CAN PERMANENT-INCOME THEORY EXPLAIN CROSS-SECTIONAL CONSUMPTION PATTERNS?

John Sabelhaus and Jeffrey A. Groen*

Abstract—The prediction that consumption-income ratios should decline as income rises in cross-sectional data is a feature of Friedman's (1957) permanent income hypothesis and other consumption-smoothing models. The theory thus provides a link between longitudinal income data and cross-sectional expenditure data: given measured income variability and a functional relationship between consumption and permanent income, we predict cross-sectional expenditure patterns and compare those predictions to actual values. Our approach cannot explain the actual skewness in consumption-income ratios under even the strictest consumption-smoothing model, which implies that income measurement error or other anomalies are affecting the data.

I. Introduction

THE prediction that consumption-income ratios will be negatively correlated with income in cross-sectional data is a well-known feature of Friedman's (1957) permanent income hypothesis (PIH) and other intertemporal consumption-smoothing theories. Some families with low annual income have higher permanent income, so average consumption in low-income groups will be greater than average income. At the top of the annual income distribution, the opposite holds: some families with high annual income have lower permanent income, so average consumption in high-income groups will be less than average income.

In this paper, we apply the consumption-smoothing principle quantitatively to investigate whether measured longitudinal income variability is consistent with actual cross-sectional spending patterns. The cross-sectional spending data we explore here are used in a broad range of applied microeconomic studies, including analyses of how tax burdens would change under a consumption tax (Caspersen & Metcalf, 1994; Feenberg, Mitrusi, & Poterba, 1997) and whether income and consumption generate different conclusions about the distribution of economic well-being (Cutler & Katz, 1991; Slesnick, 1993). If there is reason to doubt the validity of the observed relationship between spending patterns and income in the cross-sectional data, the results of those studies are in doubt.

Our main finding is that the observed correlation between spending and income in cross-sectional expenditure data cannot be reconciled with observed longitudinal income volatility (both transitory income fluctuations and random measurement errors) using even the strictest consumption-smoothing models. Simply put, there is not enough variability in household-level income over time to explain the dramatic skewness in consumption-income ratios across the income distribution, even when consumption is assumed to be completely insensitive to transitory income fluctuations.

The analysis here does not constitute a test of consumption-smoothing theories in any sense. Rather, we are testing if the deviation between current measured income and permanent income (as proxied by the ten-year average of income) can explain the observed pattern of consumption-income ratios. Those deviations between current and permanent income include both transitory fluctuations and random measurement error, both of which have the same effect on predicted consumption-income ratios. Because no version of the theory can reconcile the available longitudinal income and cross-sectional income-expenditure data, we conclude that the measured consumption-income patterns are probably significantly affected by nonrandom income measurement error or other anomalies.

II. Consumption-Income Ratios in Annual Survey Data

Although aggregate consumption is a large and relatively stable share of aggregate income over time, the PIH and other consumption-smoothing models imply that the ratio of consumption to annual income will vary predictably across income groups. This section documents the pattern of consumption-income ratios in the Consumer Expenditure Survey (CEX) for 1992, using various combinations of cash-expenditure concepts, distributional statistics, and age groupings.[1]

Several tabulations from the CEX are shown in table 1. The distributional classifier is before-tax family income divided by the family size adjustment implicit in the Census poverty thresholds. The census scale suggests, for example, that a family with two members requires only 28% more income to reach the same level of well-being as a single individual, because of shared resources and economies of scale.[2] The first column is the ratio of total expenditures to total (after-tax) income in each decile.[3] The pattern of

Received for publication April 29, 1997. Revision accepted for publication June 2, 1999.

* Congressional Budget Office and University of Michigan, respectively.

The views here do not necessarily reflect the opinion of the Congressional Budget Office. We presented an earlier version of this paper at the January 1997, American Economic Association meetings in New Orleans, Louisiana. We thank Frank Sammartino for useful comments and suggestions at every stage of this research. We also thank Eric Engen, seminar participants at the Congressional Budget Office, and two anonymous referees for helpful comments.

[1] See appendix A and Sabelhaus (1996) for a description of the CEX sample and expenditure concept used in this study. The cash-expenditure concept includes some items that are not direct consumption, such as interest payments and property taxes. On the other hand, it does not include some items that are consumption, such as the rental value of owned housing. There will be systematic divergence in the two measures over the lifecycle, but, as we discuss in the text, those do not invalidate our tests.

[2] The Census scale adjustments for family sizes 2 through 9 are, respectively, 1.28, 1.57, 2.01, 2.38, 2.68, 3.04, 3.38, and 4.04. All of the results here and throughout the paper are basically the same across three approaches we tested: no adjustment for family size, a per adult adjustment, and the Census adjustment described and used in the paper.

[3] It is important to note that these are not average consumption-income ratios. They are ratios of average consumption to average income. The former can show even more skewness, particularly if very low-income families are included.

The Review of Economics and Statistics, August 2000, 82(3): 431–438

TABLE 1.—ALTERNATIVE MEASURES OF CONSUMPTION-INCOME RATIOS

Income Decile	Ratio of Average Consumption to Average Income: Total Consumption	Ratio of Average Consumption to Average Income: Excluding Durables	Median Consumption-Income Ratio	Ratio of Average Consumption to Average Income by Age of Head: Age <40	Age 40–60	Age 60+
1	2.30	2.18	1.86	1.91	3.06	2.44
2	1.37	1.30	1.22	1.33	1.30	1.46
3	1.34	1.23	1.19	1.33	1.32	1.34
4	1.12	1.02	1.02	1.08	1.10	1.21
5	1.00	0.92	0.93	1.02	1.00	0.98
6	0.95	0.86	0.89	0.89	1.00	0.98
7	0.90	0.81	0.84	0.88	0.88	0.95
8	0.81	0.74	0.75	0.87	0.78	0.76
9	0.74	0.67	0.70	0.73	0.74	0.73
10	0.64	0.59	0.64	0.65	0.64	0.61
All incomes	0.88	0.81	0.92	0.90	0.82	0.95

Notes: Data are from 1992 Consumer Expenditure Survey. Decile rankings are based on Adjusted Family Income (AFI) measure, which adjusts for economies of scale implicit in census poverty thresholds across family size. Nondurables measure excludes furniture and motor vehicle purchases. Sample excludes top-coded observations. See appendix A for details.

expenditure-income ratios is similar to other findings using the CEX (Sabelhaus, 1993; Feenberg, Mitrusi, & Poterba, 1997; Poterba, 1989). Families in the bottom decile spend 230% of their income, while families in the top decile spend only 64%. The pattern across deciles is nonlinear and convex: the bottom five deciles have negative saving rates, while the top five have positive rates.

The distributional pattern of expenditure-income ratios across income groups is robust with respect to alternative expenditure measures. The measure in the second column excludes durable goods, because those purchases are volatile and partly represent investment. For example, someone who earns $20,000 a year and buys a $10,000 car that will last five years is consuming $2,000 worth of car per year. If the durable purchase is included in total expenditures, the ratio of consumption to income will be overstated. But, even when the expenditure concept excludes durables, the skewness in expenditure-income ratios across income groups persists. That is, the impact of durable purchases averages out across the many families within a given decile, some of whom have purchases and some of whom do not.

The third column reports median expenditure-income ratios. If average expenditure within a decile is strongly influenced by a few outliers, the median ratios are a better indicator of typical expenditure behavior within the group. Although the median expenditure-income ratio in the bottom decile is a bit lower than the mean ratio, the overall pattern remains. The typical family in the bottom decile spends 186% of disposable income, which is below the average of 230%. In the top decile, the median and average ratios of expenditures to income are identical at 64%.

The last three columns show that, after controlling for income, age has little explanatory effect. The overall expenditure-income ratio is lower for middle-aged people (forty to sixty) than it is for the young (younger than forty) or old (older than sixty). This observation seems consistent with lifecycle versions of consumption-smoothing theories: middle-aged people are at the point in their life when saving should be high. But, surprisingly, that finding disappears within any given income decile. For example, in the top decile, the young spend 65% of their income, the middle-aged 64%, and the old 61%. In general, the overall skewness in expenditure-income ratios across income groups holds within any given age group, as well as in the aggregate.[4]

III. Can Permanent-Income Theory Explain the Skewness in Consumption?

The skewness of consumption-income ratios across annual income groups is a well-known characteristic of expenditure survey data. Does that skewness result mainly from consumption-smoothing behavior? The appropriate data set for addressing that question is a panel survey with annual expenditures and annual income over a long period. Unfortunately, such a data set does not exist for the United States. So we use panel income data, then assess whether the longitudinal income variability and a functional relationship between consumption and permanent income can explain the observed skewness in the expenditure data.

The data used to measure income variability are from the Panel Study of Income Dynamics (PSID). The sample covers the period from 1982 through 1991 and includes most of the 1991 sample members who were in the survey for the ten-year period.[5] Our income measure in the PSID is the same as in the CEX: total family income adjusted for family size. Changes in family size, holding total income constant, will change a family's relative income position.[6]

Our measure of "permanent" income is the average (adjusted) annual income over this ten-year period.[7] We

[4] These results should not be interpreted as evidence against the lifecycle model for two reasons. First, as noted, we are measuring cash expenditure, which includes items that are not consumption, such as interest payments and investment in durable goods. Also, our measure omits important age-related consumption components such as the imputed rental value of housing and the value of employer- and government-provided medical services. A comprehensive measure of consumption that makes those adjustments (such as in Gokhale, Kotlikoff, and Sabelhaus (1996)) does exhibit a lifecycle pattern. Second, our measure of "income" is also limited. We do not include employer pension contributions or Social Security taxes in saving, and we do include pension and Social Security benefits in income. If we compute consumption and income in a way consistent with lifecycle theory, the expected pattern would show up in the tables. The calculations here are focused on reconciling only the cash-flow piece of saving.

[5] We excluded people who left their parents' family and established new families during the period. See appendix A for a more detailed description of the PSID sample used here.

[6] This implicitly assumes that adjusted family income is the right way to measure well-being across families. An alternative strategy is to restrict the sample to families that were intact over the ten years. That approach avoids the need to adjust income for family composition change. We prefer the approach in the paper because it uses much more of the sample. In any case, we used the restricted-sample approach in an earlier version of the paper and generated similar results.

[7] Our permanent income measure differs from the standard CEQ LC/PIH definition, which, as in Deaton (1992), is given by $(r/(1 + r))$ times net worth plus the present value of labor income. Carroll (1997) shows that the measure used here is more appropriate in a standard model with realistic labor income uncertainty.

remove the effect of economy-wide, real-income growth by indexing average incomes across the years. Thus, average income in each of the sample years is the same, but the income-variability measures capture idiosyncratic movements (including progression through age-earnings profiles) for any given family.

Table 2 shows various statistics from the annual income and permanent income distributions in our PSID sample. The annual income distribution is based on all ten years of data: each family shows up ten times, based on their annual income in each year. (The same basic result can be derived using any year's annual data, but the decile breaks and averages within deciles are slightly sensitive to the exact year chosen, especially within the thin upper tail of the distribution.) The main message of table 2 is that classifying people by annual income does indeed produce a more dispersed distribution than does classifying people by permanent income: annual decile breaks and average incomes within deciles are lower at the bottom and higher at the top in the annual distributions.

Table 3 shows the effects of cross-tabulating observations by permanent and annual income deciles. Again, each observation shows up in table 3 ten times to avoid thin distributions in any given year. The distribution of annual incomes within any permanent income decile (column) can be read by moving down the annual deciles (rows). For example, a family whose permanent income places it in the bottom decile (less than $7,600, table 2) has a 69.6% chance of being in the bottom annual decile (less than $6,420). It has a 23.8% chance of being in the second annual decile (between $7,660 and $11,750), a 4.2% chance of being in the third annual decile, and a much smaller chance of being in any of the fourth through tenth annual deciles.

A similar decomposition for annual income groups can be read off the rows of table 3. Again, a family whose annual income places them in the bottom decile in a given year has a 69.6% chance of being in the bottom permanent decile, a 17.1% chance of being in the second, 6.7% in the third, and so on. Putting the numbers in context, average consumption in the bottom annual decile in any given year will equal (0.696) × (average consumption in the bottom permanent income group) + (0.171) × (average consumption in the second permanent income group) + (0.067) × (average consumption in the third permanent income group), and so on.[8]

Table 3 suggests that families' decile rankings are relatively stable, particularly among the very poor and very rich. Assuming income is measured without error, one would conclude that about 70% of the permanent poor are annual poor, and about 70% of the permanent rich are annual rich. Almost all income variability is restricted to plus or minus one decile. There is virtually no overlap between the extremes of the permanent and annual income distributions, although it is much more likely for a permanent-rich person to have a bad year and show up in the lower annual deciles than for a permanent-poor person to have a good year and show up in the higher annual deciles.[9]

In our first experiment with these income-variability estimates, we simply compute the pattern of annual consumption-income ratios, assuming that the simple PIH holds. Consumption in a given annual income decile is the sum of consumption over all families in that annual income decile, assuming that each family's consumption is proportional to its permanent income. For annual income group k, total consumption is the sum of consumption across families whose annual income is in the appropriate range (all $i \in k$);

TABLE 2.—ADJUSTED DECILE BREAKS AND AVERAGE INCOMES IN THE PSID

Income Decile	Lower Bound of Adjusted Income Deciles: Permanent Income	Lower Bound of Adjusted Income Deciles: Annual Income	Average Adjusted Income Across Income Deciles: Permanent Income	Average Adjusted Income Across Income Deciles: Annual Income
1	na	na	$ 5,317	$ 3,623
2	$ 7,660	$ 6,420	9,668	8,322
3	11,750	10,330	13,588	12,305
4	15,470	14,240	17,433	16,198
5	19,350	18,090	21,190	20,093
6	23,080	22,100	25,039	24,201
7	27,070	26,540	29,502	29,153
8	32,070	32,010	35,039	35,467
9	38,360	39,400	43,335	44,971
10	50,360	52,470	77,067	82,918
All incomes	na	na	27,729	27,730

Notes: Data are from the Panel Survey of Income Dynamics, 1982–1991. All measures are adjusted for family size using the implicit scale economies in the census poverty thresholds across family-size groups. See appendix A for a complete description of the PSID sample used here.

TABLE 3.—CROSS-TABULATION OF PERMANENT AND ANNUAL INCOME STATUS IN THE PSID (ALL VALUES ARE PERCENTAGES.)

Annual Income Decile	Permanent Income Decile 1	2	3	4	5	6	7	8	9	10
1	**69.6**	17.1	6.7	2.7	1.6	0.9	0.4	0.4	0.2	0.3
2	23.8	**46.2**	17.5	6.3	3.1	1.5	0.8	0.5	0.3	0.3
3	4.2	25.8	**35.5**	17.7	8.6	4.2	2.1	1.0	0.5	0.3
4	0.9	7.8	26.0	**32.1**	17.9	8.0	4.5	2.2	0.7	0.1
5	0.5	1.8	8.9	24.4	**29.0**	19.8	9.5	3.8	1.8	0.4
6	0.2	1.1	3.5	10.6	22.5	**28.9**	19.7	8.9	3.5	1.0
7	0.2	0.3	1.4	4.0	10.5	22.3	**29.1**	20.8	9.7	1.9
8	0.0	0.1	0.6	1.4	4.1	10.2	22.3	**35.6**	21.2	4.4
9	0.1	0.1	0.2	0.7	2.2	3.7	9.4	22.2	**42.3**	19.3
10	0.0	0.0	0.1	0.1	0.3	0.8	1.8	5.0	20.0	**72.1**

Notes: Data are from the Panel Survey of Income Dynamics, 1982–1991. Annual and permanent decile rankings are based on Adjusted Family Income (AFI), which adjusts for economies of scale implicit in census poverty thresholds across family size. See appendix A for a complete description of the PSID sample used here. Diagonal elements are bolded.

[8] These calculations probably overstate the amount of predicted variability in consumption, because the deciles are discrete, and thus small changes in income for people close to cutoffs will cause large predicted changes in consumption.

[9] The impression from table 3 is quite different from a similar table (4–8) in Fullerton and Rogers (1993). Their table cross-classifies people by annual and lifetime incomes, and shows significantly more dispersion. The tables differ for two main reasons. Our table uses data from single families over ten years, whereas the Fullerton and Rogers table uses data from estimated age-income profiles across eleven discrete groups over entire lifetimes. Also, their age-income profiles include the effect of economy-wide real wage growth, which we eliminate by construction.

TABLE 4.—ACTUAL AND SIMULATED CONSUMPTION-INCOME RATIOS ACROSS ADJUSTED FAMILY INCOME DECILES

Income Decile	Actual (CEX)	Simulated (PSID)	
		Simple PIH	50% Transitory Offset
1	2.30	1.67	1.22
2	1.37	1.16	1.00
3	1.34	1.03	0.94
4	1.12	0.96	0.91
5	1.00	0.96	0.91
6	0.95	0.90	0.89
7	0.90	0.90	0.89
8	0.81	0.85	0.87
9	0.74	0.84	0.86
10	0.64	0.76	0.83
All incomes	0.88	0.88	0.88

Notes: The actual ratios are from 1992 Consumer Expenditure Survey (CEX), and the simulated ratios are based on the PSID, 1982–1991. Decile rankings are based on the Adjusted Family Income (AFI) measure, which adjusts for economies of scale implicit in census poverty thresholds across family size groups. See appendix A for a complete description of the data used here.

TABLE 5.—RATIOS OF PERMANENT TO ANNUAL INCOME AND CONSUMPTION TO PERMANENT INCOME ACROSS ANNUAL INCOME DECILES

Income Decile	Ratio of Average Permanent Income to Average Annual Income	Ratio of Average Consumption to Average Permanent Income
1	1.90	1.21
2	1.32	1.04
3	1.17	1.15
4	1.09	1.03
5	1.09	0.92
6	1.02	0.93
7	1.02	0.88
8	0.97	0.84
9	0.95	0.78
10	0.86	0.74
All incomes	1.00	0.88

Notes: Permanent income for any given annual income decile is computed using the PSID for 1982–1991, and consumption values are from 1992 Consumer Expenditure Survey (CEX). Decile rankings are based on the Adjusted Family Income (AFI) measure, which adjusts for economies of scale implicit in census poverty thresholds across family size groups. See appendix A for a complete description of the data used here.

that is,

$$C_k = \sum_{i \in k} \beta y_i, \tag{1}$$

where y_i is permanent (after-tax) income of family i, and β is the ratio of consumption to permanent income for all families in the population.

The simulated consumption-income ratios under the simple PIH are shown in table 4, along with actual consumption-income ratios from the CEX (reproduced from table 1). There is substantial skewness in the distribution of simulated ratios, as the PIH predicts. But that distribution is not nearly as skewed as the actual distribution. In the bottom decile, the simple PIH predicts that the annual consumption-income ratio will be 1.67, but the actual ratio in the CEX is 2.30. In the top decile, the simple PIH predicts a ratio of 0.76, but the ratio in the CEX is 0.64. In a loose sense, then, the simple PIH, together with our estimate of the relationship between annual and permanent income, is able to explain approximately half the skewness in consumption-income ratios in the annual data.[10]

It is tempting to go further and ask if minor deviations from the simple PIH can explain the remaining skewness in consumption-income ratios. For example, the PIH may generally hold, but liquidity constraints or uncertainty may cause people to adjust their consumption when their annual income varies.[11] However, any adjustments that make consumption less a function of permanent income and more a function of annual income will reduce the predicted skewness. For example, the simulated consumption-income ratios in the last column of table 4 are based on the assumption that families adjust consumption by a constant fraction (50%) of the difference between permanent income and annual income. The skewness in the simulated ratios in the 50% offset case is substantially less than in the simple PIH case.

A second way to show how the theory and data differ is presented in table 5. The first column is the ratio of average permanent income to average annual income across annual income deciles. As the PIH predicts, lower income families have (on average) substantially higher permanent income, by a factor of 1.90. At the top of the income distribution, permanent income is below annual income as the theory predicts, by a factor of 0.86. But these ratios are not enough to account for the skewness in annual consumption-income ratios. In the second column, we compute the ratio of annual consumption (CEX) to permanent income (PSID). These would be equal across income groups if the simple version of the theory held.

IV. Varying the Propensity to Consume Out of Permanent Income

The declining pattern of consumption to permanent income across annual income groups suggests one explanation for why the theory cannot reconcile the longitudinal income variability and cross-sectional spending data: long-run propensities to consume vary across permanent income groups.[12] Thus, our second set of calculations solves for a schedule of propensities to consume out of permanent income that, when combined with observed income variability, allow us to replicate the cross-sectional consumption-income ratios.

In order to solve for a pattern of consumption propensities across permanent income groups that is consistent with the actual consumption-income ratios in the CEX, we need to

[10] One immediate explanation for the divergence is systematic differences in consumption-income ratios over the lifecycle, but as table 1 showed, controlling for age does not affect the conclusion.

[11] See, for example, Carroll (1997), Hubbard, Skinner, and Zeldes (1995), and Huggett and Ventura (1996). For a comprehensive review of theory and facts about household-level saving and consumption behavior, see Browning and Lusardi (1996).

[12] A declining pattern of consumption propensities across permanent income groups is consistent with predictions of models that emphasize the redistribution built into real-world tax and transfer systems. See, for example, Hubbard, Skinner, and Zeldes (1995), Huggett and Ventura (1996), and Sabelhaus (1997).

specify a functional relationship between the consumption propensity and permanent income itself. Let y_j denote average permanent income in permanent income group j. Then, the observed curvature in consumption-permanent income ratios (table 5) suggests a quadratic in the propensity to consume out of permanent income. That is,

$$\beta_j = \gamma_0 + \gamma_1 y_j + \gamma_2 y_j^2. \quad (2)$$

Let n_{kj} denote the fraction of families in annual income group k that are also in permanent income group j. Predicted consumption in annual group k is the weighted average of consumption across permanent income groups, where the weights are just the population cross-tabs (n_{kj}). That is,

$$C_k = \sum_j n_{kj}(\gamma_0 + \gamma_1 y_j + \gamma_2 y_j^2) y_j. \quad (3)$$

The cross-tabulated distribution of families across permanent and annual income groups (n_{kj}) are from table 3, and income means across permanent income groups (y_j) are from table 2. We estimate the γ's by minimizing the sum of squared deviations between actual and simulated consumption-income ratios across annual income groups.

The result of fitting the quadratic to permanent consumption propensities is shown in table 6. The three-parameter model does a good job of replicating the actual consumption-income ratios in the annual data. (This is simply a statement about curve-fitting, not model consistency in any sense.) The main conclusion in table 6 is in the third column: the estimated pattern of propensities to consume across permanent income groups (values of the group-level β's) are much less skewed than the annual consumption-income ratios. The fitted permanent propensities to consume range from 127% in the bottom decile to 73% in the top decile, whereas the annual ratios range from 230% to 64%.

It is important to note again that these results are based on simple consumption-smoothing. If consumption depends on both permanent and transitory income, more skewness in permanent consumption propensities will be needed to replicate the annual consumption-income patterns. But it is clear that allowing differences in permanent consumption propensities can reconcile the longitudinal income variability and observed cross-sectional expenditure patterns.

TABLE 6.—ACTUAL AND SIMULATED CONSUMPTION-INCOME RATIOS USING FITTED PROPENSITIES TO CONSUME OUT OF PERMANENT INCOME

Income Decile	Ratio of Average Annual Consumption to Average Annual Income: Actual	Ratio of Average Annual Consumption to Average Annual Income: Simulated	Fitted Propensity to Consume Out Of Permanent Income
1	2.30	2.33	1.27
2	1.37	1.40	1.17
3	1.34	1.24	1.10
4	1.12	1.09	1.03
5	1.00	1.04	0.98
6	0.95	0.94	0.93
7	0.90	0.90	0.89
8	0.81	0.82	0.82
9	0.74	0.79	0.77
10	0.64	0.60	0.73
All incomes	0.88	0.88	0.88

Notes: The actual ratios are from 1992 Consumer Expenditure Survey (CEX), and the calculated ratios are based on the 1982–1991 PSID using a fitted (quadratic) propensity to consume out of permanent income. Decile rankings are based on the Adjusted Family Income (AFI) measure, which adjusts for economies of scale implicit in census poverty thresholds. See appendix A for a complete description of the data used here.

V. Reconciling Consumption-Income Ratios With Wealth-Based Saving Measures

We have shown that the longitudinal income and cross-sectional consumption data can be reconciled using the PIH in the case in which propensities to consume vary across permanent income groups, but the next step is to ask whether those implied long-run consumption propensities are reasonable. Without empirical evidence, it seems that long-run divergence between income and spending for the lowest income groups in table 6 is simply unsustainable. But we can go beyond that casual inference and compare the results to direct empirical estimates of the long-run relationship between income and spending. Various studies, including Kennickell and Starr-McCluer (1997) and Dynan, Skinner, and Zeldes (1996), use longitudinal wealth panels to estimate the distribution of saving across permanent income groups.[13] Those studies implicitly estimate consumption-income ratios because saving equals income minus consumption.

In the last section, we estimated that the (implied) saving rate in the bottom permanent decile is −27% of income, but the wealth-based studies listed above produce values near zero. We estimated that saving in the top permanent decile is 27%, but the wealth-based estimates are generally near 15%, or even a few percentage points lower.[14] These differences are too large to discount. And, as noted above, any divergence from pure consumption-smoothing will require more skewness in permanent consumption propensities than that shown in table 6.

More evidence against the possibility that consumption-smoothing underlies all the skewness in consumption-income ratios comes from looking at the distribution of wealth across and within income groups at any given point in time. Table 7 shows the distribution of wealth holdings

[13] The data sets used in these studies are the 1983–1989 panel of the Survey of Consumer Finances (SCF) and the PSID, which collected wealth information in 1984 and again in 1989.

[14] The significant variability in estimates from the wealth-change studies is due to the facts that the available panels are small, they often have substantial attrition, and even some conceptual differences exist. Some studies measure wealth change over shorter periods, but none are based on very recent data. For example, Bosworth, Burtless, and Sabelhaus (1991) show that one-year wealth change in the 1962–1963 SCF also ranges from near zero in the bottom income quintile to approximately 16% in the top quintile. They also measure a three-year wealth change for 1983–1986 and find a similar pattern, although slightly negative in the bottom of the income distribution and only approximately 12% in the top income quintile. It is interesting to note that the skewness in the one- and three-year wealth-change estimates is actually below that in the longer-run measures, which is contrary to the theory.

TABLE 7.—THE DISTRIBUTION OF NON-HOUSING WEALTH WITHIN ADJUSTED INCOME GROUPS
(ALL VALUES ARE RATIOS OF WEALTH TO AVERAGE INCOME WITHIN GROUP.)

Income Decile	Average Non-Housing Wealth	Percentiles of the Non-Housing Wealth Distribution								
		10th	20th	30th	40th	50th	60th	70th	80th	90th
1	2.5	−1.0	−0.3	−0.1	0.0	0.0	0.1	0.1	0.1	1.6
2	0.7	−0.6	−0.2	−0.1	0.0	0.0	0.0	0.1	0.4	1.4
3	1.1	−0.5	−0.2	−0.1	0.0	0.0	0.0	0.2	1.1	3.2
4	1.5	−0.5	−0.2	−0.1	0.0	0.0	0.1	0.3	1.6	4.5
5	1.5	−0.3	−0.2	0.0	0.0	0.0	0.2	0.6	1.5	4.2
6	1.6	−0.4	−0.2	−0.1	0.0	0.1	0.4	1.0	2.3	5.2
7	2.0	−0.2	−0.1	0.0	0.1	0.2	0.5	1.2	2.1	4.2
8	2.0	−0.2	−0.1	0.0	0.1	0.2	0.4	1.0	1.9	4.2
9	1.9	−0.1	0.0	0.1	0.3	0.5	0.8	1.5	2.8	4.8
10	4.6	0.0	0.1	0.4	0.6	1.0	1.6	2.7	4.7	9.7
All incomes	2.8	−0.2	−0.1	0.0	0.0	0.1	0.3	0.7	1.8	4.9

Notes: Data are from 1992 Survey of Consumer Finances. Decile rankings are based on the Adjusted Family Income (AFI) measure, which adjusts for economies of scale implicit in census poverty thresholds across family size. See appendix A for a complete description of the SCF sample used here.

within income deciles in the Survey of Consumer Finances (SCF) for 1992, the same year for which we have CEX data.[15] Average non-housing wealth divided by average income is quite high in the bottom income decile, which suggests that some people with high permanent incomes experience a negative income shock and end up in the bottom decile. If those families are smoothing consumption, they may be pulling up the overall average consumption in the CEX, as well as the wealth average in the SCF.[16]

But the distribution of wealth within deciles does not support that argument. In the bottom income decile, very few families have large wealth holdings relative to their income. Thus, financing consumption by drawing down wealth may explain why average consumption in the decile is high, but not why the median consumption-income ratio (1.86, table 1) is also high. Also, very few families have large negative wealth balances, which would be consistent with borrowing to finance consumption. Reconciling the CEX flow and SCF balance data requires that many families with high permanent income dropped into the bottom annual income decile in 1992, had only approximately 1.5 year's income in the bank, and spent down those balances so that their measured wealth at the end of 1992 was zero.

VI. Measurement Errors in the Survey Data

Given the results of wealth-based saving studies, it is impossible to reconcile the observed longitudinal income variability with the skewness in consumption-income ratios, because the implied permanent consumption propensities are not realistic. It seems that one or more of the components of the standard asset-accumulation identity—wealth at two points in time, and income and consumption during the intervening period—are being measured with error in survey data. Because the CEX focuses on expenditures, not incomes, it seems reasonable to consider whether the skewness in the annual consumption-income ratios arises because income is poorly measured.

Income-measurement validation studies have been undertaken for the PSID (Duncan & Hill; 1985; Bound et al., 1994) and the CPS (Bound & Krueger, 1991), but not for the CEX. The results of those studies indicate the potential for significant measurement error: perhaps 25% of the variance in observed income is attributable to measurement error. That degree of measurement error would cause significant skewness in consumption-income ratios even if transitory income fluctuations are insignificant. Randomly distributed measurement error in the PSID is implicitly treated the same as transitory income in our analysis, because it represents a deviation from permanent income, so the effect shows up in the predicted cross-sectional consumption-income ratios.

But the issue at hand is differences in measurement error across the data sets, not the level of measurement error in any given data set. If there is simply more randomly distributed income-measurement error in the CEX, the skewness in actual consumption-income ratios will exceed predicted skewness based on the PSID, but only because income is more poorly measured in the CEX. A second possibility to consider is that there is nonrandom measurement error in both data sets.

There are reasons to believe that the PSID is better, in general, at measuring income than the CEX. The survey focuses on collecting income data and is conducted in March when families are more aware (for tax purposes) of what they earned in the previous year. Also, unlike other surveys, the CEX does not attempt to impute incomes when respondents are unable or unwilling to report their incomes. One offsetting factor is that the CEX does indicate which respondents are "incomplete" income reporters, and we use this indicator (as does BLS) to exclude observations from our sample. (See appendix for details.)

Ultimately, it is impossible to identify whether there is greater randomly distributed income-measurement error in

[15] See appendix for a description of the SCF sample used here.

[16] In the case of the CEX, it is interesting to ask whether people who are very rich but had a bad cash-flow year (probably by hiring good accountants) are even in the survey. In the SCF and Statistics of Income (SOI) individual tax file, it is clear that some very rich people are in the bottom income groups, but those samples target the very rich. The CEX poorly captures the top end of the income distribution, and it may not capture those rich people who show up at the bottom of the income distribution in those targeted surveys.

TABLE 8.—COMPARISON OF SURVEY-LEVEL INCOMES ACROSS INCOME DECILES

Income Decile	Average Income			
	CPS-SOI	CEX	PSID	SCF
1	$ 5,033	$ 4,937	$ 3,751	$ 3,827
2	10,481	9,358	9,172	8,492
3	15,627	13,603	14,590	12,433
4	21,275	18,457	20,442	16,761
5	27,082	24,025	26,378	22,562
6	33,736	30,526	32,666	28,950
7	41,472	37,997	40,406	36,703
8	51,471	48,013	50,148	46,671
9	65,817	62,181	65,621	62,727
10	143,471	95,510	121,352	144,213
All incomes	41,547	34,447	38,500	38,775

Notes: See data appendix for a description of the data sets used here. The CPS-SOI data set is the result of merging micro data from the Current Population Survey (CPS) with data from individual tax returns in the Statistics of Income (SOI) data set, as described in CBO (1987) and Kasten and Sammartino (1990). The PSID distribution comes from the cross-sectional file for 1992.

the CEX without a separate validation study, but there is one inconsistency in the explanation. If both data sets are drawn from the same true population, and one has greater random measurement errors, then the tails of that sample should be farther away from the true tails, because families near the true tails are further dispersed. We would thus expect to see lower average income at the bottom of the income distribution in the sample with greater measurement error.

Table 8 compares income across all the data sets used here, as well as the Congressional Budget Office (CBO) CPS-SOI merge file, which combines data from the Current Population Survey (CPS) and the Statistics of Income (SOI) tax return sample.[17] Relative to the CPS-SOI benchmark, the CEX average incomes within deciles are about the same as the PSID and SCF, except at the highest income levels. In the top decile, the CEX average income is much lower than any other: the data for high-income families are simply not there.[18]

Although table 8 indicates that the bottom tail of the CEX income distribution is similar to the other data sets, that evidence is not sufficient to completely discount the hypothesis that the CEX has more random income-measurement error. But it opens the door to consider a second explanation that is also consistent with the observed patterns: there may be nonrandom measurement error across all the data sets. Consider a case in which half the people in the bottom decile report an income of zero in all of the surveys, and the other half report the true value. The observed income distributions in all the data sets will be equal, and the observed consumption-income ratio in the bottom decile will be double both the predicted and true values.

It should be emphasized that the issue may not even be income measurement. There may be significant unrecorded private transfers paying for the extra consumption. But there is evidence that all the surveys suffer from some underreporting of income. Even the IRS's annual income "survey," where nonrespondents are fined or imprisoned, only finds about 82% of the income that should be reported.[19] Thus, solving the consumption-income puzzle in survey data may coincidentally change our views on how income is distributed in the population.

[17] By including the incomes of tax filers and non-filers, the CPS-SOI data set provides the most accurate income-distribution estimates available.

[18] As described in appendix A, the CEX removes top-coded observations because top-coded incomes along with unconstrained consumption will bias the consumption-income ratios up for the rich. But the fraction top-coded (incomes above $100,000) is nowhere near the fraction of people above that cutoff in the other data sets.

VII. Conclusion

We have shown that intertemporal consumption-smoothing cannot explain the discrepancy between measured income variability and the skewness in cross-sectional consumption-income ratios. Even the simplest version of the smoothing model in which consumption varies just with permanent income can explain only a part of the skewness in the consumption-income ratios found in expenditure data. Allowing the permanent propensity to consume to vary across permanent income groups does not solve the problem, because the derived long-term consumption rates are at odds with long-term wealth-change studies.

The most likely explanation for the differences between the cross-sectional data and theoretical predictions is some sort of measurement error. Either the income in the cross-sectional data is measured with more randomly distributed error than income in the other data sources, or there is similar nonrandom measurement error across all the data sets. Even though the income distribution in the expenditure data is similar to that in other surveys, all of the surveys miss a large share of aggregate income, a disproportionate share of which may belong at the bottom of the income distribution. Thus, developing better income-distribution estimates might be the key to reconciling cross-sectional consumption and wealth-change patterns.

[19] See Park (1996).

REFERENCES

Antoniewicz, Rochelle L., "A Comparison of the Household Sector from the Flow of Funds Accounts and the Survey of Consumer Finances," Federal Reserve Board (April 1996).

Bosworth, Barry, Gary Burtless, and John Sabelhaus, "The Decline in Saving: Some Microeconomic Evidence," *Brookings Papers on Economic Activity* 1 (1991), 183–256.

Bound, John, and Alan B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics* 9 (1) (1991), 1–24.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers, "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics* 12 (3) (1994), 345–368.

Browning, Martin, and AnnaMaria Lusardi, "Household Saving: Micro Theory and Micro Facts," *Journal of Economic Literature* 34 (December 1996), 1797–1855.

Carroll, Chris, "Buffer Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics* 112 (1) (February 1997), 1–56.

Caspersen, Erik, and Gilbert Metcalf, "Is a Value Added Tax Regressive? Annual Versus Lifetime Incidence Measures," *National Tax Journal* 47 (4) (December 1994), 731–746.

Congressional Budget Office, *The Changing Distribution of Federal Taxes* (1987).
Cutler, David, and Lawrence Katz, "Macroeconomic Performance and the Disadvantaged," *Brookings Papers on Economic Activity* 2 (1991).
Deaton, Angus S., *Understanding Consumption* (New York: Oxford University Press, 1992).
Duncan, Greg J., and Daniel H. Hill, "An Investigation into the Extent and Consequences of Measurement Error in Labor-Economic Survey Data," *Journal of Labor Economics* 3 (4) (1985), 508–532.
Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes, "Do The Rich Save More?" Federal Reserve Board, Dartmouth College, and Columbia University (July 1996).
Feenberg, Daniel R., Andrew W. Mitrusi, and James M. Poterba, "Distributional Effects of Adopting A National Retail Sales Tax," NBER working paper no. 5885N (January 1997).
Friedman, Milton, *A Theory of the Consumption Function* (Princeton, N.J.: Princeton University Press, 1957).
Fullerton, Don, and Diane Lim Rogers, *Who Bears the Lifetime Tax Burden?* (Washington, D.C.: Brookings Institution, 1993).
Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus, "Understanding the Post-War Decline in U.S. Saving," *Brookings Papers on Economic Activity* 1 (1996), 315–407.
Hill, Martha S., *The Panel Study of Income Dynamics: A User's Guide* (Newbury Park, CA: Sage Publications, 1992).
Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (1995), 360–399.
Huggett, Mark, and Gustavo Ventura, "Understanding Why High Income Households Save More Than Low Income Households," Mimeograph, University of Illinois (October 1998).
Kasten, Richard A., and Frank J. Sammartino, "A Method for Simulating the Distribution of Combined Federal Taxes Using Census, Tax Return, and Expenditure Microdata," in Gordon H. Lewis and Richard C. Michel (Eds.), *Microsimulation Techniques for Tax and Transfer Analysis* (Washington, D.C.: The Urban Institute Press, 1990).
Kennickell, Arthur B., and Martha Starr-McCluer, "Changes in Family Finances from 1989 to 1992: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin* (October 1994), 861–882.
——— "Household Saving and Portfolio Change: Evidence from the 1983–89 SCF Panel," *Review of Income and Wealth* 43 (4) (December 1997), 381–399.
Park, Thae S., "Relationship between Personal Income and Adjusted Gross Income," *Survey of Current Business* (May 1996), 78–93.
Poterba, James M., "Lifetime Incidence and the Distributional Burden of Excise Taxes," *American Economic Review* 79 (May 1989), 325–330.
Sabelhaus, John, "What is the Distributional Burden of Taxing Consumption?" *National Tax Journal* 46 (3) (September 1993), 331–344.
——— "Consumer Expenditure Survey: Family-Level Extracts, 1980:1–1994:1." Congressional Budget Office, (November 1996). Available at www.nber.org.
——— "Public Policy and Saving in the U.S. and Canada," *Canadian Journal of Economics* 30 (2) (May 1997), 253–275.
Slesnick, Daniel T., "Gaining Ground: Poverty in the Postwar United States," *Journal of Political Economy* 101 (February 1993), 1–38.

APPENDIX

Throughout the paper, we have discussed the Consumer Expenditure Survey (CEX) as though it was a single interview about the previous year. However, the CEX is a quarterly panel survey; we construct annual records by merging data from four quarterly interviews. We create annual expenditures by summing reported expenditures across the four quarters and take annual income as reported in the last quarter. We measure consumption-income ratios on an annual rather than a quarterly basis, because the quarterly expenditure data are more volatile and there are no true quarterly income data. The CEX interviews each family for four quarters; each quarter, one-fourth of the sample is in its first interview and one-fourth of the sample is in its last interview. We construct the 1992 sample from cohorts that entered the survey between the second quarter of 1991 and the first quarter of 1993, because a portion of their annual expenditures occurred in 1992. We use the monthly CPI to adjust all nominal values to a benchmark survey starting point (April, 1992) so that all expenditure and income data then apply to calendar year 1992.

Approximately 1,500 families enter the CEX each quarter, but only about half of them complete all four interviews and answer all the income questions, which we need to compute consumption-income ratios. The final sample size is 6,124. We adjust sample weights differentially by age and home-ownership status, because those two variables are highly correlated with attrition from the sample. We also replace reported income taxes with a calculated estimate, because the effective tax rates in the reported data are significantly biased across income groups. We impute the 5% of expenditures—mostly personal care items and other small purchases not tracked well in the three-month recall basic interview—using data from the separate CEX diary survey. The expenditure concept used is a "cash"-based measure; it differs from the usual definition of personal consumption (as in the NIPA) because we do not count imputed items like the rental value of owner-occupied housing, but we add actual cash expenditures on things like owned housing. Those expenditures include mortgage interest, maintenance, and property taxes, but not principal repayments, which are a form of saving. The other big difference between our measure and NIPA consumption is that we count only out-of-pocket health expenditures; the consumption value of government and employer-provided health care are not available in the data. See Sabelhaus (1996) for more details.

As noted in the text, our sample excludes the CEX "incomplete income reporters" who refused to answer questions about their incomes during the interview. We do not, however, exclude any families that BLS believes answered the income questions properly, no matter how low their actual incomes were. The other major exclusion is for top-coded observations: the tables are based on a sample that excludes the 157 observations whose incomes were above the threshold. Each component of a given person's income is top coded at $100,000, but family incomes can be higher than that if there are multiple sources of income or multiple earners. It does not make sense to use top-coded observations, however, because income is capped but consumption is not. So we drop those observations, but reweight observations in the top decile by the fraction dropped to keep constant the other decile breaks in the income distribution. (Just dropping the top-coded observations would pull the whole percentage distribution to the right.)

We construct our Panel Study of Income Dynamics (PSID) sample from the longitudinal individual file and cross-sectional family files for 1982 through 1991. (For a detailed description of the PSID, see Hill (1992).) We construct longitudinal family records by attaching family information to the longitudinal record of one individual from each family in the 1991 cross section. We start with the entire core sample in 1991, then discard the individuals who were not in the sample in any period and individuals who left their parents' family at any point during the period, so the final sample size is 4,987. Because we have defined the longitudinal family records around an individual in each 1991 family, changes in family composition over time are represented in the data set. Our income measure accounts for those changes in family size. We use the 1991 family weight in all of our calculations.

The Survey of Consumer Finances (SCF) data set used to estimate the distribution of wealth is based on all 3,906 observations in the 1992 survey, which consists of a standard area-probability sample and a high-income sample. (For a general overview of the SCF, see Kennickell and Starr-McCluer (1994).) We create a measure of non-housing wealth that matches the household-level concept in the Flow of Funds Accounts (FFA). This task was made easier because of the reconciliation between the SCF and FFA in Antoniewicz (1996).

The Current Population Survey-Statistics of Income (CPS-SOI) file used here underlies the extensive research on income distribution at the Congressional Budget Office (CBO) in the last decade. (See, for example, CBO (1987) or Kasten and Sammartino (1990).) The file is constructed by adjusting income values in the March CPS so that the distribution of (taxable) incomes in the resulting file matches the distribution in the SOI, which reflects amounts reported on tax returns. The bottom, non-filer part of the income distribution is based on CPS data, and the rest of the distribution, including the high-income families whose incomes are top-coded on the CPS, is consistent with the taxable values on the SOI. To construct the file, CBO splits CPS families into tax-filing units comparable to those on the SOI. After the tax-unit data are adjusted, we reassemble families to match the unit of observation in the other data sets.