Technical Paper Series
Macroeconomic Analysis and Tax Analysis Divisions
Congressional Budget Office
Washington, D.C.

Can Permanent-Income Theory Explain Cross-Section Consumption Patterns?

John Sabelhaus
Jeff Groen

July 1997

1997-3

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

John Sabelhaus
Jeff Groen*

April 1997

Abstract

The prediction that consumption-income ratios will decline as income rises in cross-section data is a feature of Friedman's [1957] Permanent Income Hypothesis and other intertemporal consumption-smoothing models. That prediction underlies our approach; we use longitudinal income data to predict consumption-income ratios across income groups, then compare those predictions to actual values. We show that models with a fixed propensity to consume out of permanent income cannot explain the skewness in annual consumption-income ratios. Allowing the long-run propensity to consume to decline with permanent income is the key to replicating actual consumption-income ratios.

JEL Codes: D12 Consumer Economics; Empirical Analysis. E21 Consumption; Saving.

*Tax Analysis Division, Congressional Budget Office, Washington, D.C. 20515. 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, La. We thank Frank Sammartino for useful comments and suggestions at every stage of this research. We also thank Eric Engen and seminar participants at the Congressional Budget

I. Introduction

The prediction that consumption-income ratios will be negatively correlated with income in cross-section 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.

The PIH and other consumption-smoothing models are important building blocks in economic theory, but the predicted relationship between annual consumption and permanent income cannot be tested directly because no data set exists that includes both family-level consumption and incomes over time. This paper therefore investigates the relationship between consumption and permanent income indirectly. We use longitudinal income data to predict the distribution of annual consumption-income ratios under various versions of the PIH and then compare those predictions to actual cross-section data. Thus we can discern which versions of the theory are able to reconcile the longitudinal income and cross-section expenditure data.

We show first that a simple version of the PIH in which the propensity to consume is constant across permanent income groups is able to explain only about half the skewness in cross-section consumption-income ratios. Basically, given a fixed propensity to consume out of permanent income, there is not enough variability in family-level incomes over time to explain why consumption-income ratios are so high for lower-income families, and so low for higher-income families. More realistic intertemporal models with uncertainty suggest that consumption is

determined by both permanent income and annual income, and therefore those models predict even less cross-section skewness than in the simple certainty-equivalence case.[1]

However, a reasonable extension of the simple case is able to replicate the observed spending patterns. If propensities to consume decrease with permanent income, the predicted skewness in annual consumption-income ratios increases, because families in the lowest permanent income groups (who are most likely to be in the lowest annual income groups) have higher consumption propensities. We use this logic in reverse and solve for the schedule of permanent consumption propensities that replicates the observed cross-section consumption-income ratios. As expected, the fitted permanent consumption propensities across permanent income groups are much less skewed than the annual consumption propensities, because (as in the simple case) some of the observed pattern of consumption-income ratios can be attributed to income variability.[2]

We then compare the estimated propensities to consume out of permanent income to other (implied) estimates from various wealth-based saving measures. Although we find much less skewness in permanent consumption propensities than in the annual estimates, there is still more skewness than would be consistent with measuring changes in wealth over time. The fitted propensity to consume for the bottom income decile is 127 percent, but wealth-change studies find values closer to 100 percent. For the top permanent income decile, the fitted propensity to consume is 73 percent, but the wealth-change studies find values of 85 percent or more.

---

[1] 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].

[2] The 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].

Thus, the PIH and longitudinal income variability cannot completely explain the pattern of cross-section consumption-income ratios. One remaining (and likely) explanation is systematic income measurement error; people at the bottom of the income distribution may under report their income on surveys, thus overstating consumption-income ratios at the bottom. That proposition is supported by cross-section wealth holdings, because there is insufficient wealth (plus income) to account for reported consumption at low income levels. Therefore, either consumption is over reported, or income is under reported, and the latter seems more likely. If valid, that explanation also suggests problems with existing estimates of how income is distributed across the population, because the distribution of income reported in the cross-section expenditure survey is similar to the distribution in other (including income-oriented) surveys except at the highest income levels.

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 predict that the ratio of consumption to annual income will vary systematically across income groups. This section documents the pattern of consumption-income ratios in the Consumer Expenditure Survey (CEX) for 1992, using various combinations of consumption concepts, distributional statistics, and age groupings.[3]

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

[3] See the data appendix and Sabelhaus [1996] for a description of the CEX sample used in this study.

28 percent more income to reach the same level of well-being as a single individual, because of shared resources and economies of scale.[4] The first column is the ratio of total consumption to total (after-tax) income in each decile.[5] The pattern of consumption-income ratios is similar to other findings using the CEX (Sabelhaus [1993]; Feenberg, Mitrusi, and Poterba [1997]; Poterba [1989]). Families in the bottom decile spend 230 percent of their income, while families in the top decile spend only 64 percent. The pattern between across deciles is non-linear and convex—the bottom five deciles have negative saving rates, while the top five have positive rates.

The distributional pattern of consumption-income ratios across income groups is robust with respect to alternative measures of consumption. 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 consumption, the ratio of consumption to income will be overstated. But even when the expenditure concept excludes durables, the skewness in consumption-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 consumption-income ratios. If average consumption within a decile is strongly influenced by a few outliers, the median ratios are a better indicator of typical

---

[4] 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, 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.

[5] 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.

consumption behavior within the group. Though the median consumption-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 percent of disposable income, which is below the average of 230 percent. In the top decile the median and average ratios of consumption to income are identical at 64 percent.

The last three columns show that, after controlling for income, age has little explanatory effect. The overall consumption-income ratio is lower for middle-aged people (40 to 60) than it is for the young (<40) or old (60+). That observation seems consistent with life-cycle 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 percent of their income, the middle-aged 64 percent, and the old 61 percent. In general, the overall skewness in consumption-income ratios across income groups holds within any given age group, as well as in the aggregate. [6]

---

[6] These results should not be interpreted as evidence against the life-cycle model, for two reasons. First, we use the term "consumption" loosely—our cash expenditure concept includes items that are not consumption, such as 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 life-cycle 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 life-cycle theory, the expected pattern would show up in the tables. The calculations here are focussed on reconciling only the cash-flow piece of saving.

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 consumption and annual income over a long period. Unfortunately, that data set does not exist. So we use panel data on incomes to measure income variability over time, then assess whether the pattern of variability in the income data together with various versions of consumption-smoothing behavior can explain the observed consumption-income skewness in the expenditure data.

The data used to measure income variability are from the Panel Survey 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.[7] 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.[8]

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

---

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

[8] 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.

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 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 percent chance of being in the bottom annual decile (less than $6,420). It has a 23.8 percent chance of being in the second annual decile (between $7,660 and $11,750), a 4.2 percent 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 percent chance of being in the bottom permanent decile, a 17.1 percent chance of being in the second, 6.7 percent 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) plus (0.171)*(average consumption in the second permanent income group) plus (0.067)*(average consumption in the third permanent income group), and so on.

Table 3 suggests that families' decile rankings are relatively stable, particularly among the very poor and very rich. About 70 percent of the permanent poor are annual poor, and about 70 percent 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, though 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 $i \in k$), that is,

$$(1) \qquad C_k = \sum_{i \in k} \beta \, y_i \, ,$$

[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. There are a few reasons why the tables differ. 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.

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

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.

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 Exenditure 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. Non-Durables measure excludes furniture and motor vehicle purchases. Sample excludes topcoded observations. See data appendix for details.

Table 2

Adjusted Decile Breaks and Average Incomes in the PSID

Income Decile	Lower Bound of Adjusted Income Deciles		Average Adjusted Income Across Income Deciles	
	Permanent Income	Annual Income	Permanent Income	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, 1982-91. All measures are adjusted for family-size using the implicit scale economies in the Census poverty thresholds across family-size groups. See the data appendix 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 In Percents)

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, 1982-91. 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 the data appendix for a complete description of the PSID sample used here. Diagonal elements are in **bold**.

Table 4

Actual and Simulated Consumption-Income Ratios Across Adjusted Family Income Deciles

Income Decile	Actual (CEX)	Simulated (PSID) Simple PIH	Simulated (PSID) 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 Exenditure Survey (CEX), and the simulated ratios are based on the PSID, 1982-91. 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 data appendix 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 1982-91, and consumption values are from 1992 Consumer Exenditure 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 data appendix for a complete description of the data used here.

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		Fitted Propensity to Consume Out Of Permanent Income
	Actual	Simulated	
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 Exenditure Survey (CEX), and the calculated ratios are based on the 1982-91 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 data appendix for a complete description of the the data used here.

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 the data appendix for a complete description of the SCF sample used here.

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-section file for 1992.