *Journal of Economic Literature* 34, 1797-1855.
- [9] Campbell, John Y, 1987. Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis. *Econometrica* 55, 1249-1273.
- [10] Campbell, John Y, 1993. Intertemporal Asset Pricing without Consumption Data. *American Economic Review* 83, 487-512.

- [11] Campbell, John Y. and Deaton, Angus, 1989. Why is Consumption So Smooth? *Review of Economic Studies* 56, 357-374.
- [12] Campbell, John Y. and Mankiw, Gregory N., 1989. Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. In Olivier Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual* (Cambridge, NBER) 185-245.
- [13] Carroll, Chris D., 1992. The Buffer Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity* 1992(2), 61-156.
- [14] Carroll, Chris D., 1994. How Does Future Income Affect Current Consumption? *Quarterly Journal of Economics* 109, 111-148.
- [15] Carroll, Chris D., 1997. Buffer Stock Saving and the Life-Cycle Permanent Income Hypothesis. *Quarterly Journal of Economics* 112, 1-55.
- [16] Carroll, Chris D. and Summers, Lawrence H., 1989. Consumption Growth Parallels Income Growth: Some New Evidence. In B. Douglas Bernheim and John B. Shoven, eds., *National savings and Economic Performance* (Chicago: University of Chicago Press) 305-343.
- [17] Chen, Nai-Fu, 1991. Financial Investment Opportunities and the Macroeconomy. *Journal of Finance* 46, 529-554.
- [18] Cochrane, John H., 1994. Permanent and Transitory Components of GNP and Stock-Prices. *Quarterly Journal of Economics* 109, 241-265.
- [19] Deaton, Angus, 1992. Saving and Income Smoothing in the Cote d'Ivoire. *Journal of African Economies* 1, 1-24.
- [20] Deaton, Angus, 1992. *Understanding Consumption* (Oxford, Clarendon Press).

- [21] Deaton, Angus, 1997. *The Analysis of Household Surveys* (The World Bank, Johns Hopkins University Press).
- [22] Fama, Eugene F., 1990. Stock Returns, Expected Returns, and Real Activity. *Journal of Finance* 45, 1089-1108.
- [23] Flavin, Margorie, 1981. The Adjustment of Consumption to Changing Expectations about Future Income. *Journal of Political Economy* 89, 974-1009.
- [24] Friedman, Milton, 1957. *A Theory of the Consumption Function* (Cambridge, NBER).
- [25] Gourinchas, Pierre-Olivier and Parker, Jonathan A., 2002. Consumption Over the Lifecycle. *Econometrica* 70, 47-89.
- [26] Hansen, Lars P., Roberds, William, and Sargent, Thomas J., 1991. Time Series Implications of Present Value Budget Balance and of Martingale Models of Consumption and Taxes. In Lars P. Hansen and Thomas J. Sargent, eds., *Rational Expectations Econometrics* (Boulder, Westview) 121-161.
- [27] Hansen, Lars Peter and Singleton, Kenneth J., 1982. Generalized Instrumental Variables Estimation of Non-linear Rational Expectations Models. *Econometrica* 50, 1269-1286.
- [28] Lettau, Martin and Ludvigson, Sydney, 2001. Consumption, Aggregate Wealth, and Expected Stock Returns. *Journal of Finance* 56, 815-849.
- [29] Lettau, Martin and Ludvigson, Sydney, 2002. Time-varying risk premia and the cost of capital: An alternative implication of the Q theory of investment. *Journal of Monetary Economics* 49, 31-66.
- [30] Lettau, Martin and Ludvigson, Sydney, 2002. Tay's as good as cay: Reply. Working paper, New York University.

- [31] Nalewaik, Jeremy J., 2003. The Response of Consumption and Wives' Leisure to Future Income Growth: Theory and Synthetic Panel Evidence. PhD Dissertation, University of Chicago Graduate School of Business.
- [32] Newey, Whitney K. and West, Kenneth D., 1987. A Simple, Positive Semi-definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55, 703-708.
- [33] Pakes, Ariel and Griliches, Zvi, 1984. Estimating Distributed Lags in Short Panels with an Application to the Specification of Depreciation Patterns and Capital Stock Contracts. *Review of Economic Studies* 51, 243-262.
- [34] Paxson, Christina H., 1992. Using Weather Variability To Estimate the Response of Savings to Transitory Income in Thailand. *American Economic Review* 82, 15-33.
- [35] Pischke, Jorn-Steffen, 1995. Individual Income, Incomplete Information, and Aggregate Consumption. *Econometrica* 63, 805-840.
- [36] Sargent, Thomas J., 1978. Rational Expectations, Econometric Exogeneity, and Consumption. *Journal of Political Economy* 86, 673-700.
- [37] U.S. Department of Labor, Bureau of Labor Statistics, 1980-1999. Consumer Expenditure Survey, Interview Survey Files.
- [38] U.S. Department of Labor, Bureau of Labor Statistics, 1981-2000. Current Population Survey, Annual Demographic Files.

Appendix A: A Derivation of the Consumption Function

Household h 's liquid asset holdings follow: $A_{t+1}^h = (1+r)(A_t^h + X_t^h - C_t^h)$, where C_t^h the household's level of consumption at time t , A_t^h is asset holdings, r is the (assumed constant) interest rate, and X_t^h is the level of time t exogenous "labor" income. The present value version of the budget constraint is:

$$(A1) \quad A_t^h + \sum_{j=0}^{\infty} \frac{X_{t+j}^h}{(1+r)^j} = W_t^h = \sum_{j=0}^{\infty} \frac{C_{t+j}^h}{(1+r)^j},$$

where we've defined wealth W_t^h in the standard way. Employing log-linearizations, the expected value of this budget constraint (with respect to the information set of the household) can be converted into a decomposition of consumption growth.

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:

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

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

$$(A2) \quad c_t^h \approx w_t^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \Delta c_{t+j}^h - \left(\frac{1+r}{r} \right).$$

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 (A1) 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:

$$(A3) \quad w_t^h \approx y_t^h + \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \Delta y_{t+j}^h + \ln \left(\frac{1+r}{r} \right).$$

Next take the time t conditional expectation of (A2); this equation and its lead are:

$$(A4) \quad \begin{aligned} c_{t+1}^h &\approx E_{t+1} w_{t+1}^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} E_{t+1} \Delta c_{t+1+j}^h - \ln \left(\frac{1+r}{r} \right) \\ c_t^h &\approx E_t w_t^h - \sum_{j=1}^{\infty} \frac{1}{(1+r)^j} E_t \Delta c_{t+j}^h - \ln \left(\frac{1+r}{r} \right). \end{aligned}$$

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): $w_{t+1}^h = r + k + (\frac{1}{\lambda})(w_t^h) + (1 - \frac{1}{\lambda})(c_t^h)$, where $\lambda = 1 - \overline{C/W}$ is one minus the mean consumption-wealth ratio.²⁵ Notice that if we take the unconditional mean of the right hand side of (A1), we get $\frac{1}{1+r} = \lambda$. So these are interchangeable; the paper largely employs the λ notation. Use this law of motion to substitute out w_{t+1}^h from the first equation of (A4), then multiply the second equation of (A4) 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, and rearranging using $r + \ln(\lambda) \approx 0$, we have:

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

Finally, use (A3) to substitute income growth terms for the innovation in w_t^h . This yields

²⁵The constant $k = \ln(\lambda) - (1 - \frac{1}{\lambda}) \ln(1 - \lambda)$.

our decomposition of consumption growth:

$$\begin{aligned}
 \Delta c_{t+1}^h &= E_t \Delta c_{t+1}^h - \sum_{j=1}^{\infty} \lambda^j (E_{t+1} - E_t) \Delta c_{t+j+1}^h \\
 &+ \sum_{j=1}^{\infty} \lambda^{j-1} (E_{t+1} - E_t) \Delta y_{t+j}^h + \nu_{t+1}^{h,BC}.
 \end{aligned}
 \tag{A5}$$

The last term here, $\nu_{t+1}^{h,BC}$, is a summary statistic representing the approximation errors or remainders from taking the first order approximations (A2) and (A3).

If consumption growth is unpredictable, so that conditional expectations of consumption growth are constant, (A5) essentially reduces to (3).²⁶ To make this unpredictability assumption more concrete, consider a utility function for household h that 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 the N_{t+j}^h are taste shifters that impact the household's utility from consumption, such as household size and possibly leisure. For illustrative purposes, assume that consumption and the taste shifters are jointly log-normal,²⁷ so we can write (see Hansen and Kenneth J. Singleton (1982)):

$$E_t (\Delta c_{t+j}^h) = \frac{r - \delta^h}{\gamma^h} + \frac{E_t (\Delta n_{t+j}^h)}{\gamma^h} + \frac{1}{2\gamma^h} \text{var}_t (-\gamma^h \Delta c_{t+j}^h + \Delta n_{t+j}^h),
 \tag{A6}$$

where $-\delta^h = \ln(\beta^h)$. Then a constant conditional expectation of consumption growth amounts to assuming constant conditional expectations for the taste shifters and a con-

²⁶Higher order moments of consumption growth and other variables entering the budget constraint (in $\nu_{t+1}^{h,BC}$) must be constant as well for (A5) to reduce to (3).

²⁷If the normality condition is not met we can still write a first order approximation to the non-linear Euler equation similar to (A6), with the conditional variance term replaced by a composite approximation error including all higher order moments.

stant conditional variance for the marginal utility of consumption.

Appendix B: The HRS Derivation of Equation (1), and a Two Dimensional Example

HRS start by constructing the matrix Q , whose first row is $\frac{\rho^h(\lambda)}{|\rho^h(\lambda)|}$, where $|\rho^h(\lambda)| = \sqrt{\rho^h(\lambda) \rho^h(\lambda)'}$, and whose $n-1$ other 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 its 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:

$$\begin{aligned}
 \Delta c_{t+1}^h &= \rho^h(\lambda) \varepsilon_{t+1} \\
 &= \rho^h(\lambda) Q'Q\varepsilon_{t+1} \\
 \text{(B1)} \qquad &= |\rho^h(\lambda)| \varepsilon_{1,t+1}^+.
 \end{aligned}$$

$\varepsilon_{1,t+1}^+$ is the one dimension of the consumer's information set revealed to the econometrician by consumption growth; as noted in the text it is a linear combination of elements of ε_{t+1} .

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

$$\begin{aligned}
 \Delta y_{t+1}^h &= \kappa_{t+1}^h + \rho^h(\mathbf{L}) \varepsilon_{t+1} \\
 &= \kappa_{t+1}^h + \rho^h(\mathbf{L}) Q'Q\varepsilon_{t+1} \\
 \text{(B2)} \qquad &= \kappa_{t+1}^h + \rho_1^{h,+}(\mathbf{L}) \varepsilon_{1,t+1}^+ + \rho_2^{h,+}(\mathbf{L}) \varepsilon_{2,t+1}^+.
 \end{aligned}$$

The polynomial coefficients $\rho_1^{h,+}(\mathbf{L}) = \frac{\rho^h(\lambda)}{|\rho^h(\lambda)|} \rho^h(\mathbf{L})'$ are the moving average coefficients corresponding to $\varepsilon_{1,t+1}^+$.

We can use the expression for consumption growth (B1) to substitute $\varepsilon_{1,t+1}^+$ out of

(B2), yielding:

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

where in the second line we have rewritten the regression coefficients in the notation of equation (1), where $\beta(\mathbf{L}) = \frac{\rho_1^{h,+}(\mathbf{L})}{|\rho^h(\lambda)|} = \frac{\rho^h(\lambda)}{|\rho^h(\lambda)|^2} \rho^h(\mathbf{L})'$. The $\beta(\mathbf{L})$ in (1) are the normalized moving average coefficients corresponding to $\varepsilon_{1,t+1}^+$. The other $n - 1$ dimensions of household information about income growth end up in the error term: $e_{t+1}^h = \kappa_{t+1}^h + \rho_2^{h,+}(\mathbf{L}) \varepsilon_{2,t+1}^+$.

To clarify this general HRS derivation, we'll work through an example where $n = 2$. Assume the income process of the household takes the following form:

$$\begin{aligned}\Delta y_{t+1} &= \varepsilon_{1,t+1} + \phi \varepsilon_{1,t} + \varepsilon_{2,t+1} \\ &= \begin{bmatrix} 1 + \phi \mathbf{L} & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{bmatrix},\end{aligned}$$

where we've dropped the h superscripts and ignore the linearly deterministic component of income growth. The innovations are uncorrelated and each has a variance of unity, so $\text{var } \Delta y_{t+1} = 2 + \phi^2$. The econometrician would forecast Δy_{t+1} with $R_{MAX}^2 = \frac{\phi^2}{2 + \phi^2}$ if the two shocks were separately observable; this is the maximum possible R^2 attainable in forecasting Δy_{t+1} . Hamilton (1994, Ch. 4) derives the univariate representation of this income process: $\Delta y_{t+1} = \varepsilon_{t+1} + \theta \varepsilon_t$. Not surprisingly, $|\theta| < |\phi|$, and the precision with which the econometrician forecasts Δy_{t+1} deteriorates when only the composite univariate shock process is observed.

For consumption growth, the benchmark model predicts: $\Delta c_{t+1} = (1 + \phi \lambda) \varepsilon_{1,t+1} +$

$\varepsilon_{2,t+1}$, and it can be verified that a valid Q matrix is:

$$Q = \begin{bmatrix} \frac{1+\phi\lambda}{\sqrt{(1+\phi\lambda)^2+1}} & \frac{1}{\sqrt{(1+\phi\lambda)^2+1}} \\ \frac{-1}{\sqrt{(1+\phi\lambda)^2+1}} & \frac{1+\phi\lambda}{\sqrt{(1+\phi\lambda)^2+1}} \end{bmatrix}.$$

Decomposing the income process as in (B1):

$$\begin{aligned} \Delta y_{t+1} &= \begin{bmatrix} 1 + \phi L & 1 \end{bmatrix} Q' Q \begin{bmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{bmatrix} \\ &= \frac{(1 + \phi\lambda)(1 + \phi L) + 1}{(1 + \phi\lambda)^2 + 1} [(1 + \phi\lambda) \varepsilon_{1,t+1} + \varepsilon_{2,t+1}] \\ &\quad + \frac{-(1 + \phi L) + (1 + \phi\lambda)}{(1 + \phi\lambda)^2 + 1} [-\varepsilon_{1,t+1} + (1 + \phi\lambda) \varepsilon_{2,t+1}]. \end{aligned}$$

Use Δc_{t+1} to substitute $(1 + \phi\lambda) \varepsilon_{1,t+1} + \varepsilon_{2,t+1}$ out of this equation; after some rearrangements and simplifications, we have:

$$\begin{aligned} \Delta y_{t+1} &= \frac{(2 + \phi\lambda)}{(1 + \phi\lambda)^2 + 1} \Delta c_{t+1} + \frac{\phi(1 + \phi\lambda)}{(1 + \phi\lambda)^2 + 1} \Delta c_t \\ &\quad + \frac{\phi^2 \lambda^2}{(1 + \phi\lambda)^2 + 1} \left(1 - \frac{1}{\lambda} L\right) [\varepsilon_{1,t+1} + \varepsilon_{2,t+1}], \end{aligned}$$

where it can be verified that $\beta(\lambda) = 1$. The fraction of variance of Δy_{t+1} captured by consumption growth lagged one period, Δc_t , can be written as $R_{\Delta c_t}^2 = \left(\frac{(1+\phi\lambda)^2}{(1+\phi\lambda)^2+1}\right) \frac{\phi^2}{2+\phi^2}$, and is strictly less than R_{MAX}^2 .²⁸ This is sensible: Δc_t is a weighted average of two shocks, one white noise, and the part of Δc_t reflecting this shock will contribute nothing towards predicting Δy_{t+1} . The residual from a regression of Δy_{t+1} on $(\Delta c_{t+1}, \Delta c_t)$ is itself partially predictable. Write this residual, $\rho_2^+(L) \varepsilon_{2,t+1}^+$, in its invertible representation

²⁸For $\phi > 0$, Δc_t has more explanatory power for Δy_{t+1} than does the history of the univariate representation of Δy_{t+1} . For $\phi < 0$, the history of the univariate representation of Δy_{t+1} has more explanatory power than Δc_t .

(see Hamilton, 1994, Ch. 3):

$$\rho_2^+ (\mathbf{L}) \varepsilon_{2,t+1}^+ = \frac{\phi^2 \lambda^2}{(1 + \phi \lambda)^2 + 1} \left(\frac{1}{\lambda} \right) (1 - \lambda \mathbf{L}) \varepsilon_{t+1},$$

where we now let ε_{t+1} denote $\varepsilon_{1,t+1} + \varepsilon_{2,t+1}$. One thing to notice is that the regressors $(\Delta c_{t+1}, \Delta c_t)$ draw out of the income process all of its persistence, leaving behind only transitory variation in levels. In the more general n -dimensional case, the variation left behind in $\rho_2^+ (\mathbf{L}) \varepsilon_{2,t+1}^+$ will exhibit less persistence than the variation in the income process as a whole; empirical results bear out this implication.

Appendix C: Standard Errors for the Weighted Least Squares Estimator

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 ; let \mathbf{X} denote the $N\bar{T} \times K$ matrix of regressors; let $\mathbf{\Omega}$ represent the $N\bar{T} \times N\bar{T}$ variance-covariance matrix of regression residuals; finally let \mathbf{W} denote the $N\bar{T} \times N\bar{T}$ diagonal matrix with the vector of weights on the diagonal.²⁹ Then, following standard practice, the variance-covariance matrix of the OLS parameter estimates is computed as:

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

$\mathbf{X}'\mathbf{W}\mathbf{\Omega}\mathbf{W}\mathbf{X}$ is the sum of two additive components. The first component (which we shall call $\mathbf{X}'\mathbf{W}\mathbf{\Omega}_0\mathbf{W}\mathbf{X}$) is robust to both arbitrary heteroskedasticity and cross correla-

²⁹The weights in this paper are the square root of the average of the synthetic cohort cell counts of the cells used to produce the explanatory consumption variables - i.e. in regression specification (1) if $w_{i,t+1}$ is the weight on the i th synthetic person at time $t + 1$ and $f_{i,t+1}$ is the number of households in the sample at time $t + 1$ who meet the qualifications to belong to the i th synthetic cohort, we have $w_{i,t+1} = \sqrt{\frac{\sum_{k=0}^{q+1} f_{i,t+1-k}}{q+2}}$.

tion. The matrix can be represented as:

$$\mathbf{X}'\mathbf{W}\mathbf{\Omega}_0\mathbf{W}\mathbf{X} = \sum_{i=1}^N \sum_{j=1}^N \sum_{t=1}^{T_{ij}} (x'_{i,t} w_{i,t} u_{i,t} u_{j,t} w_{j,t} x_{j,t} + x'_{j,t} w_{j,t} u_{j,t} u_{i,t} w_{i,t} x_{i,t}),$$

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:

$$\begin{aligned} \mathbf{X}'\mathbf{W}\mathbf{\Omega}_0\mathbf{W}\mathbf{X} &= \sum_{t=1}^T \sum_{i=1}^{N_t} \sum_{j=1}^{N_t} (x'_{i,t} w_{i,t} u_{i,t} u_{j,t} w_{j,t} x_{j,t} + x'_{j,t} w_{j,t} u_{j,t} u_{i,t} w_{i,t} x_{i,t}) \\ &= \sum_{t=1}^T \mathbf{X}'_t \mathbf{W}_t \mathbf{\Omega}_{0,t} \mathbf{W}_t \mathbf{X}_t. \end{aligned}$$

Here N_t denotes the number of cross sectional observations in year t , and \mathbf{X}_t denotes the N_t rows of \mathbf{X} that correspond to the year t cross sectional observations. Similarly, $\mathbf{\Omega}_{0,t}$ denotes the matrix $\mathbf{u}_t \mathbf{u}'_t$, the outer product of the vector of time t regression 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 (ignoring any autocorrelation for the moment), with each 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.

A second component of the estimated matrix $\mathbf{X}'\mathbf{W}\mathbf{\Omega}\mathbf{W}\mathbf{X}$ corrects for autocorrelation

as suggested by Whitney K. Newey and Kenneth D. West (1987):

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

In this paper k' is set to seven, so the matrix corrects for seventh order autocorrelation.

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

$$\mathbf{X}'\mathbf{W}\Omega\mathbf{W}\mathbf{X} = \mathbf{X}'\mathbf{W}\Omega_0\mathbf{W}\mathbf{X} + \mathbf{X}'\mathbf{W}\Omega_k\mathbf{W}\mathbf{X}.$$

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	26	56	121	106	130	269

Notes to Table 1: Summary statistics on the number of households in each of the 524 group-year cells (i.e each of the 524 synthetic panel observations). Each CEX household is assigned to a unique year to avoid double counting, although its monthly consumption data may span two years.

Table 2: Summary Statistics

Variable	Data Source	Standard Deviation	Autocorrelations						
			1	2	3	4	5	6	7
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)	-0.19 (0.08)
Δc	CEX	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)	0.09 (0.08)

Notes to Table 2: Standard deviations are computed weighting each observation by the average cell counts of the two observations used to compute each growth rate. Autocorrelations are computed by weighted least squares, where the weights are the average cell counts of the explanatory variables. Asymptotic standard errors are in parentheses; details of their computation are in Appendix A. Control variables are quartic age polynomials and birth cohort fixed effects, both interacted with education fixed effects.

Table 3:**Panel A: Estimates 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							0.07	
	(0.04)								484
1	0.16	0.12						0.07	
	(0.04)	(0.03)							444
2	0.13	0.12	0.16					0.11	
	(0.04)	(0.03)	(0.03)						408
3	0.13	0.14	0.18	0.12				0.15	
	(0.04)	(0.03)	(0.04)	(0.03)					372
4	0.10	0.13	0.19	0.14	0.09			0.16	
	(0.04)	(0.03)	(0.04)	(0.03)	(0.03)				336
5	0.06	0.09	0.18	0.16	0.13	0.12		0.18	
	(0.04)	(0.04)	(0.04)	(0.03)	(0.05)	(0.04)			300
6	0.08	0.07	0.16	0.15	0.16	0.14	0.04	0.17	
	(0.05)	(0.04)	(0.04)	(0.03)	(0.05)	(0.04)	(0.05)		268
6			0.14	0.16	0.18	0.16	0.06	0.16	
			(0.04)	(0.03)	(0.05)	(0.04)	(0.05)		268

Panel B: Univariate Estimates of $\Delta y_t = \beta_k \Delta c_{t-k} + e_t$

	β_0	β_1	β_2	β_3	β_4	β_5	β_6
β	0.09	0.05	0.13	0.10	0.12	0.13	0.06
	(0.04)	(0.03)	(0.04)	(0.04)	(0.06)	(0.04)	(0.04)
R^2	0.02	0.00	0.04	0.02	0.03	0.03	0.01

Table 4:**Panel A: Estimates 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								0.07	
	(0.09)									484
1	0.26	0.35							0.10	
	(0.07)	(0.06)								444
2	0.23	0.36	0.02						0.09	
	(0.07)	(0.08)	(0.07)							408
3	0.24	0.35	0.01	0.01					0.08	
	(0.07)	(0.09)	(0.09)	(0.07)						372
4	0.20	0.31	0.03	0.01	0.11				0.07	
	(0.08)	(0.10)	(0.09)	(0.09)	(0.10)					336
5	0.18	0.28	0.02	-0.03	0.10	0.07			0.06	
	(0.08)	(0.11)	(0.09)	(0.10)	(0.11)	(0.11)				300
6	0.21	0.28	0.05	-0.06	0.01	-0.06	0.22		0.07	
	(0.09)	(0.11)	(0.09)	(0.10)	(0.15)	(0.11)	(0.10)			268
7	0.20	0.26	0.09	-0.08	0.03	-0.07	0.22	0.02	0.07	
	(0.11)	(0.12)	(0.09)	(0.10)	(0.17)	(0.15)	(0.13)	(0.17)		240

Panel B: Univariate Estimates of $\Delta c_t = \beta_k \Delta y_{t-k} + e_t$

	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7
β	0.20	0.24	0.13	-0.02	0.09	-0.02	0.20	0.01
	(0.12)	(0.12)	(0.10)	(0.08)	(0.13)	(0.14)	(0.12)	(0.16)
R^2	0.02	0.02	0.01	0.00	0.00	0.00	0.01	0.00

Notes to Tables 3 and 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-k} terms are income growth computed from the CPS; the Δc_{t-k} terms are consumption growth 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 A.

Table 5: Estimates 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:	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	$\sum_{k=0}^q \lambda^k \beta_k$	Implied r	$R_{\Delta c}^2$
	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
Age, Cohorts	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	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
Age, Cohorts, Years	-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-k}^{CPS}	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-k}^{CPS} , Age, Cohorts	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

Notes to Tables 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 dependent variable Δy_t is income growth computed from the CPS. The explanatory variables Δc_{t-k} are consumption growth 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-k}^{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 A.

In computing the discounted sum of the regression coefficients $\sum_{k=0}^q \lambda^k \beta_k$, $\lambda = \frac{1}{1+r}$ is set to $\frac{1}{1.025}$. The implied r is the value of r for which the discounted sum of the regression coefficients sums to one. $R_{\Delta c}^2$ is fraction of variance the dependent variable explained by consumption growth, after orthogonalizing the data with respect to the control variables.

Table 6:

**Estimates 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**

Cutoff	β_0	β_1	β_2	β_3	β_4	β_5	β_6	$\sum_k \lambda^k \beta_k$
Age ≤ 33	0.48 (0.18)	0.29 (0.16)	0.37 (0.12)	0.33 (0.09)	0.28 (0.13)	0.18 (0.09)	0.11 (0.08)	1.93 (0.64)
Age ≤ 38	0.22 (0.09)	0.08 (0.06)	0.19 (0.08)	0.13 (0.07)	0.10 (0.05)	0.11 (0.07)	0.09 (0.06)	0.88 (0.29)
Age ≤ 43	0.14 (0.08)	0.17 (0.06)	0.25 (0.07)	0.22 (0.06)	0.15 (0.07)	0.13 (0.06)	0.07 (0.06)	1.05 (0.22)
Age > 33	0.07 (0.06)	0.11 (0.06)	0.20 (0.06)	0.20 (0.06)	0.24 (0.08)	0.21 (0.07)	0.07 (0.07)	1.01 (0.27)
Age > 38	0.03 (0.06)	0.09 (0.07)	0.20 (0.06)	0.24 (0.05)	0.31 (0.09)	0.25 (0.06)	0.05 (0.08)	1.09 (0.24)
Age > 43	0.05 (0.09)	0.05 (0.09)	0.22 (0.08)	0.19 (0.10)	0.41 (0.09)	0.37 (0.08)	0.13 (0.10)	1.30 (0.43)

Table 7:

**Estimates 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**

Yrs of School	β_0	β_1	β_2	β_3	β_4	β_5	β_6	$\sum_k \lambda^k \beta_k$
Less than 12	0.09 (0.09)	0.09 (0.08)	0.14 (0.09)	0.08 (0.13)	-0.11 (0.17)	0.08 (0.12)	-0.02 (0.13)	0.33 (0.55)
12	0.04 (0.11)	0.08 (0.12)	0.33 (0.10)	0.25 (0.09)	0.23 (0.10)	0.14 (0.09)	0.12 (0.08)	1.10 (0.36)
13-15	0.18 (0.08)	0.18 (0.11)	0.19 (0.07)	0.21 (0.06)	0.29 (0.11)	0.18 (0.08)	0.11 (0.07)	1.25 (0.28)
16 or More	0.03 (0.12)	0.02 (0.11)	0.09 (0.09)	0.20 (0.10)	0.36 (0.10)	0.28 (0.06)	0.06 (0.10)	0.94 (0.37)

Notes to Tables 6 and 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 growth computed from the CPS; the explanatory variables Δc_{t-k} are consumption growth 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 consumption growth variables. Asymptotic standard errors are in parentheses; details of their computation are in Appendix A.

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 a quartic age polynomial and birth cohort fixed effects.

Table 8: Estimates of $\Delta w_t = \beta_0 \Delta c_t + \beta_1 \Delta c_{t-1} + \dots + \beta_q \Delta c_{t-q} + \text{controls} + e_t$

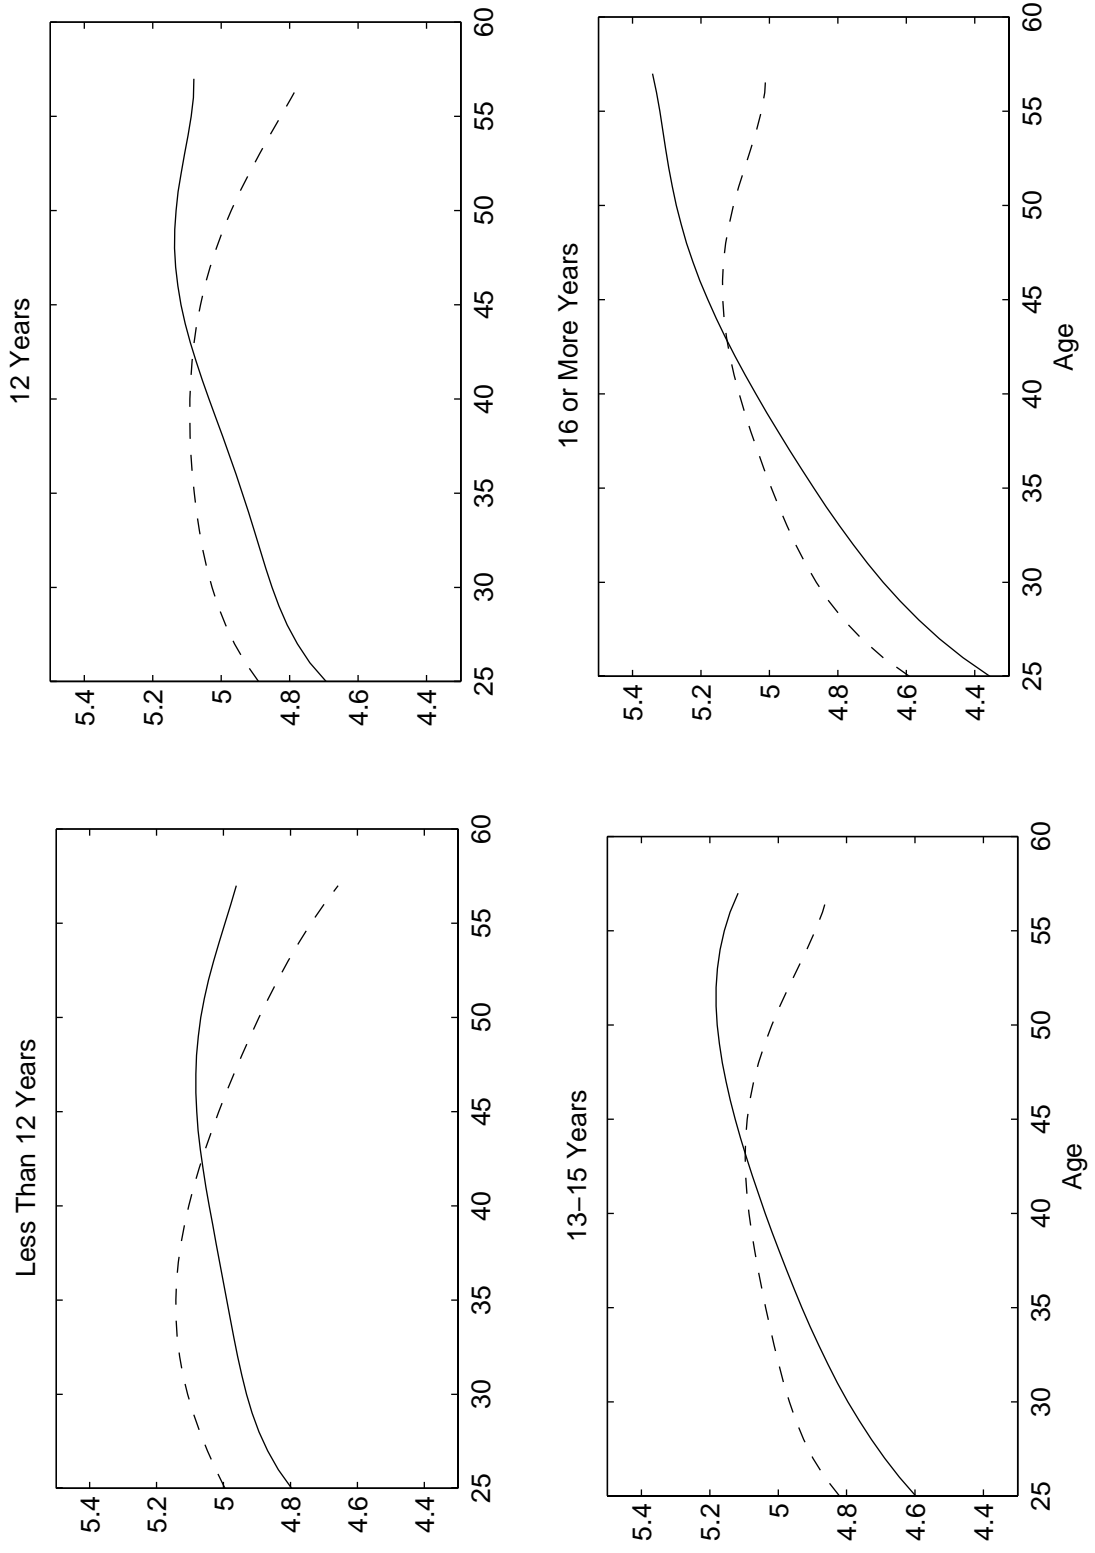
Controls for:	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	$R^2_{\Delta c}$
	-0.01	0.10	0.14	0.15	0.14	0.05	-0.00		0.14
	(0.04)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)	(0.06)		
Age, Cohorts	0.07	0.22	0.27	0.26	0.22	0.08	0.03		0.13
	(0.05)	(0.05)	(0.07)	(0.07)	(0.05)	(0.06)	(0.05)		
Years	-0.00	0.06	0.03	0.00	0.08	0.08	0.14	0.07	0.15
	(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.04)	(0.03)	
Age, Cohorts, Years	0.00	0.08	0.07	0.05	0.14	0.13	0.17	0.07	0.07
	(0.04)	(0.03)	(0.03)	(0.03)	(0.05)	(0.05)	(0.06)	(0.03)	

Notes to Table 8: 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 dependent variable Δw_t is the wage growth of male household heads, computed from the CPS. The explanatory variables Δc_{t-k} are consumption growth 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 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 A.

In computing the discounted sum of the regression coefficients $\sum_{k=0}^q \lambda^k \beta_k$, $\lambda = \frac{1}{1+r}$ is set to $\frac{1}{1.025}$. The implied r is the value of r for which the discounted sum of the regression coefficients sums to one. $R_{\Delta c}^2$ is fraction of variance the dependent variable explained by consumption growth, after orthogonalizing the data with respect to the control variables.

Figure 1: Life Cycle Log Income (solid) and Log Consumption (dashed) Profiles, by Education of Male Household Head



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 productsp (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 accomodate 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)

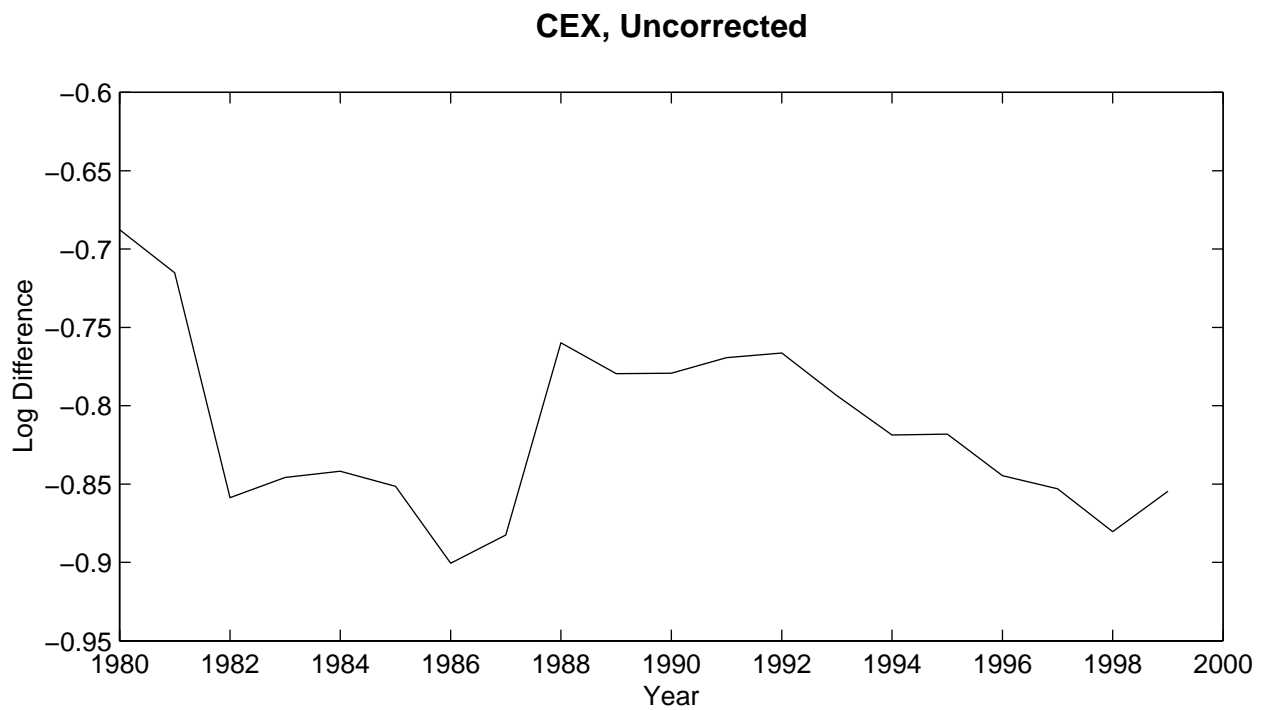
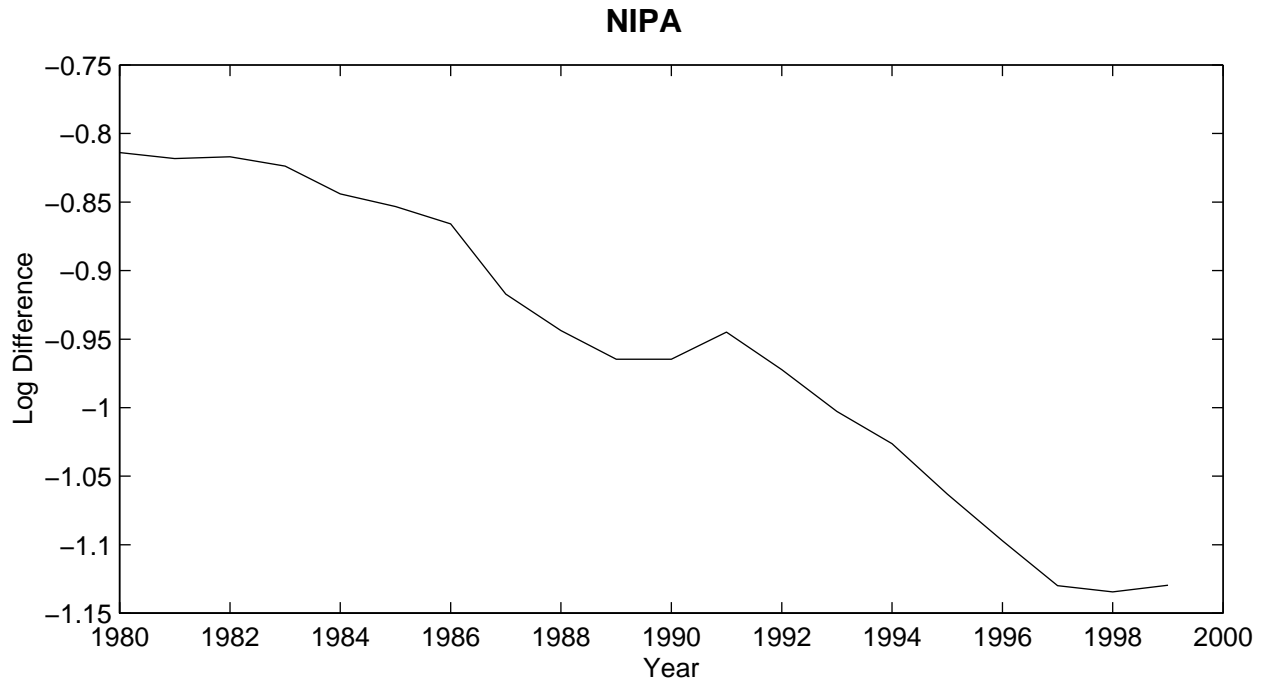


Figure B.3: Corrected (Dashed) and Uncorrected (Solid) CEX Real Food at Home Data

