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 about half the skewness in consumption-income ratios in the annual data.

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. 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 percent) of the difference between permanent income and annual income. The skewness in the simulated ratios in the 50 percent 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 incomes

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

Let n_{kj} denote the fraction of families in annual income group k who 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,

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

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 percent in the bottom decile to 73 percent in the top decile, whereas the annual ratios range from 230 percent to 64 percent.

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 moves the theory and data much closer together.

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

We have shown that the longitudinal income and cross-section consumption data can be reconciled using the PIH in the case where propensities to consume vary across permanent income groups, but it is worth going one step further and asking whether those implied consumption propensities are consistent with estimates based on wealth-change studies. Various studies, including Kennickell and Starr-McCluer [1996] and Dynan, Skinner, and Zeldes [1996], use longitudinal wealth panels to estimate the distribution of saving across permanent income groups.[10] Those studies implicitly estimate consumption-income ratios because saving equals income minus consumption.

We estimate that the (implied) saving rate in the bottom permanent decile is -27 percent of income, whereas the wealth-based studies produce values near zero. We estimate that saving in the top permanent decile of 27 percent, but the wealth-based estimates are generally near 15 percent, or even a few percentage points lower.[11] These differences are too large to discount. And, as noted

---

[10] The data sets used in these studies are the 1983-89 panel of the Survey of Consumer Finances (SCF) and the PSID, which collected wealth information in 1984 and again in 1989.

[11] There is significant variability in estimates from the wealth-change studies because the available panels are small, they often have substantial attrition, and even some conceptual differences exist. There are studies which 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 1962-63 SCF also ranges from near zero in the bottom income quintile to about 16 percent in the top quintile. They also measure a three-year wealth change for 1983-86 and find a similar pattern, though slightly negative in the bottom of the income distribution and only about 12 percent 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

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 within income deciles in the Survey of Consumer Finances (SCF) for 1992, the same year for which we have CEX data.[12] 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.[13]

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

---

contrary to the theory.

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

[13] 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 there are very rich people 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 is may not capture those rich people who show up at the bottom of the income distribution in those targeted surveys.

with high permanent income dropped into the bottom annual income decile in 1992, had only about one and a half-year's income in the bank, and spent down those balances so that their measured wealth at the end of 1992 was zero.

Thus, if the wealth-based studies are anywhere near accurate, the PIH is numerically incapable of explaining all of 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 poorly measured in survey data. Because the CEX focusses on expenditures, not incomes, it seems reasonable to consider whether the skewness in the annual consumption-income ratios arises because income is poorly measured.[14] This is a potentially quick answer to the question at hand: if the CEX incomes are much lower than in other surveys, the high consumption-income ratios at the bottom of the income distribution can be attributed to systematic income-measurement error in the CEX.

However, the CEX income data are comparable to the data from other surveys, except at the top income deciles. 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 Statistics of Income (SOI) tax return sample. By including the incomes of tax filers and non-tax-filers, the CPS-SOI data set provides the most accurate income distribution estimates available.

Relative to the CPS-SOI benchmark, the CEX average incomes within deciles are about the

---

[14] Also, unlike other surveys, the CEX does not attempt to impute incomes when respondents are unable or unwilling to report their incomes. However, the survey does indicate which respondents are "incomplete" income reporters, and we dropped those observations from our sample. See the data appendix for details.

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.[15] Notice that the SCF, which supplements its stratified sample with a high-income sub-sample, measures average income in the top decile that is close to the CPS-SOI value. The PSID estimate is somewhat lower.

In any case, the data indicate that, throughout most of the distribution, the CEX does not suffer from any worse income reporting than the other data sets. So if income is under-reported at the low income levels in the CEX it is probably under-reported in all of the surveys. In fact, all of the income surveys suffer from some type of under-reporting—even the IRS's annual income "survey," where non-respondents are fined or imprisoned, only finds about 82 percent of the income that should be reported.[16] Thus, solving the consumption-income puzzle in survey data may coincidentally change our views on how income is distributed in the population.

VI. Conclusion

We have shown that intertemporal consumption-smoothing cannot explain the relationship between consumption and income in cross-section survey data. Even the simplest version of the smoothing model in which consumption varies just with permanent income can only explain perhaps half the skewness in the consumption-income ratios found in expenditure data. We show that the key to

---

[15] As described in the data appendix, the CEX has top-coded observations removed 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.

[16] See Park [1996].

replicating the cross-section consumption-income ratios is allowing the permanent propensity to consume to vary across permanent income groups. But, those propensities are much more skewed than comparable estimates implied by long-run wealth-based saving measures, so it is fair to say that the theory and data are unreconciled.

The most likely explanation for the differences between the cross section data and theoretical predictions is the measurement of income itself. Even though the expenditure data are not fundamentally worse at measuring income than other surveys, all of the surveys miss a large share of aggregate income, a disproportionate share of which probably belongs at the bottom of the income distribution. Thus, developing better income distribution estimates is probably the key to reconciling cross-section consumption and wealth-change patterns.

These findings raise concerns about tax and welfare studies that rely on cross-section consumption-income patterns. In particular, the cross-section expenditure data used here have also been used to predict how tax burdens would change under a consumption tax (e.g., Caspersen and Metcalf [1994], Feenberg, Mitrusi, and Poterba [1997]) or to test whether income and consumption generate different conclusions about the distribution of economic well-being (e.g., Cutler and Katz [1991], Slesnick [1993]). Not surprisingly, given the skewness in consumption-income ratios, the tax studies find significant distributional consequences from shifting to a consumption tax (when burdens are measured using an annual income classifier) and the resource-distribution studies find significant differences in levels (but not necessarily trends) of inequality across the population when measured using consumption instead of income. Our findings suggest that these conclusions are probably overstated, because income is measured poorly in the first place.

VII. References

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

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

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

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

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

Congressional Budget Office (1987) *The Changing Distribution of Federal Taxes*.

Cutler, David, and Lawrence Katz (1991) "Macroeconomic Performance and the Disadvantaged." *Brookings Papers on Economic Activity* (2).

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes (1996) "Do The Rich Save More?" Federal Reserve Board, Dartmouth College, and Columbia University (July).

Feenberg, Daniel R., Andrew W. Mitrusi, and James M. Poterba (1997) "Distributional Effects of Adopting A National Retail Sales Tax," NBER Working Paper No. 5885N (January).

Friedman, Milton (1957) *A Theory of the Consumption Function*. Princeton, N.J.: Princeton University Press.

Fullerton, Don and Diane Lim Rogers (1993) *Who Bears the Lifetime Tax Burden?* Washington, D.C.: Brookings Institution.

Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus (1996) "Understanding the Post-War Decline in U.S. Saving." *Brookings Papers on Economic Activity* 1: 315-407.

Hill, Martha S. (1992). *The Panel Study of Income Dynamics: A User's Guide.* Newbury Park, Calif.: Sage Publications.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1995) "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103: 360-399.

Huggett, Mark, and Gustavo Ventura (1996) "Understanding Why High Income Households Save More Than Low Income Households." University of Illinois (October).

Kasten, Richard A., and Frank J. Sammartino (1990) "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.* The Urban Institute Press, Washington, D.C.

Kennickell, Arthur B., and Martha Starr-McCluer (1996) "Household Saving and Portfolio Change: Evidence From the 1983-89 SCF Panel." Federal Reserve Board (December).

_____(1994) "Changes in Family Finances from 1989 to 1992: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin* (October): 861-882.

Park, Thae S. (1996) "Relationship Between Personal Income and Adjusted Gross Income," *Survey of Current Business* (May).

Poterba, James M. (1989) "Lifetime Incidence and the Distributional Burden of Excise Taxes," *American Economic Review* 79: 325-330 (May).

Sabelhaus, John (1997) "Public Policy and Saving in the U.S. and Canada." *Canadian Journal of Economics* (Forthcoming, May).

_____ (1996) "Consumer Expenditure Survey: Family-Level Extracts, 1980:1-1994:1." Congressional Budget Office, (November).

_____ (1993). "What is the Distributional Burden of Taxing Consumption?" *National Tax Journal* 46:3, 331-344 (September).

Slesnick, Daniel T. (1993) "Gaining Ground: Poverty in the Postwar United States." *Journal of Political Economy* 101: 1-38 (February).

VIII. Data 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 not covered in the basic interview using data from the separate CEX diary survey. See Sabelhaus [1996] for more details.

The CEX tables are based on a sample that does not include the 157 top-coded observations. 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 re-weight 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 Survey of Income Dynamics (PSID) sample from the longitudinal individual file and cross-section 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. Since 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