

Do the Rich Save More?

Karen E. Dynan

Federal Reserve Board

Jonathan Skinner

Dartmouth College and National Bureau of Economic Research

Stephen P. Zeldes

Columbia University and National Bureau of Economic Research

The question of whether higher-lifetime income households save a larger fraction of their income was the subject of much debate in the 1950s and 1960s, and while not resolved, it remains central to the evaluation of tax and macroeconomic policies. We resolve this long-standing question using new empirical methods applied to the Panel Study of Income Dynamics, the Survey of Consumer Finances, and the Consumer Expenditure Survey. We find a strong positive relationship between saving rates and lifetime income and a weaker but still positive relationship between the marginal propensity to save and lifetime income. There is little support for theories that seek to explain these positive correlations by relying solely on time preference rates,

We thank Dan Bergstresser, Jeffrey Brown, Wynn Huang, Julie Kozack, Geng Li, Stephen Lin, Byron Lutz, and Marta Noguera for research assistance; Orazio Attanasio for help with the CEX data; Joe Lupton for help with the PSID data; and Andrew Samwick for help with the Social Security data and calculations. Zeldes is grateful for financial support from Columbia Business School and a TIAA-CREF Pension and Economic Research grant, and Skinner for financial assistance from the National Institute on Aging (PO1 AG19783). We thank two anonymous referees, Don Fullerton, David Laibson, Steve Levitt, James Poterba, John Sabelhaus, James Smith, Mark Warshawsky, David Weil, and seminar participants at the NBER Summer Institute, the NBER Inter-American Conference, the Universities of Michigan, Maryland, Pennsylvania, California at Berkeley, Chicago, Georgetown, Harvard, Johns Hopkins, Northwestern, New York, Stanford, and Yale for helpful comments; we are especially grateful to Casey Mulligan for detailed suggestions. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or its staff.

[*Journal of Political Economy*, 2004, vol. 112, no. 2]
© 2004 by The University of Chicago. All rights reserved. 0022-3808/2004/11202-0002\$10.00

nonhomothetic preferences, or variations in Social Security benefits. There is more support for models emphasizing uncertainty with respect to income and health expenses, bequest motives, and asset-based means testing or behavioral factors causing minimal saving rates among low-income households.

I. Introduction

It would be easy to convince a room full of noneconomists that higher lifetime income levels lead to higher saving rates. Noneconomists would tell you that low-income people cannot afford to save. Certainly a room full of journalists would need little convincing: Examples include “a sales tax would shift the tax burden from the rich to the middle class, since affluent people save a much larger portion of their earnings” (Passell 1995, p. 1) and “the poor and middle class spend a higher percentage of their income on goods than do the rich, and so, according to most economists’ studies, a value-added tax is regressive” (Greenhouse 1992, p. 1).

A room full of economists would be less easily persuaded that higher lifetime income levels lead to higher saving rates. The typical economist would point out that people with temporarily high income will tend to save more to compensate for lower future income, and people with temporarily low income will tend to save less in anticipation of higher future income. Thus, even if the saving rate is invariant with regard to *lifetime* income, we would observe people with high *current* incomes saving more than their lower-income brethren (Friedman 1957).

Moreover, one can point to several stylized facts that do not seem to support a positive correlation between saving rates and income. First, there has been no time-series increase in the aggregate saving rate during the past century despite dramatic growth in real per capita income. Also, the increasing concentration of income toward the top income quintile during the 1980s and early 1990s did not lead to higher aggregate saving rates.¹ Looking across countries, Schmidt-Hebbel and Serven (2000) found no evidence of a statistically significant link between measures of income inequality and aggregate saving rates. Turning to household data, several recent studies that estimate roughly constant wealth-income ratios across lifetime income groups (Gustman and Steinmeier 1997, 1999; Venti and Wise 1999) pose a serious challenge to the view that saving rates rise with lifetime income.

¹ Blinder (1975) finds little connection between shifts in the income distribution and the aggregate saving rate but argues that the changes in the income distribution present in postwar U.S. data are unlikely to correspond to the type of pure redistribution required by the theory.

Despite an outpouring of research in the 1950s and 1960s, the question of whether the rich save more has since received little attention. Much of the early empirical work favored the view that high-income people did in fact save a higher fraction of their income (e.g., Mayer 1966, 1972). However, studies by Milton Friedman and others that reached the opposite conclusion, together with suggestive evidence like that described above, have left “reasonable doubt” about the alleged propensity of high-lifetime income households to save more.

We return to the relationship between saving rates and lifetime income for two reasons. First, there are a wide variety of newer data sources such as the Panel Study of Income Dynamics (PSID), the Survey of Consumer Finances (SCF), the Consumer Expenditure Survey (CEX), and work on imputed saving from Social Security and pension contributions (Gustman and Steinmeier 1989; Feldstein and Samwick 1992). Second, we believe that the topic has important implications for the evaluation of economic policy. If Friedman and his collaborators did not earn a clear-cut victory in the empirical battles of the 1960s, they won the war. “Representative agent” models and many other models used for macroeconomic or microeconomic policy evaluation *assume* that saving rates are at best invariant to proportional increases in the sum of human and physical wealth (Auerbach and Kotlikoff 1987; Fullerton and Rogers 1993; Leeper and Sims 1994; Altig et al. 2001).

The link between income and saving rates bears on a number of specific issues. First, differences across income groups in saving rates and marginal propensities to save imply that the effects on aggregate consumption of shocks to aggregate income or wealth depend on the distribution of the shock across income groups (see Stoker 1986; Shapiro and Slemrod 2003); this issue figured prominently in the debate surrounding the 2001 tax “rebates.” Second, the incidence and effectiveness of reform proposals that shift taxation away from saving, such as consumption or value-added taxes or expanded 401(k) or individual retirement accounts, depend on how much saving is done by each income group. Third, the link between income and saving rates influences how changes in income inequality alter saving and possibly aggregate growth. Finally, the question of whether higher-income households save at higher rates than lower-income households has important implications for the distribution of wealth.

We find first, like previous researchers, a strong positive relationship between current income and saving rates across all income groups, including the very highest income categories. Second, and more important, we continue to find a positive correlation when we use proxies for permanent income such as education, lagged and future earnings, and measures of consumption. Estimated saving rates range from zero for the bottom quintile of the income distribution to more than 25 percent

of income for the top quintile. The positive relationship is equally strong or even more pronounced when we include imputed Social Security saving and pension contributions. Among the elderly, we find little evidence of dissaving and some suggestive evidence of slightly higher saving rates among high-income households. In sum, our results suggest strongly that the rich do save more; more broadly, we find that saving rates increase across the entire income distribution. In addition, we present evidence suggesting that the *marginal* propensity to save is greater for higher-income households than for lower-income households.

After documenting the basic patterns of saving, we consider why the rich save more. Our findings are not consistent with the predictions of standard homothetic life cycle models. Nor, as we show below, are they consistent with explanations that range from differences in time preference rates or subsistence parameters to variation in Social Security replacement rates. Alternative models of saving that can explain the empirical regularities include those with hyperbolic discounting (e.g., Laibson, Repetto, and Tobacman 1998) or those with differential asset accumulation against out-of-pocket health expenditures late in life (e.g., Smith 1999*b*; Dynan, Skinner, and Zeldes 2002).

II. The Empirical and Theoretical Background

Many economists in previous generations used both theory and empirics to assess whether people with high incomes save more than those with low incomes (e.g., Fisher 1930; Keynes 1936; Vickrey 1947; Duesenberry 1949; Hicks 1950; Pigou 1951; Friend and Kravis 1957; Modigliani and Ando 1960). In his pioneering work on the permanent income hypothesis, Friedman (1957) argued that the positive correlation between income and saving rates observed in cross-sectional data reflected individuals changing their saving in order to keep consumption smooth in the face of temporarily high or low income. He presented empirical evidence consistent with the proportionality hypothesis that individuals with high permanent income consume the same fraction of income as individuals with low permanent income. Many studies of this hypothesis followed, some supporting Friedman and some not. Evans (1969) summarized the state of knowledge about consumption in 1969, concluding that "it is still an open question whether relatively wealthy individuals save a greater proportion of their income than do relatively poor individuals" (p. 14).

In a comprehensive examination of the available results and data, Mayer (1972) disagreed, claiming strong evidence against the proportionality hypothesis. For example, using five-year income and spending measures, he found the elasticity of consumption with respect to per-

manent income to be significantly below one (0.905) and not much different from the elasticity based on one year of income.

Despite the abundance of early studies on this important question, little work has been done since the 1970s. The lull in part reflects the influential work of Lucas (1976) and Hall (1978), which shifted interest away from learning about levels of consumption or saving and toward “Euler equation” estimation techniques that implicitly examine first differences in consumption (see Browning and Lusardi 1996).

Some studies have found that *wealth* levels are disproportionately higher among households with high lifetime income (Diamond and Hausman 1984; Bernheim and Scholz 1993; Hubbard, Skinner, and Zeldes 1995). While this result could be explained by higher saving rates among higher-income households, it could also be explained by higher rates of return (on housing or the stock market, e.g.) or the receipt of proportionately more intergenerational transfers by these households. Others (Gustman and Steinmeier 1999; Venti and Wise 1999) have argued that ratios of wealth including Social Security and pension wealth to lifetime income are no higher among high-income households, consistent with the Friedman hypothesis.

To help make our question more precise, consider a life cycle/permanent income model with a bequest motive. At each age t , households maximize expected lifetime utility

$$U_{it} \equiv E_t \sum_{s=t}^T \left[\frac{U(C_{is}^*) D_{is}}{(1 + \delta_t)^{s-t}} + (D_{is-1} - D_{is}) V(B_{is}) \right], \quad (1)$$

where E is the expectation operator, C_{is}^* is nonmedical consumption for household i at age s , δ_t is the household-specific rate of time preference, B_{is} is the bequest left in the event of death, and $V(\cdot)$ is the utility of leaving a bequest. To allow for mortality risk, D_{is} is a state variable that is equal to one if the household is alive through period s and zero otherwise, and T is the maximum possible length of life.

The family begins period s with net worth (exclusive of human wealth) $A_{is-1}(1 + r_{is-1})$, where r_{is-1} is the real after-tax rate of return on non-human wealth between $s - 1$ and s . We assume that there are no private annuity markets. The family first learns about medical expenses (M_{is}), which we treat as necessary consumption that generates no utility. It next receives transfers (TR_{is}) from the government. It then learns whether it survives through the period. If not, it leaves to heirs a non-negative bequest

$$B_{is} = A_{is-1}(1 + r_{is-1}) - M_{is} + TR_{is}. \quad (2)$$

If it survives, the household receives after-tax earnings (E_{is}) and chooses

nonmedical consumption. We define total consumption as $C_{is} \equiv C_{is}^* + M_{is}$. End-of-period wealth (A_{is}) is thus

$$A_{is} = A_{is-1}(1 + r_{is-1}) + TR_{is} + E_{is} - C_{is}. \quad (3)$$

We define real annual income as $Y_{is} \equiv r_{is-1}A_{is-1} + E_{is} + TR_{is}$ and saving as $S_{is} \equiv Y_{is} - C_{is} = A_{is} - A_{is-1}$.²

Define lifetime resources (as of age t) as age t non-human wealth plus the expected present value of future earnings and transfers. Under what circumstances will saving rates be identical across households with different lifetime resources? In a life cycle model without uncertainty, there are two alternative assumptions that will cause consumption to be proportional to lifetime resources. Either we can assume that the rate of time preference (δ_t) and the rate of return (r_{it}) are constant and equal to each other (in which case the result holds for *any* time-separable utility function), or we can assume that preferences are homothetic (i.e., $U(C) = [C^{1-\gamma} - 1]/[1 - \gamma]$, in which case the result holds for any rate of interest or time preference). If all households face the same interest rates and have the same preference parameters and life span, then the ratio of consumption to lifetime resources will be the same for all households (of a given age). If one further assumes that the initial wealth and the age-earnings and age-transfers profiles of rich households are simply scaled-up versions of those for poor households, then the proportionality in consumption implies that the saving rate (at a given age) will be the same across lifetime resource groups.³

How then could saving rates differ across income groups? We consider three general classes of models: one encompasses certainty models without a bequest motive, the second allows for uncertainty with respect to future income or health expenses (but no bequest motive), and the third includes an operative bequest motive. To provide illustrative calculations of how saving rates differ across income groups in these classes of models, we present results from a simple three-period version of the model above. We think of period 1 (“young”) as ages 30–60, period 2

² Note that since r includes the total return to non-human wealth including capital gains, saving measured as income minus consumption is identical to saving measured as the change in wealth. We return to this issue below.

³ Adding uncertainty complicates the model, but again two sets of assumptions will generate the result that consumption rises proportionately with the scale factor for earnings. First, if the utility function is quadratic and δ_t and r_{it} are constant and equal to each other, consumption will be proportional to the expected value of lifetime resources, as defined above. If the utility function is not quadratic, then there is no single summary statistic that defines consumption. However, if one assumes that the utility function is isoelastic and initial wealth and all possible realizations of earnings are scaled up by a constant factor, then consumption will also be scaled up by that factor and saving rates will be identical (Bar-Ilan 1995).

(“old”) as ages 60–90, and period 3 as the time around death (when old) when medical expenditures are paid and bequests are left.⁴

A. Consumption Models with No Uncertainty and No Bequest Motive

We begin with a model with no bequest motive and no uncertainty other than about the length of life.⁵ We use an isoelastic utility function with $\gamma = 3$, a value consistent with previous studies. We examine low-income and high-income groups, with respective average first-period incomes of \$14,634 and \$78,666, based on the top and bottom quintiles of five-year average income from the 1984–89 PSID (discussed in more detail below). We assume that in the second period Social Security and pension income replace 60 percent of preretirement (or first-period) income, consistent with the overall replacement rate in Gustman and Steinmeier (1999). The (annual) rates of time preference and interest are 0.02 and 0.03, respectively, which, together with uncertain life span, result in a roughly flat pattern of nonmedical consumption over the lifetime.

As shown in table 1, for both working-age (young) and retirement age (old) households, the predicted saving rate for the low-income group is identical to that of the high-income group. The saving rate (exclusive of pension and Social Security) while young equals 12.5 percent and while old equals –16.0 percent.

There are two approaches to generating higher saving rates for higher-income households in the standard life cycle model: differences in the timing of income for these households and differences in the timing of consumption. We consider each in turn.

Differences in the timing of earnings and transfers across lifetime income groups do not change the slope of the consumption path but yield different patterns of saving. For example, Social Security programs typically provide a higher replacement rate for low-income households and thus reduce the need for these households to save for retirement (e.g., Smith 1999*b*; Huggett and Ventura 2000).⁶ As shown in table 1, increasing the replacement rate for the low-income households from 60 percent to 75 percent (and increasing their first-period Social Security taxes such that the present value of lifetime resources is unaffected) reduces saving while young—to 6.7 percent. However, it also increases saving while old to –7.5 percent. The higher replacement rate

⁴ The only purpose of the third period is to allow for medical expenses late in life; we assume no nonmedical consumption in this period.

⁵ We set the probability of living to old age (period 2) at 82 percent, on the basis of U.S. life tables from the Berkeley-Wilmoth data set (<http://www.demog.berkeley.edu/wilmoth/mortality/>). Since period 3 represents the very end of life, all households that survive to period 2 die in period 3.

⁶ Other examples include differences in age-earnings profiles, differences in life expectancy, and differences in retirement age.

TABLE 1
SIMULATED SAVING PATTERNS

	Saving Rate of Young	Saving Rate of Old
Benchmark:		
Low income	12.5	-16.0
High income	12.5	-16.0
Income replacement rate:		
Low income (75%)	6.7	-7.5
High income (60%)	12.5	-16.0
Time preference rate:		
Low income (5%)	5.4	-8.0
High income (2%)	12.5	-16.0
Income and medical expense uncertainty:		
Low income	29.9	24.2
High income	16.8	-7.9
Income and medical expense uncertainty with consumption floor (\$10,000):		
Low income	.0	.0
High income	16.8	-7.9
Bequest motive ($\mu = 1.5$):		
Low income	12.5	-16.1
High income	17.2	-2.9
Income and medical expense uncertainty, consumption floor, and bequest mo- tive:		
Low income	.0	.0
High income	18.3	-1.8

NOTE.—Default parameters are 2 percent time preference rate, 82 percent chance of surviving to be "old," 60 percent replacement rate, and 3 percent interest rate.

for the low-income households leads to less saving while working and less *dissaving* while retired.

To properly account for Social Security saving, we construct a more comprehensive first-period saving rate that includes implicit Social Security saving (equal to the present value of future Social Security benefits accrued as a result of contributions) in both the numerator (saving) and the denominator (income), an approach we also use in the empirical section. This comprehensive saving rate (not shown in the table) is identical for low-income and high-income households. In other words, adding back implicit Social Security saving "undoes" the substitution between private and public saving that might otherwise make low-income households appear to save less. Furthermore, were the Social Security program in the model progressive, providing a net present value transfer to lower-income households, their comprehensive saving rate would be greater than that of high-income households.⁷

⁷ If low-income households were unable to fully offset higher Social Security saving with lower private saving, then the comprehensive saving rate would again be higher for low-income households.

Next consider differences in the timing of consumption. If high-income households choose more rapid growth rates in consumption, they will have higher saving rates, at least at younger ages.⁸ This might happen in a world with imperfect capital markets because households with lower time preference rates would have a greater inclination toward saving (when young) and would also be more likely to have higher earnings because of greater investment in education and other forms of human capital.⁹ Alternatively, Becker and Mulligan (1997) suggest that a higher level of income might encourage people to invest in resources that make them more farsighted, steepening their consumption paths. Finally, differences across lifetime income groups in the number and timing of children could also generate differences in the timing of consumption (Attanasio and Browning 1995).

Table 1 shows that when low-income households have an (annual) time preference rate of 0.05 instead of 0.02, saving patterns look similar to those when the income replacement rate is higher for these households. Their saving rate while young drops to 5.4 percent, and their saving rate while old increases to -8.0 percent. Once again, we see higher saving by higher-income households while young and more dis-saving while old.¹⁰

A “subsistence” or necessary level of consumption will also produce differences in consumption growth rates across income levels. Informal arguments are sometimes made that subsistence levels imply that poor households have lower saving rates because they cannot “afford to save” after buying the necessities. However, this result requires that $r > \delta$; if $r < \delta$, a subsistence level of consumption causes rich households to save *less* than poor households.¹¹ Closely related are models in which the intertemporal elasticity of substitution is larger for high-income house-

⁸ See especially Samwick (1998) for a model in which δ is correlated with income. Lawrance (1991) offers empirical evidence to this effect, although Dynan (1994) shows that the patterns are not pronounced after controlling for ex post shocks to income.

⁹ With perfect capital markets, households with a high time preference would borrow to finance their education, yielding no relationship between time preference and years of schooling or earnings (see, e.g., Cameron and Taber 2000).

¹⁰ A similar logic holds if survival probabilities increase with lifetime income. Higher-income households will have higher saving rates when young but lower rates when old, conditional on surviving to that age (Skinner 1985).

¹¹ Although the need to meet the current subsistence level depresses the saving rate of lower-income households, the need to meet future requirements boosts the saving rate of those households. Because of the subsistence level, poor households will be on a more steeply sloped portion of their utility functions than rich households. As a result, they will be less willing to substitute consumption over time and will have flatter consumption paths. If $r > \delta$, the consumption paths of both rich and poor households will slope upward, and the flatter paths of poor households will be associated with lower saving rates when young. If $r < \delta$, the reverse is true: consumption paths will slope downward, and the flatter path of the poor will be associated with a higher saving rate when young. A different way to generate the result that higher-income households have higher saving rates is to assume that subsistence levels decline with age.

holds (Attanasio and Browning 1995; Atkeson and Ogaki 1996; Ogaki, Ostry, and Reinhart 1996). Finally, higher-income households may save more if they enjoy better access to investment opportunities, such as equity markets, pensions, housing, and businesses, potentially providing a higher rate of return (Yitzhaki 1987; Gentry and Hubbard 2000).¹² In sum, differences in the timing of income and differences in the timing of consumption can explain higher saving among higher-income households while young, but they also imply that these households have higher dissaving rates when old.

B. Consumption Models with Uncertainty but No Bequest Motive

Does the precautionary motive for saving imply that high-income households should save more? To answer this question, we incorporate two additional sources of uncertainty in the model. First, we allow for risk to second-period income that might be associated with earnings shocks, forced early retirement, or the loss of a spouse. We assume a discretized distribution with an equal chance of earnings either one-quarter higher or one-quarter lower than in the case of perfect certainty.¹³

Second, we allow for the possibility of large medical expenses, especially near death. For example, Hurd and Wise (1989) found a decline in median wealth of \$103,134 (in 1999 dollars) for couples following the death of a husband, and Smith (1999*a*) estimated that wealth fell following severe health shocks, by \$25,371 for households above median income and by \$11,348 for families below median income. Covinsky et al. (1994) found that 20 percent of a sample of families experiencing a death from serious illness reported that the illness had essentially wiped out their assets.¹⁴ We include uncertainty about health care expenditures that is revealed only in the final period, at the very end of life. In the model, the bad state of health occurs 10 percent of the time, and when it does, out-of-pocket expenditures are \$8,000, or one-quarter second-period income averaged across high- and low-income groups.¹⁵ We compute the average saving rate in period 2 as average saving divided by average income.

¹² This result presumes that substitution effects dominate income effects (see Elmendorf 1996). Note also that higher-income households face higher marginal tax rates, lowering their after-tax return.

¹³ This degree of uncertainty is consistent with empirical parameterizations of earnings variability (e.g., Hubbard et al. 1994).

¹⁴ On the other hand, Hurd and Smith (2001) find smaller median changes in wealth near death.

¹⁵ Crystal et al. (2000) found that elderly patients in poor health spend 28.5 percent of their income in out-of-pocket health care expenditures. See also French and Jones (2003), who find that a very small fraction of individuals experience health shocks with a present value of \$125,000 or more. In good health, health expenditures are assumed to be zero.

Table 1 shows that when these types of uncertainty are added, the saving rate while young for low-income households (29.9 percent) is *larger* than that for high-income households (16.8 percent), and low-income households continue to save even in the second period against the possibility of large final-period medical expenses. While income risk is proportional to lifetime income, health expenditures represent a greater proportional risk for low-income households. The introduction of these factors alone implies that low-income households should save *more* than high-income households.

A difficult part of fitting theoretical models to actual saving patterns is to explain why low-income households save so little. At a minimum, one requires a Medicaid program (or charity care) to avoid the implication that every elderly household needs to save against the “worst-case” health expense outcome. One additional approach is to specify hyperbolic preferences among some households but allow financial institutions such as pension plans and home ownership to be available differentially to higher-income households. This would leave many low-income households in a hyperbolic saving trap with little or no wealth accumulation (Thaler 1994; Laibson 1997; Laibson et al. 1998; Harris and Laibson 2001). For analytic convenience, we consider in this model a mechanism that relies instead on the presence of asset-based means-testing, like Medicaid or supplemental security income, combined with a consumption “floor,” to explain low saving among the bottom income group (Hubbard et al. 1995; Powers 1998; Gruber and Yelowitz 1999). We guarantee to households a consumption floor of \$10,000 in period 2 plus payment of all period 3 medical expenses; to be eligible, households must hand over all available period 2 assets. Table 1 shows that these programs lead low-income households to have zero saving when young and to dissave nothing when older.

C. *Consumption Models with a Bequest Motive*

Thus far, our model has produced only bequests that do not generate utility for the household—sometimes referred to as unintended or accidental bequests. Here we consider an operative bequest motive as in Becker and Tomes (1986) or Mulligan (1997). Suppose that individuals value the utility of their children and that earnings are mean-reverting across generations. In this case, Friedman’s permanent income hypothesis effectively applies across generations: a household with high lifetime income will save a higher fraction of its lifetime income in order to leave a larger bequest to its offspring, who are likely to be relatively worse off.¹⁶

¹⁶ An alternative model is one in which wealth per se gives utility above and beyond the flow of consumption it enables (Carroll 2000).

We implement this model by specifying an operative bequest function $V(B_{is}) = \mu[(B_{is} + YL_{is}^c)^{1-\gamma} - 1]/(1 - \gamma)$, where μ is the trade-off parameter between own consumption and bequests, and YL_{is}^c is the value of the next generation's lifetime earnings. We assume complete mean reversion of earnings, so that earnings of the children are equal to the average earnings of parents, and $\mu = 1.5$.

Saving rates in this bequest model (without income or medical expense uncertainty) are shown in table 1. Saving rates while young and old are higher for the higher-income group, where the bequest motive is operative. By contrast, lower-income households expect their children to have earnings higher than theirs and so consume their overall resources, yielding saving rates that are the same as for the life cycle model.

Finally, we consider a model with income and medical expense uncertainty, a consumption floor, and an operative bequest motive. Here, bequests are conditional on the health and income draws, so in the good states of the world, the family leaves a much larger bequest than in the bad states of the world (Dynan et al. 2002). For the high-income household, the saving rate is 18.3 percent when young and just -1.8 percent when old. For the low-income household, the saving rate is zero for both periods because saving is discouraged by the presence of an asset-tested consumption floor.

III. Empirical Methodology

Three key issues arise in designing and implementing empirical tests. The first is how to define saving. One approach is to consider all forms of saving, including realized and unrealized capital gains on housing, financial assets, owned businesses, and other components of wealth. (These capital gains should also be added to income to be consistent with the Haig-Simon definition of full income.) An alternative is to examine a definition of saving that focuses on the "active" component—that is, the difference between income *exclusive of capital gains* and consumption.

Neither saving concept is clearly superior for our purposes. Measures that include capital gains are more comprehensive, in that they include all wealth accumulation regardless of the form it takes. However, if capital gains are unanticipated as of the time the saving decision is made, then the true intentions of households may be better captured by active saving measures that exclude these capital gains. The appropriate saving concept may also depend on the question of interest. For example, capital gains should be included when measuring the adequacy of saving (ex post) for retirement. On the other hand, active saving corresponds to the supply of loanable funds for new investment and therefore may be helpful in gauging the effect of a redistribution of income on eco-

conomic growth (Gale and Potter 2002). We thus consider both measures that include capital gains (from the SCF and PSID) and active saving measures (from the CEX and PSID).

The second and third key issues are how to distinguish those with high lifetime income from those whose income is high only transitorily and how to correct for measurement error in income. As Friedman (1957, p. 29) pointed out, these issues are intertwined: “in any statistical analysis errors of measurement will in general be indissolubly merged with the correctly measured transitory component [of income].”

When we measure saving as the residual between income and consumption, measurement error in income (Y) will, by construction, show up as measurement error of the same sign in saving ($Y - C$).¹⁷ Therefore, measurement error in income, like transitory income, can induce a positive correlation between measured income and saving rates even when saving rates do not actually differ across groups with different lifetime resources. A bias arises in the other direction when we define saving as the change in wealth: measurement error in income enters only in the denominator, inducing a negative correlation between measured income and the saving rate.

To reduce the problems associated with measurement error and transitory income, we use proxies for permanent income—an approach with a long history (Mayer 1972). We consider four instruments: consumption, lagged labor income, future labor income, and education. A good instrument for permanent income should satisfy two requirements. First, it should be highly correlated with true “permanent” or anticipated lifetime income at the time of the saving decision. Second, the instrument should be uncorrelated with the error term, which includes measurement error and transitory income, so that it affects saving rates only through its influence on permanent income.

All our instruments are likely to satisfy the first requirement. What about the second? Since consumption reflects permanent income in standard models, it should be uncorrelated with transitory income and thus be an excellent instrument (see, e.g., Vickrey 1947).¹⁸ However, transitory consumption will bias the estimated relationship between saving rates and permanent income toward being negative, and measurement error will reinforce the bias when saving is defined as the difference between income and consumption. Although these (highly probable) forms of bias likely make the resulting point estimates not useful for policy analysis, we can still potentially use the point estimates to address our main question. Specifically, a finding that measured sav-

¹⁷ This assumes that the measurement error in Y is uncorrelated with that in C .

¹⁸ If some households face binding liquidity constraints, however, consumption may be correlated with transitory income.

ing rates rise with measured consumption, despite the induced bias in the opposite direction, would represent strong evidence that saving rates do rise with permanent income.

Consider next the use of lagged and future labor earnings. The longer the lags used and the less persistent transitory income is, the more likely that lagged and future labor income will be uncorrelated with transitory income. MaCurdy (1982) and Abowd and Card (1989) find that the transitory component of earnings shows little persistence. Note that if households' forecasts of future labor income are superior to what they would be on the basis solely of the income history, then using future labor income as an instrument will (as with consumption) tend to bias the estimated relationship between income and saving rates toward being negative; households predicting higher income in the future (and hence more likely allocated to a higher instrumented income quintile) will tend to save less in anticipation.

Finally, education is typically constant over time and therefore has little correlation with transitory income; there is a long tradition of using it to proxy for permanent income (Modigliani and Ando 1960; Zellner 1960). However, it may be correlated with tastes toward saving (see Mayer 1972) or have an independent effect on the ability to plan for retirement (e.g., Lusardi 1999).

A related set of issues arises with respect to our choice of denominator for the saving rate. Our theoretical analysis focused on consumption and saving relative to lifetime resources; since we cannot measure lifetime resources, we must use an imperfect proxy in our empirical work. We use "current" income as the denominator; that is, we use a five-year average for the PSID, a two-year average for the SCF, and a one-year figure for the CEX. Although these income measures are likely more influenced by transitory income and measurement error than some of the instruments for permanent income mentioned above, measurement error in the denominator is unlikely to bias median estimates of saving ratios as long as the permanent income quintiles are determined accurately. In practice, we have not found evidence that the choice of denominator is important; the patterns we find are not sensitive to switching between a one-year and two-year average of income in the SCF and switching between a one-year, five-year, and ten-year (or more) average of income in the PSID (results not reported).

Most of our results are based on a two-stage estimation procedure. In the first stage, we regress current income on proxies for permanent income and age dummies. We then use the fitted values from the first-stage regression to place households into predicted permanent income quintiles and create an indicator dummy for each predicted income quintile. The quintiles of predicted permanent income are created separately for each age group. In the second stage, we estimate a median

regression, with the saving rate as the dependent variable and the predicted permanent income quintile dummies and age dummies as the independent variables. We follow this procedure to allow for nonlinearities in the saving rate–income relationship. We construct standard errors for the estimated saving rates by bootstrapping the entire two-step process. Separately, we also use fitted permanent income, instead of fitted quintiles, as the independent variable in the second stage, both to summarize the relationship between the variables and to provide a simple test of whether the relationship is positive.

IV. Data

Using the CEX, the SCF, and the PSID not only allows for different measures of saving but also ensures that our conclusions are not unduly influenced by the idiosyncracies of a single data source. Most of our results are based on the saving patterns of working-age households—those between the ages of 30 and 59 (as of the midpoint of their participation in each sample), with younger households excluded because they are more likely to be in transitional stages. Examination of the saving behavior of older households is complicated by the noncomparability of households that are on the verge of retirement and those that are beyond retirement. To increase comparability, our CEX and SCF analyses of older households are based on those aged 70–79 (most of whose heads should already be retired). For the PSID, we examine households aged 62 and older, but we focus on the subset of retired households. In this section we describe briefly each of the data sets; further details are in a data appendix available from the authors on request.

A. *Consumer Expenditure Survey*

The CEX has the best available data on total household consumption.¹⁹ In each quarter since 1980, about 5,000 households have been interviewed; a given household remains in the sample for four consecutive quarters and then is rotated out and replaced with a new household. The survey asks for information about consumption, demographics, and income.

We define the saving rate for a CEX household as the difference between consumption and after-tax income divided by after-tax income. Consumption equals total household expenditures *plus* imputed rent for home owners *minus* mortgage payments, expenditures on home

¹⁹ Attanasio (1994) provides a comprehensive analysis of U.S. saving rate data based on the CEX.

capital improvements, life insurance payments, and spending on new and used vehicles. This definition includes expenditures for houses and vehicles as part of saving, in part in order to make the measure of saving in the CEX closer to those in the PSID and SCF. We use Nelson's (1994*a*) reorganization of the CEX, which sums consumption across the four interview quarters for households in the 1982–89 waves. After-tax income equals pretax income for the previous year less taxes for this period, as reported in each household's final interview.²⁰ We deflate both income and consumption with price indexes based in 1994.

We exclude households with income below \$1,000, as well as households with invalid income or missing age data and households that did not participate for all the interviews. We are left with 13,054 households for our working-age analysis and 2,970 households for our older-household analysis.

B. Survey of Consumer Finances

The 1983–89 SCF panel contains information on 1,479 households that were surveyed in 1983 and then again in 1989. The sample has two parts: households from an area-probability sample and households from a special high-income sample selected on the basis of tax data from the Internal Revenue Service. The SCF contains very high quality information about assets and liabilities, as well as limited data on demographic characteristics and income in the calendar year prior to the survey.

The saving rate variable used for the SCF calculations equals the change in real net worth between 1983 and 1989 divided by six times the average of 1982 and 1988 total real household income. Because it spans several years, this variable is likely to be a less noisy measure of average saving than a one-year measure. Net worth is calculated as the value of financial assets (including the cash value of life insurance and the value of defined-contribution pension plans), businesses, real estate, vehicles, and other nonfinancial assets *minus* credit card and other consumer debt, business debt, real estate debt, vehicle debt, and other debt.

We exclude households with 1982 or 1988 income less than \$1,000. We also eliminate households in which the head or spouse changed between 1983 and 1989 because such changes tend to have dramatic and idiosyncratic effects on household net worth. The resulting sample

²⁰ Nelson (1994*b*) warns that the data on household tax payments are quite poor. Inaccurate tax data will bias our results only if the degree of inaccuracy is correlated with our instruments for permanent income.

contains information on 728 households for our working-age analysis and 154 households for our older-household analysis.²¹

C. Panel Study of Income Dynamics

The PSID is the longest-running U.S. panel data set, and, as such, it provides a valuable resource unavailable to researchers in the 1950s and 1960s. The long earnings history for each household helps us disentangle transitory and permanent income shocks, thus facilitating the key issue of identification. Our baseline analysis explores saving between 1984 and 1989; these years correspond to the first two wealth supplements. We also confirm the robustness of our results by examining saving between 1989 and 1994.

Net worth is calculated as the sum of the value of checking and savings accounts, money market funds, certificates of deposit, government savings bonds, Treasury bills, and individual retirement accounts; the net value of stocks, bonds, rights in a trust or estate, cash value of life insurance, valuable collections, and other assets; the value of the main house, net value of other real estate, net value of farm or business, and net value of vehicles; *minus* remaining mortgage principal on the main home and other debts. Net worth does not include either defined-benefit or defined-contribution pension wealth.

We consider several different measures of the saving rate. First, we use the change in real net worth divided by average real after-federal tax money income for the period. Second, we use an “active saving” measure. We start with a measure designed by the PSID staff—the change in wealth *minus* capital gains for housing and financial assets, inheritances and gifts received, and the value of assets less debt brought into the household *plus* the value of assets less debt taken out of the household—and then modify it, following Juster et al. (2001), to correct for inflation and to account for likely reporting error in whether a family moved.²² The saving rate is computed by dividing active saving by five times the average real income measure described above. This active

²¹ Our SCF samples actually contain 2,184 and 462 observations, respectively, because each household’s data are repeated three times with different random draws of imputed variables in order to more accurately represent the variance of the imputed variables. Thus the standard errors in our analysis must be corrected for the presence of replicates. We do so by multiplying them by 1.73—the square root of the number of replicates (three).

²² Capital gain in housing is set equal to the change in the value of the main home during years in which the family did not move and to zero during years in which the family did move less the cost of additions and repairs made to the home. We also tried a third measure of saving, equal to the change in net worth adjusted for assets and debt brought into and out of the household and inheritances. The results were very close to those based on the change in net worth without adjustments.

saving measure should more closely match the traditional income minus consumption measure of saving.

We also create variants of the two PSID saving measures above that include estimates of saving through Social Security and private pensions. Feldstein and Samwick (1992) used then-current (1990) Social Security legislation to determine how much of the payroll tax is reflected in higher marginal benefits at retirement and how much constitutes redistribution. We count the former part as the implicit saving component of the 11.2 cents in total Social Security (Old Age and Survivors Insurance) contributions per dollar of net income. In addition, if a household worker is enrolled in a defined-contribution plan, we count his or her own contribution as saving (we have no data on employer contributions). If a household worker is enrolled in a defined-benefit plan, we include imputations of saving for representative defined-benefit plans, as provided by Gustman and Steinmeier (1989).

We drop households that had active saving greater in absolute value than \$750,000 (1994 dollars) and households that, during relevant years (either 1984–89 or 1989–94), had missing data, a change in head or spouse, or real disposable income less than \$1,000. For the regressions that include lagged or future earnings, we drop households for which there was a change in head or spouse during the relevant years.

D. Summary Statistics from the Three Data Sources

Table 2 shows summary measures of saving and income from the CEX, SCF, and PSID. All saving rates are given on an annual basis, and all income figures are given in 1994 dollars. To avoid undue influence from extreme values of the saving rate when income is close to zero, the “average” saving rates are calculated as average saving for the group divided by average income for the group.

The PSID “active” saving rates are generally the lowest in the table. By contrast, the estimates from the CEX—where saving is also based on the “active” concept—are among the highest. The high levels of CEX saving have been noted by previous authors (e.g., Bosworth, Burtless, and Sabelhaus 1991) and probably reflect measurement error: both income and consumption are understated by respondents, but consumption is thought to be understated by a greater amount, lending an upward bias to saving (Branch 1994). We also calculate the saving rate averaged over the entire sample in each data set, including younger and older respondents, to correspond most closely to an aggregate rate of saving. For comparison, the average national income and product

TABLE 2
SUMMARY SAVING AND INCOME MEASURES

	CEX	SCF	PSID			
	Y-C (1)	Δ Wealth (2)	Δ Wealth (3)	Δ Wealth + Pension (4)	Active+ Pension (6)	
Ages 30-39						
Median saving rate	.27	.07	.05	.13	.05	.14
Average saving rate	.30	.04	.17	.25	.12	.20
Median income	33,807	39,566	36,438	38,670	36,438	38,670
Ages 40-49						
Median saving rate	.26	.10	.03	.15	.05	.16
Average saving rate	.29	.32	.19	.28	.14	.23
Median income	38,810	38,425	45,626	49,586	45,626	49,586
Ages 50-59						
Median saving rate	.26	.06	.03	.16	.02	.16
Average saving rate	.30	.25	.31	.41	.07	.19
Median income	34,087	38,186	33,478	36,828	33,478	36,828
Aggregate (All Ages)						
Average saving rate	.25	.22	.21	NA*	.11	NA*
Aggregate (NIPA)						
Saving rate for period covered by data set	.09 (1982-89)	.08 (1983-89)		.08 (1984-89)		

NOTE.—For each column, the top heading indicates the data set used and the second heading indicates the saving measure used. Measures of income: CEX: disposable income; SCF: pretax income; PSID Δ wealth and active: disposable income; PSID " Δ wealth+pension" and "active+pension": sum of disposable income and employer contributions to Social Security and pensions. All income data are expressed in 1994 dollars. Median saving rate equals median of the ratio of saving to income. Average saving rate equals average saving divided by average income. The sample is the same as the one used in table 3.

* Estimates of pension saving not available for the elderly.

accounts (NIPA) saving rate is shown in the final row; conceptually, this rate is closest to the average "active" saving rate.²³

²³ Note, though, that our active saving measures include purchases of motor vehicles, which should boost them relative to the NIPA concept.

TABLE 3
 MEDIAN REGRESSIONS OF SAVING RATE ON CURRENT INCOME

	CEX	SCF	PSID (1984–89)			
	Y–C (1)	Δ Wealth (2)	Δ Wealth (3)	Δ Wealth+ Pension (4)	Active (5)	Active+ Pension (6)
Quintile 1	–.227 (.017)	.014 (.019)	.000 (.002)	.068 (.007)	.010 (.006)	.090 (.008)
Quintile 2	.150* (.008)	.090* (.029)	.015* (.007)	.110* (.010)	.033* (.008)	.135* (.009)
Quintile 3	.269* (.007)	.111 (.032)	.055* (.009)	.162* (.011)	.063* (.007)	.172* (.011)
Quintile 4	.346* (.006)	.173 (.027)	.077 (.012)	.195* (.011)	.067 (.009)	.192* (.010)
Quintile 5	.455* (.006)	.236 (.040)	.185* (.018)	.307* (.016)	.137* (.011)	.244* (.011)
Top 5%	NA	.372 (.098)	NA	NA	NA	NA
Top 1%	NA	.512 (.111)	NA	NA	NA	NA
Ages 30–39	.006 (.006)	–.041 (.028)	.000 (.004)	–.012 (.008)	–.005 (.006)	–.025 (.008)
Ages 50–59	–.001 (.007)	–.012 (.027)	.002 (.005)	.017 (.016)	–.010 (.006)	.005 (.011)
Pseudo R^2	.142	.047	.028	.045	.030	.044
Sample size	13,054	728	2,854	2,854	2,854	2,854
Coefficient on income/ 10^4	.068 (.002)	.018 (.002)	.022 (.002)	.028 (.002)	.015 (.001)	.019 (.001)

NOTE.—For each column, the top heading indicates the data set used and the second heading indicates the saving measure used. Bootstrapped standard errors are shown in parentheses. The SCF and PSID quintiles are weighted; all regressions are unweighted. Definitions of income: CEX: current income; SCF: average of 1982, 1988 income; PSID: average of 1984–88 income. The PSID active saving and wealth data cover 1984–89.

* The coefficient is significantly greater than that for the previous quintile, on the basis of a one-sided 5 percent test.

V. Empirical Results

A. Saving Rates and Current Income

We begin our empirical inquiry by documenting the well-accepted fact that saving rates increase with current income. Table 3 summarizes how the saving rate varies with respect to the current income quintile for households between the ages of 30 and 59.²⁴ We estimate median regressions, with the saving rate as the dependent variable and dummies for income quintiles and age categories as independent variables. In each case, we suppress the constant term and include dummies for all

²⁴ Income quintiles were calculated (on a weighted basis for the SCF and PSID) for each 10-year age group separately to ensure comparability across data sets and within the U.S. population. We did not use population weights in the regression analysis because the SCF weights—especially those for the top of the income distribution—ranged by orders of magnitude, causing considerable instability in the estimated coefficients. For example, just three of the 107 households in the top 1 percent of the income distribution accounted for 38 percent of the total population weights of these 107 households.

five income quintiles and the 30–39 and 50–59 age groups so that the estimated coefficient for a given income quintile corresponds to the saving rate for households in that quintile with heads between 40 and 49 years old. (Regressions that include interaction terms between age and income variables are similar.) Bootstrapped standard errors for the coefficients, based on 500 replications, are shown in parentheses.

Column 1 of table 3 shows that the saving rate increases dramatically with measured current income in the CEX. Among households with heads between 40 and 49, median saving rates range from –23 percent in the lowest income quintile to 46 percent in the highest. We also calculate (but do not report) bootstrapped standard errors for the difference in the saving rate of quintiles i and $i - 1$ and use an asterisk to indicate a statistically significant difference, based on a 95 percent confidence level and a one-sided test. All the differences in this column are statistically significant. We also report the coefficient from a regression of saving rates on the level of income. This coefficient suggests that a \$10,000 increase in income is associated with a seven-percentage-point increase in the saving rate. Consistent with previous research based on the CEX, we estimate an extremely low saving rate for the lowest income quintile; this reflects bias from measurement error in income and possibly transitory income, since households in this quintile could not sustain such a high rate of dissaving for very long (see Sabelhaus 1993).

Column 2 shows results from similar regressions using SCF data, including saving rate estimates for households in the ninety-fifth and ninety-ninth percentiles of the income distribution.²⁵ The slope of the relationship between the saving rate and measured current income is smaller than in the CEX. This result is not surprising: the change-in-wealth saving rate is not subject to the upward bias associated with measurement error in income, and many transitory movements in income likely wash out over the six-year period covered by the SCF panel. Nevertheless, we see the estimated (annualized) median saving rate rising significantly from 1 percent for households in the bottom quintile to 24 percent for households in the top quintile.²⁶ Saving rates are even larger for the richest households: 37 percent for those in the top 5 percent of the income distribution and 51 percent for those in the top

²⁵ We are able to estimate fairly precise saving rates for households in the highest part of the income distribution because the SCF disproportionately samples high-income households: out of a total of 728 households in the age 30–59 sample, 193 have income above the ninety-fifth percentile and 86 have income above the ninety-ninth percentile.

²⁶ Because median saving rates within quintiles are calculated using population weights, the saving rates of the fifth quintile in the SCF (in this and subsequent tables) will be biased upward because of the very high income households in this group. Since median income in this quintile will be higher as well, however, the slope of the line in fig. 2 below will convey the true relationship between saving rates and income.

1 percent.²⁷ The coefficient from the linear regression suggests that the saving rate rises two percentage points for each additional \$10,000 in income.²⁸

The remaining columns of the table show the relationship between saving and income in the PSID. As in the SCF, the several-year period over which saving is measured reduces the importance of transitory income. The (annualized) change-in-wealth saving rate—shown in column 3—is similar to that in the SCF for the lowest-income households and rises as income moves up. For the highest income quintile, the estimated saving rate of 19 percent is somewhat lower than the comparable figure from the SCF. Still, the coefficient from the linear regression is similar to that from the SCF.

Adding estimates of saving through Social Security and private pensions to the PSID change-in-wealth saving rate (col. 4) both raises the levels and somewhat steepens the trajectory of saving rates across the income distribution. How can this be, given the higher Social Security rates of return and replacement rates among households with lower earnings? The main answer is that saving through private pensions increases disproportionately with income, more than offsetting the decline in Social Security saving rates.²⁹ The column 4 results suggest that the low observed rates of financial saving among lower-income households cannot be explained by higher implicit Social Security and pension wealth accumulation.

Columns 5 and 6 of table 3 show the results for the active PSID saving measures. For the first four income quintiles, active saving rates are very similar to change-in-wealth saving rates. But for the highest income group, the change-in-wealth saving rates are five to six percentage points higher than the active ones, owing to “passive” capital gains.

Looking at averages rather than medians does not alter the basic

²⁷ The top quintile includes the top 5 percent, and the top 5 percent includes the top 1 percent.

²⁸ Because of nonlinearities at very high levels of income, in this SCF regression (but not in the quintile regressions) we exclude households with income in excess of \$500,000. In the corresponding subsequent SCF two-stage regressions, we exclude households from the second-stage regression if their fitted values of income exceed \$500,000.

²⁹ For example, among households 40–49, median Social Security saving as a percentage of disposable income declines from 6.5 percent for the lowest income quintile to 3.9 percent for the highest income quintile. However, Social Security plus pension saving rises from 7.6 percent in the lowest income quintile to 11.1 percent in the highest income quintile. One reason why Social Security saving is not more strongly progressive is that we measure saving relative to disposable income rather than earnings. Because transfer payments such as Aid to Families with Dependent Children, disability insurance, and unemployment insurance are concentrated among lower-income households, Social Security saving is a smaller fraction of disposable income among these households. Accounting for additional factors such as the correlation between income and life span would presumably attenuate the progressivity of Social Security further (Coronado, Fullerton, and Glass 2000; Gustman and Steinmeier 2001; Liebman 2002).

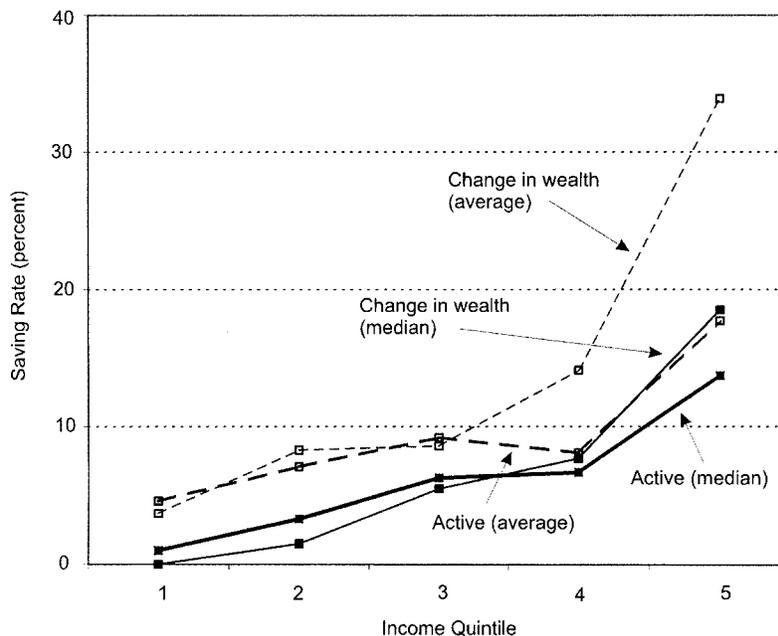


FIG. 1.—Median and average saving rates: active saving and change in wealth, PSID, 1984–89. Estimates of the median saving rate are taken from table 3. Average saving rates were calculated by dividing the average level of saving for the quintile by the average level of income; to improve statistical power, the full sample (ages 30–59) was used for the calculation.

picture formed by table 3. Figure 1 shows average saving rates and median saving rates (from table 3) by income quintile for both the change-in-wealth measure of saving and active saving. As in table 2, average saving rates are defined to be average saving for each quintile divided by average five-year income. Although the mean saving rates are somewhat higher than the medians, the patterns are generally similar except in the top quintile, where the mean change-in-wealth saving rate jumps to 34 percent of income.

B. Saving Rates and Permanent Income

We now turn to the relationship between saving rates and permanent income, using the two-stage procedure described earlier. We first focus on consumption as an instrument. Recall that transitory consumption will bias the estimated slope toward a negative number in all three data sets, as will measurement error in the case of the CEX.

Column 1 of table 4 shows results from the CEX. The estimated

TABLE 4
 MEDIAN INSTRUMENTAL VARIABLE REGRESSIONS OF SAVING RATE ON INCOME USING CONSUMPTION AS AN INSTRUMENT

	CEX	SCF	PSID			
	$Y-C$	Δ Wealth	Δ Wealth		Δ Wealth + Pension	Active
	Nonauto Consumption (1)	Vehicles (2)	Food Consumption (3)	Weighted Consumption (4)	Weighted Consumption (5)	Weighted Consumption (6)
Quintile 1	.211 (.010)	.028 (.030)	.000 (.005)	.000 (.004)	.077 (.010)	.011 (.006)
Quintile 2	.288* (.008)	.140* (.044)	.021* (.008)	.029* (.008)	.124* (.011)	.045* (.008)
Quintile 3	.278 (.008)	.134 (.039)	.033 (.010)	.044 (.010)	.159* (.013)	.053 (.007)
Quintile 4	.283 (.007)	.173 (.027)	.057* (.010)	.078* (.013)	.191* (.016)	.076* (.008)
Quintile 5	.246 (.007)	.286* (.049)	.139* (.018)	.260* (.032)	.344* (.027)	.121* (.013)
Top 5%	NA	.505 (.113)	NA	NA	NA	NA
Top 1%	NA	.356 (.158)	NA	NA	NA	NA
Ages 30–39	.007 (.007)	–.051 (.029)	.005 (.007)	.000 (.005)	–.012 (.011)	–.003 (.006)
Ages 50–59	–.001 (.009)	–.016 (.034)	.006 (.008)	.002 (.007)	.042 (.018)	–.011 (.007)
Pseudo R^2	.003	.031	.013	.034	.040	.022
Sample size	13,054	728	2,793	2,793	2,793	2,793
Coefficient on income/ 10^4	–.003 (.002)	.040 (.026)	.022 (.003)	.035 (.004)	.041 (.004)	.018 (.002)

NOTE.—For each column, the top heading indicates the data set used, the second heading indicates the saving measure used, and the third heading indicates the instrument used. Bootstrapped standard errors are shown in parentheses. The SCF and PSID quintiles are weighted; all regressions are unweighted.

* The coefficient is significantly greater than that for the previous quintile, on the basis of a one-sided 5 percent test.

median saving rate rises from the predicted first to second quintile but then remains fairly flat. One interpretation is that the results favor the Friedman proportionality hypothesis, but a more likely explanation is that the negative bias associated with transitory consumption and consumption measurement error is approximately offset by a positive correlation between saving rates and permanent income.

We next consider data from the SCF and PSID, where saving is derived from the change in wealth and is thus likely uncorrelated with consumption measurement error. The SCF does not contain direct consumption flow measures, but it does include the reported value of owned vehicles. As shown in column 2, the results when the average of 1983 and 1989 vehicle values is used as an instrument are surprisingly similar to those in table 3, with saving rates rising from 3 percent in the lowest quintile to 29 percent in the top quintile. Saving rates in the top 5 percent are even higher. The estimated linear impact of income on saving rates is roughly four percentage points per \$10,000 in income but is not significant at the 5 percent level.

Column 3 of table 4 shows that when PSID food consumption is used as an instrument, the estimated change-in-wealth saving rate rises consistently with income.³⁰ Indeed, the step-up in the saving rate is significant for every quintile but the third. Columns 4–6 of table 4 use a more comprehensive measure of consumption from the PSID: a weighted average of food at home, food away from home, rental payments, and imputed housing flows, with weights derived from the CEX (Hamermesh 1984; Skinner 1987).³¹ The estimated gradients of the saving rate with respect to income are similar to (and in some cases slightly larger than) those in table 3.³²

Our next approach instruments current income with lagged and future earnings. For the CEX, we have no data on lagged or future earnings. For the SCF, we redefine current income (for the purposes of putting households into quintiles) as 1988 income and instrument with 1982 income. Column 1 of table 5 shows that this procedure yields a strong relationship between predicted income and saving rates. Only

³⁰ In the first stage, we predict average current disposable income (1984–88) using food consumption for each year from 1984 to 1987. (The food questions were temporarily suspended in 1988, so that we have no consumption measures for that year.) See Zeldes (1989) for further details on the construction of the food consumption variable.

³¹ More specifically, we calculate the measure for each year from 1984 to 1987, using weights from Bernheim, Skinner, and Weinberg (2001), so that

$$C^{\text{weighted}} = 1.930(\text{food at home}) + 2.928(\text{food away from home}) \\ + 1.828(\text{rental payments if renter}) + .1374(\text{value of house if home owner}).$$

³² To save space, we do not report active plus pension and Social Security saving estimates in this and subsequent tables. Full sets of regression results are available on request.

TABLE 5
 MEDIAN INSTRUMENTAL VARIABLE REGRESSIONS OF SAVING RATE ON INCOME USING LAGGED AND/OR FUTURE EARNINGS AS
 INSTRUMENTS

	SCF		PSID				
	Lagged Income		Lagged Earnings			Future Earnings	
	Δ Wealth (1)	Δ Wealth (2)	Δ Wealth+ Pension (3)	Active (4)	Δ Wealth (5)	Δ Wealth+ Pension (6)	Active (7)
Quintile 1	.022 (.025)	.000 (.004)	.065 (.010)	.003 (.005)	.000 (.003)	.063 (.008)	.003 (.005)
Quintile 2	.094* (.027)	.024* (.012)	.136* (.019)	.024* (.009)	.025* (.008)	.140* (.013)	.028* (.009)
Quintile 3	.106 (.036)	.059 (.019)	.174 (.020)	.052* (.011)	.057* (.010)	.170* (.011)	.064* (.008)
Quintile 4	.167 (.028)	.081 (.019)	.195 (.019)	.073 (.013)	.072 (.014)	.201* (.015)	.062 (.009)
Quintile 5	.246 (.035)	.115 (.035)	.239 (.030)	.080 (.016)	.170* (.017)	.289* (.020)	.119* (.010)
Ages 30–39	-.057 (.026)	.000 (.008)	-.012 (.016)	.006 (.009)	.000 (.004)	-.020 (.010)	-.001 (.006)
Ages 50–59	-.016 (.027)	.001 (.007)	.029 (.018)	-.003 (.006)	.000 (.004)	.010 (.015)	-.003 (.006)
Pseudo R^2	.041	.013	.022	.017	.025	.041	.030
Sample size	728	1,359	1,359	1,359	2,471	2,471	2,471
Coefficient on income/ 10^4	.020 (.005)	.014 (.002)	.022 (.003)	.011 (.002)	.026 (.002)	.031 (.003)	.017 (.002)

NOTE.—For each column, the top heading indicates the data set used, the second heading indicates the saving measure used, and the third heading indicates the instrument used. Bootstrapped standard errors are shown in parentheses. The SCF and PSID quintiles are weighted; all regressions are unweighted. The SCF results use 1988 income as current income and 1982 income as lagged income. The PSID results use 1974–78 for lagged earnings and 1989–91 for future earnings. The SCF coefficients for the top 5 percent and top 1 percent are .397 (.115) and .455 (.088), respectively.

* The coefficient is significantly greater than that for the previous quintile, on the basis of a one-sided 5 percent test.

TABLE 6
 MEDIAN REGRESSIONS OF SAVING RATE ON EDUCATION

	CEX	SCF	PSID		
	Y-C (1)	Δ Wealth (2)	Δ Wealth (3)	Δ Wealth+ Pension (4)	Active (5)
No high school diploma	.155 (.009)	.057 (.043)	.000 (.003)	.090 (.009)	.020 (.006)
High school diploma	.284* (.006)	.131 (.031)	.039* (.006)	.148* (.009)	.052* (.007)
College degree+	.342* (.007)	.323* (.027)	.123* (.015)	.236* (.014)	.102* (.010)
Ages 30-39	-.004 (.007)	-.063 (.033)	.002 (.005)	-.021 (.009)	-.006 (.007)
Ages 50-59	.017 (.009)	.021 (.046)	.008 (.007)	.033 (.013)	-.017 (.008)
Pseudo R^2	.017	.025	.014	.020	.014
Sample size	13,054	728	2,840	2,840	2,840
Coefficient on income/ 10^4	.060 (.003)	.009 (.002)	.028 (.003)	.033 (.003)	.021 (.002)

NOTE.—For each column, the top heading indicates the data set used and the second heading indicates the saving measure used. Bootstrapped standard errors are shown in parentheses. The regressions are unweighted. Definitions of income: CEX: current income; SCF: average of 1982, 1988 income; PSID: average of 1984-88 income.

* The coefficient is significantly greater than that for the next lower education, on the basis of a one-sided 5 percent test.

one of the differences between quintile estimates is statistically significant, but the estimate from the linear equation (a two-percentage-point increase for each \$10,000 in predicted income) is highly significant.

For the PSID, we use as instruments labor earnings of the head and wife (combined) for each year from 1974 to 1978—10 years before the period over which saving is measured.³³ Columns 2, 3, and 4 of table 5 show the results of this approach for the three PSID saving measures. In all cases, saving rates rise with predicted permanent income. The magnitudes of the differences are in fact quite close to those from the uninstrumented results in table 3, suggesting that the simple five-year average of current income eliminates much of the effects of transitory income. Columns 5-7 of table 5 show that when future earnings (1989-91) are used as instruments, we again see saving rates increasing with predicted income. This is true whether one looks at the quintile coefficients (ranging, for the change-in-wealth plus pension saving measure, from 6 percent to 29 percent) or the coefficient from the regression on predicted income.

In table 6, we next turn to education as an instrument—a proxy for permanent income that is generally fixed over the life cycle. For the

³³ We have earnings information back to 1967, but when we condition on earnings in more recent years, the earlier earnings added little or no predictive power for income in 1984-88.

top of the table, we do not use a two-stage procedure but simply report the “reduced-form” estimates of saving rate by education group. We suppress the constant and include dummies for two age groups (30–39 and 50–59) and all education groups; the excluded age group is 40–49. At the bottom of the table, we report the coefficient on predicted income from a two-stage regression.

As shown in column 1 of table 6, estimated median saving rates in the CEX range from 16 percent for high school dropouts to 34 percent for college graduates, with the differences statistically significant. The range for the SCF (col. 2) is 6–32 percent. The PSID also shows a positive correlation, with the range depending on the saving rate measure; for example, the change-in-wealth saving rate with imputed Social Security and pension income ranges from 9 percent for high school dropouts to 24 percent for college graduates. Almost all the differences between education groups are statistically significant. The coefficients on predicted income for the PSID runs indicate that the saving rate rises by between two and three percentage points for each \$10,000 increase in income.

We summarize the results presented so far in figure 2. For each fitted income quintile or education group, we plot the median saving rate against median income.³⁴ The results are striking. While the CEX (fig. 2*a*) shows a much flatter line for the consumption-based regression than for the others, the SCF (fig. 2*b*) and PSID (figs. 2*c* and *d*) show upward-sloping lines that are essentially the same across all choices of instruments and when no instrumenting is done. In sum, the results presented thus far strongly suggest that saving rates rise with lifetime income among working-age households.

C. *Is It Just High-Income Entrepreneurs Who Save More?*

Quadrini (1999, 2000), Gentry and Hubbard (2000), and Hurst and Lusardi (2004 [in this issue]) have emphasized the importance of entrepreneurs in wealth accumulation, particularly at the top of the income distribution. To what extent are our results driven by the saving behavior of entrepreneurs? To examine this question, we restrict our SCF sample to nonentrepreneurs using the Gentry-Hubbard definition: households for which the value of businesses in which they have an active management role is less than \$5,000. We continue to estimate a strong positive correlation between saving rates and income. For example, the specification corresponding to column 2 of table 3 yields

³⁴ The median saving rate numbers plotted are the coefficients from the regressions in tables 3–6. The median income numbers are coefficients from median regressions of current income on the same variables included in the corresponding saving rate regression.

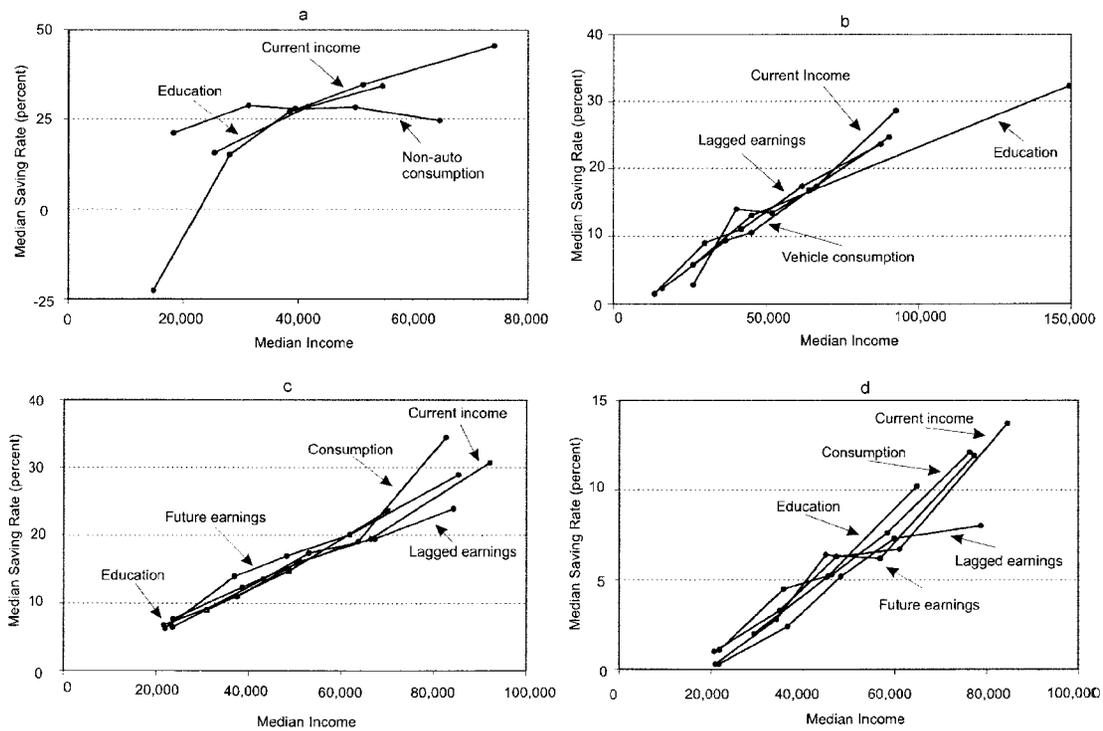


FIG. 2.—Summary of regressions: *a*, CEX saving rates; *b*, SCF saving rates; *c*, PSID saving rates, 1984–89, change in wealth plus pension; *d*, PSID saving rates, 1984–89, active.

estimated median saving rates that range from zero for the first quintile to 22 percent for the fifth quintile to 49 percent for the top 5 percent of the sample. We also find a significant positive relationship for non-entrepreneurs defined as those not self-employed.

D. Saving Rates and Children

The dynastic smoothing model implies that higher-income households smooth consumption over generations by leaving bequests to their children (e.g., Becker and Tomes 1986). One would expect that, on average, households with children would save more for bequests than childless households, with income held constant. Of course, the null hypothesis of equal saving rates is overly simplistic, given the importance of inter vivos transfers to children that may appear either as saving (for college) or as expenditures and the presence of philanthropic bequest motives among childless households. Nonetheless, we find no evidence that the saving rates of households with children are higher or have a higher gradient with respect to income than those of households without children (see also Hurd 1987; Altonji and Villanueva 2003).³⁵ This suggests that dynastic consumption smoothing through bequests alone is not the primary explanation for why high-income households save more.

E. Saving Rates among Older Households

In this subsection, we consider how the relationship between saving and permanent income changes at older ages. Estimating the relationship for older households is difficult because observable measures of income may not be good indicators of lifetime income, particularly if we pool together people who are still working and those who have already retired. For the CEX and SCF, we attempt to avoid at least the noncomparability issue by restricting the sample to just the age group 70–79, where fewer than 15 percent of households have heads or spouses who are still working more than 20 hours a week. For the PSID, in table 7, we use a larger sample of retired households aged 62 and over, and for columns 4 and 6 we use the sum of the averages of real after-tax earnings of the head and of the spouse (if present) during all their available working years (age 62 and below). For the estimates based on this average, we restrict the sample to those with earnings data going back at least 10 years for the head and five years for the spouse. We examine

³⁵ Households with children were determined on the basis of whether either head or spouse had been a biological or adoptive parent, even if the child was not currently living in the household. The specification and the sample (working-age households) were the same as in table 3. If anything, the saving rates of households without children are higher and have a slightly higher gradient with respect to income.

TABLE 7
 MEDIAN REGRESSIONS OF SAVING RATE FOR OLDER HOUSEHOLDS BY INCOME QUINTILE

	CEX (Ages 70–79)		PSID (Ages 62+, Retired)			
	$Y-C$	Δ Wealth	Δ Wealth		Active	
	Current Income (1)	Pension/Social Security Income (2)	Current Income (3)	10-Year+ Earnings (4)	Current Income (5)	10-Year+ Earnings (6)
Quintile 1	-.485 (.039)	.012 (.098)	-.005 (.010)	.000 (.015)	.000 (.004)	-.004 (.009)
Quintile 2	-.341* (.032)	.601* (.238)	-.005 (.010)	-.022 (.024)	.000 (.005)	-.010 (.016)
Quintile 3	-.136* (.018)	-.019 (.109)	-.027 (.036)	-.013 (.042)	-.013 (.018)	.004 (.029)
Quintile 4	.049* (.018)	.158 (.155)	.001 (.029)	.053 (.040)	.006 (.017)	.019 (.018)
Quintile 5	.319* (.013)	-.090 (.111)	.144* (.060)	.130 (.120)	.027 (.027)	.023 (.068)
Pseudo R^2	.116	.032	.004	.006	.002	.002
Sample size	2,970	154	636	262	636	262
Coefficient on income/ 10^4	.144 (.007)	.027 (.019)	.012 (.012)	.007 (.009)	.004 (.004)	.001 (.004)

NOTE.—For each column, the top heading indicates the data set used, the second heading indicates the saving measure used, and the third heading indicates the income measure used to form quintiles. Bootstrapped standard errors are shown in parentheses. In the SCF regression, dummy variables for the top 5 percent and top 1 percent are included but not reported in the table. In the PSID regressions, age dummy variables (62–69 and 80+) are included but not reported in the table. Denominator for saving rate: CEX: current disposable income; SCF: average of 1982, 1988 total income; PSID: average of 1984–88 disposable income. Ten-year+ earnings for PSID equals the average of (present value of) all available lagged earnings for the head during years the head was aged 62 or younger, plus a similar average for the spouse. The sample is restricted to households with earnings histories for the head and spouse of 10 years and five years, respectively.

* The coefficient is significantly greater than that for the next lower quintile, on the basis of a one-sided 5 percent test.

TABLE 8
 MEDIAN REGRESSIONS OF SAVING RATE FOR OLDER HOUSEHOLDS BY EDUCATION

	CEX (Ages 70–79)	SCF (Ages 70–79)	PSID (Age 62+, Retired)	
	<i>Y</i> – <i>C</i> (1)	Δ Wealth (2)	Δ Wealth (3)	Active (4)
No high school	–.069 (.014)	.016 (.035)	–.003 (.006)	.000 (.003)
High school graduate	.027* (.032)	.134 (.170)	.005 (.029)	.003 (.014)
College graduate+	.128 (.047)	.182 (.160)	.020 (.111)	–.031 (.039)
Pseudo R^2	.092	.005	.000	.001
Sample size	2,970	154	630	630

NOTE.—For each column, the top heading indicates the data set used and the second heading indicates the saving measure used. Bootstrapped standard errors are shown in parentheses. In the PSID regressions, age dummy variables (62–69 and 80+) are included but not reported in the table. Denominator for saving rate: CEX: current disposable income; SCF: average of 1982, 1988 total income; PSID: average of 1984–88 disposable income.

* The coefficient is significantly greater than that for the next lower education group, on the basis of a one-sided 5 percent test.

the change-in-wealth and active measures of saving, but we do not attempt to impute the “dissaving” that occurs by the drawing down of Social Security and pension wealth.

Table 7 presents median saving rates by income quintile. Not surprisingly, given the apparent bias from measurement error and transitory income, there is a strong correlation between saving rates and income in the CEX (col. 1). However, there is little or no correlation in the SCF (col. 2) or, for most measures of saving, in the PSID (cols. 3–6). One exception is the change-in-wealth saving rate in the PSID, where the estimate for the fifth quintile is at least 14 percentage points higher than those for any of the lower quintiles and statistically different both from zero and from the coefficient for the fourth quintile.³⁶ Table 8 shows corresponding results by educational attainment. The CEX results show a positive association between education and saving among older groups. The SCF is suggestive of a positive education–saving rate gradient, although the sample size is too small to establish statistical significance. However, results from the PSID suggest no relationship. Taken together, the results certainly provide no evidence that older high–lifetime income households dissave at a faster rate than older low–lifetime income households; if anything, they provide some weak evidence to the contrary.

³⁶ When working households aged 62 or older were added to the sample, the results showed a stronger and more consistently positive relationship between saving rates and income.

F. Do the Results Hold Up for a Later Time Period?

Using the third wealth supplement of the PSID (holdings as of 1994), we constructed estimates of both change in wealth and active saving between 1989 and 1994 and reran most of the estimations reported in the tables above (these results are available on request). The estimated median active saving rate for the later period rose from 0.7 percent for the lowest current income quintile to 11.2 percent for the highest quintile, similar to the estimates of 1.0 percent and 13.7 percent for the earlier period. The estimated medians for the change-in-wealth saving rate were also similar to those of the earlier period for the bottom four quintiles, although the estimate for the top income decile, 12.5 percent, was below the estimate from the earlier time period, 18.5 percent. Results were generally similar across periods when our sets of instruments were used, except for a downturn in the later period in active (but not change-in-wealth) saving when weighted consumption was used as an instrument.

G. The Marginal Propensity to Save

For some questions, the relevant relationship between income and saving is not that pertaining to the *average* propensity to save but instead that pertaining to the *marginal* propensity to save (MPS) (see Blinder 1975). We consider two approaches to estimating the MPS. The simplest approach, based on a cross section, is to use the variation in average saving rates across income classes to trace out a “marginal” saving schedule. To do this, we divide the change in the level of saving across income groups by the change in income across these groups. The dots in figure 3 correspond to the predicted (change-in-wealth) saving from the cross-sectional regressions, and the slope of each line connecting the dots is the estimated MPS.³⁷ These MPS estimates range from 3 cents per dollar of income between quintiles 1 and 2 up to 43 cents between quintiles 4 and 5.

These estimates may be problematic, however, if there is a third factor that varies across the population and is driving differences in both income and saving rates. For example, if households with high rates of time preference tend to have both a lower permanent income (perhaps due to lack of investment in education) and a low saving rate (e.g., Evans and Montgomery 1995), then exogenously increasing the permanent income of those in the low-income group would not cause them to save at the rate observed in the high-income group (Mayer 1972).

Our second approach to estimating the MPS relates the change in

³⁷ The implied level of saving is equal to median income multiplied by the corresponding coefficient from col. 3 in table 3.

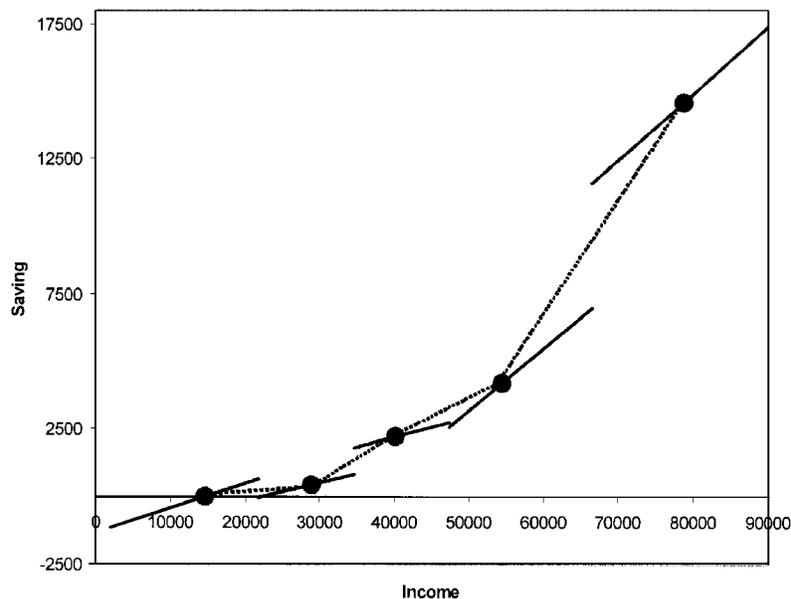


FIG. 3.—Cross-section and time-series estimates of the marginal propensity to save. (1) Estimates are based on the change-in-wealth measure of saving. All saving and disposable income measures are given in 1994 dollars. (2) Each dot represents the predicted median level of saving for an income quintile. Its horizontal position is equal to the median household income for the quintile. Its vertical position is equal to the product of median income and the median saving rate for that quintile, estimated in col. 3 of table 3. (3) The connecting dotted lines therefore trace out an implicit saving function, where the slope of each line is the implicit MPS for that income bracket. The slopes of the implicit MPS segments are 3.0 percent (quintiles 1 and 2), 15.8 percent (quintiles 2 and 3), 13.9 percent (quintiles 3 and 4), and 42.9 percent (quintiles 4 and 5). (4) The slopes of the solid lines are equal to the quintile-specific MPS estimated directly from the time-series changes in saving rates (col. 1 of table 9 below). For ease of comparison, these solid lines are drawn through the corresponding median income and saving dots described above.

saving across time for a given household to the change in its income. We use the PSID and define the change in saving as the difference between 1989–94 saving and 1984–89 saving. We define the change in income as the difference between average income for 1989–93 and average income for 1984–88 and estimate the median regression as

$$\Delta S_{iq} = a_q + b_q \Delta Y_{iq} + \zeta_3 \text{age}_3 + \zeta_5 \text{age}_5 + e_{qi} \quad (4)$$

where ΔS_{iq} equals the change in saving for household i in age-specific income quintile q (in 1984–89), a_q is the quintile-specific trend term over this period, b_q is the quintile-specific coefficient corresponding (under appropriate assumptions) to the MPS, ΔY_{iq} is the household's change in average income, and the remaining terms reflect age dummy

TABLE 9
 MEDIAN REGRESSION ESTIMATES OF THE MARGINAL
 PROPENSITY TO SAVE

	PSID	
	Δ Wealth (1)	Active Saving (2)
Quintile 1	.089 (.033)	.080 (.040)
Quintile 2	.068 (.044)	.035 (.043)
Quintile 3	.075 (.057)	.088 (.087)
Quintile 4	.227 (.086)	.183 (.083)
Quintile 5	.248 (.111)	.039 (.081)
Pseudo R^2	.006	.004
Sample size	2,907	2,068

NOTE.—For each column, the top heading indicates the data set used and the second heading indicates the saving measure used. The dependent variable is the change in household saving between 1989–94 and 1984–89. Included as independent variables are the changes in household income between 1989–93 and 1984–88, interacted with 1984–88 income quintile dummy variables. Also included (coefficients are not reported in the table) are age dummies (30–39 and 50–59, with ages 40–49 the excluded value) and a set of dummy variables for the initial quintile of income in 1984–88.

variables (age_3 for ages 30–39 and age_5 for ages 50–59 in 1987) multiplied by their coefficients and the error term. Note that the validity of this procedure depends on the assumption that the changes in five-year average income reflect permanent rather than transitory changes.

Table 9 presents estimates of equation (4) for both active saving and change-in-wealth saving from the PSID. For both types of saving, the estimated coefficients b_q are modest and not statistically significant for quintiles 2 and 3. For the change-in-wealth definition of saving, the coefficients for quintiles 4 and 5 are 0.227 and 0.248, respectively, and are significantly different from zero. For the active saving measure, the quintile 4 coefficient, 0.183, is large and significant, but surprisingly, the predicted estimate of the quintile 5 coefficient is neither.³⁸ Figure 3 displays these direct estimates; the slopes of the solid lines equal the quintile-specific estimated MPS, with each line drawn for convenience

³⁸ The substantial difference in the two estimates for quintile 5 could reflect statistical noise, since the 95 percent confidence intervals overlap. Alternatively, if households in the top quintile had experienced both rapid wage growth and rapid appreciation of house prices (because of regional growth), their estimated MPS based on the change-in-wealth saving measure would be too high (because it would include the appreciation in asset values) and their estimated MPS based on the active saving measure would be too low (because it would include the consumption response both to the income gain and to the capital gain; see, e.g., Juster et al. [2001]). Note also that the results were somewhat sensitive to variations (not reported) in the specification of the equation.

through the corresponding median income and saving dots described above. While the two MPS estimates do not match up exactly, suggesting heterogeneity in saving behavior across income groups in saving behavior, both approaches suggest that the MPS (and the marginal propensity to consume) differs across income groups. This means that the impact of a change in taxes or other fiscal variable on the aggregate economy will depend on its distribution across income groups and thus suggests that a representative agent model cannot fully predict the results of a given fiscal policy (e.g., Stoker 1986; Caselli and Ventura 2000; see also Harris and Laibson 2001).

H. Comparison with Earlier Results on Wealth-Earnings Ratios

Our results are seemingly at odds with several recent studies using the Health and Retirement Survey (HRS) that find that the ratio of wealth to lifetime earnings is highest for households with the lowest lifetime earnings, falls off sharply at slightly higher earnings levels, and is roughly constant thereafter (Gustman and Steinmeier 1997, 1999; Venti and Wise 1999). How can our results be reconciled with theirs? An important difference is that Gustman and Steinmeier's and Venti and Wise's measure of pension and Social Security wealth is more accurate than our measure of pension and Social Security saving, because they use information from actual Social Security records to compute Social Security wealth and because the HRS has more comprehensive data on defined-benefit pension coverage. This does not appear to be the explanation, however. Gustman and Steinmeier find that Social Security and pension wealth as a percentage of lifetime earnings is roughly flat between the tenth and ninetieth percentiles, which is consistent with our finding above in table 3. When they examine a measure of wealth that excludes Social Security and pension wealth (leaving financial, business, and housing wealth), the resulting ratio of wealth to lifetime earnings continues to be either declining or flat across the income distribution (see Gustman and Steinmeier 1997, tables 9 [p. 75], 12 [p. 78]).

Perhaps the discrepancy can instead be explained by the fact that Gustman and Steinmeier and Venti and Wise focus on wealth stocks whereas we measure saving flows. This also does not appear to be the explanation, however, since in our data sets wealth-income ratios display a pattern similar to that of saving-income ratios. Table 10 shows estimates of wealth-income ratios by income quintile from the PSID and SCF, based on samples that match the ages (51–61) of the HRS households considered by Gustman and Steinmeier and Venti and Wise. For the 1983–89 SCF panel (col. 1), the wealth-income ratio rises from 1.5 in the lowest quintile to 3.4 in the highest quintile. Similar results are found in the 1992 SCF (col. 2) and the PSID (col. 3) using quintiles

TABLE 10
 MEDIAN REGRESSION ESTIMATES OF WEALTH-INCOME RATIOS BY INCOME QUINTILE:
 HOUSEHOLD HEADS AGED 51–61 AT TIME OF SURVEY

	1983–89 SCF PANEL	1992 SCF CROSS SECTION	1989 PSID	
	Average 1982, 1988 Income (1)	1991 Income (2)	Average 1984–89 Income (3)	10 Year+ Average Earnings (4)
Quintile 1	1.50 (.61)	.86 (.62)	.58 (.28)	.37 (.14)
Quintile 2	1.59 (.66)	2.93* (.51)	1.75* (.30)	1.78* (.21)
Quintile 3	1.48 (.70)	3.18 (.36)	2.19 (.30)	2.79* (.26)
Quintile 4	1.88 (.39)	2.82 (.18)	2.57 (.22)	2.76 (.28)
Quintile 5	3.37* (.48)	4.01 (.66)	3.85* (.34)	4.04* (.26)
Pseudo R^2	.031	.001	.056	.085
Sample size	294	684	717	677

NOTE.—Wealth includes housing equity. Wealth in col. 1 is the average of 1983 and 1989 wealth. For each column, the top heading indicates the data set used and the second heading indicates the income measure used both to form quintiles and as the denominator in the wealth-income ratio. The SCF and PSID quintiles are weighted; the SCF regressions are weighted. The 10-year+ earnings for PSID equals the average of (the present value of) all available lagged earnings (back to 1987) for the head while the head was aged 62 or younger, plus a similar average for the spouse. The sample is restricted to households with at least 10 years of earnings history for the head and at least five years for the spouse (even if earnings are zero).

* The coefficient is significantly greater than that for the next lower quintile, on the basis of a one-sided 5 percent test.

computed as we have previously. Another possible explanation for the different findings is the definition of lifetime income; for the most part, we have used proxies that average over fewer years than the Gustman-Steinmeier and Venti-Wise measure. For column 4 of table 10, the variable used to define quintiles and for the denominator of the ratio is the very long average of lagged earnings (stretching back to 1967 when possible), described in subsection *E* above. However, we again see a distinct positive relationship between income and the wealth-income ratio.

What are the remaining possible explanations for why their results might differ from ours? One is measurement error: when permanent income is measured imperfectly, wealth-income ratios for households misclassified to the lowest income groups will tend to be biased upward (because the denominator is too small), whereas wealth-income ratios for households misclassified to the highest income groups will be biased downward (since the denominator is too big). The greater the error in classification is, the larger the negative bias in the estimated relationship

between income and wealth-income ratios.³⁹ Venti and Wise and Gustman and Steinmeier use Social Security administrative records to compute lifetime earnings.⁴⁰ While this will likely result in an accurate measure of lifetime earnings for some types of households, for others, such as households with “off-the-books” earnings, it is likely to be biased downward.⁴¹ Indeed, the Venti-Wise and Gustman-Steinmeier estimates imply a present value of lifetime earnings (in 1992 dollars) for the bottom decile of about \$36,000 and \$86,000, respectively. On an annual (rather than accumulated) basis, these figures imply average total household income of less than \$3,500; moreover, the wealth of these groups *exceeds* their lifetime earnings—an implausible result. These patterns suggest a downward bias in the bottom decile of lifetime earnings. If this mismeasurement carries over to other deciles, it could help explain the flatness in the Venti-Wise/Gustman-Steinmeier estimates.

VI. Discussion

This paper revisits an old question: Do high-lifetime income households save a larger fraction of their income than low-lifetime income households? This question was the topic of heated and largely inconclusive debates in the 1950s and 1960s. We have approached it with three data sources: the Consumer Expenditure Survey, the Survey of Consumer Finances, and the Panel Study of Income Dynamics. For households aged 30–59, we consistently find that higher-lifetime income households save a larger fraction of their income than lower-income households.

³⁹ Of course, the bias could run in the other direction if a household experiences a transitory income shock that induces a percentage change in wealth greater than that of income (and thus moves the wealth-income ratio in the direction of the transitory shock). But we view this latter effect as likely to be limited since it will occur only when a household has a low level of starting wealth and experiences a large shock to transitory income.

⁴⁰ For individuals not covered by Social Security or with earnings above the Social Security earnings cap, Gustman and Steinmeier also use survey questions about earnings at current and past jobs.

⁴¹ Another example is immigrants with few years of covered earnings (Gustman and Steinmeier 2000). We average earnings only over years in which households are in the survey (and do not include any new immigrants after 1968), whereas Gustman-Steinmeier/Venti-Wise average across all years, including pre-immigration years in which Social Security earnings equal zero. Yet another example is divorced spouses who did not work outside the house in earlier years: we average only across years subsequent to divorce whereas the Gustman-Steinmeier/Venti-Wise average would include the zeros from earlier years as well. Finally, for households above the Social Security earnings cap, Gustman and Steinmeier estimate earnings based on recollective questions about earnings at past jobs. This is likely to be a noisier estimate of earnings than the contemporary questions asked in the PSID.

A. Which Model Is Right?

Our results clearly rule out models that imply that saving is proportional to permanent income. But, as shown in Section II, a variety of economic models deliver the prediction that saving rates are positively related to permanent income. Is there any way to distinguish among them?

First, our results are not consistent with life cycle explanations based on differences in the timing of income. We find that including imputed Social Security and pension saving does not alter our basic result that high-income households save a larger fraction of their income. The decline in Social Security saving rates as income rises is more than offset by a rise in pension saving rates, so that including imputed Social Security and pension saving leads to a *steeper* relationship between saving rates and income within each age group. Thus models such as Huggett and Ventura (2000), which appeal to differential Social Security replacement rates, are not capable of explaining why high-income households save more.

Second, our results are inconsistent with life cycle explanations based solely on differences in time preference rates or differing elasticities of intertemporal substitution, which lead to differences in the timing of consumption. Our data sets show no evidence of the “switching” pattern at later ages implied by these explanations: Households with higher saving rates when young do not exhibit higher dissaving rates when old. This is consistent with Altonji and Villanueva’s (2003) finding that the marginal propensity to bequeath out of permanent income rises with permanent income. Indeed, we find no evidence of dissaving across all income groups, a result that differs from Hurd (1987) but is consistent with more recent studies of wealth changes among the elderly (e.g., Alessie, Lusardi, and Aldershof 1997; Alessie, Lusardi, and Kapteyn 1999; Hurd 1999; Attanasio and Hoynes 2000; Feinstein and Ho 2000; Hildebrand 2001). That active saving—netting out capital gains in equity or housing—is nonnegative among the elderly suggests that our results are not simply the consequence of the buoyant equity market during the 1980s and 1990s.

Third, precautionary saving models with uncertain medical expenses alone are not consistent with observed saving patterns because lower-income households typically face more uncertainty about health costs relative to their income. Fourth, higher saving rates for higher-income groups are consistent with an operative bequest motive as in Becker and Tomes (1986) or models in which wealth itself is an argument in the utility function (Carroll 2000); the very high saving rates of the top 1 percent or top 5 percent are difficult to explain any other way. However, we did not find significant differences in saving behavior between households with children and those without, casting doubt on the ability of

the dynastic smoothing model alone to explain the correlation between income and saving rates. Nor would a simple bequest model alone explain why observed levels of saving (and dissaving) in the lowest income quintiles are so small.⁴²

In sum, we suggest that the minimum components of a model to capture the empirical regularity that the rich save more should include a precautionary saving motive against uncertain expenditures late in life, thus explaining the nondissaving behavior among the elderly, coupled with a bequest motive. The different motives need not be exclusive: Households save for precautionary reasons but with a reasonable expectation that they will be able to pass along unspent balances to their children (Smith 1999*b*; Dynan et al. 2002).⁴³ As well, the empirical patterns of the data are consistent with an institutional or behavioral mechanism that systematically leads to low levels of saving among the poor. This may be caused by the absence of financial institutions such as pension plans or home ownership necessary to overcome time-inconsistent saving behavior. It may also be the consequence of asset-based means testing of government programs, or even a result of very high time preference rates.⁴⁴ These results can be consistent with the lack of a times-series increase in the aggregate saving rate (noted in the Introduction), as long as the theoretical model implies that saving rates are invariant to constant proportional increases in the relevant parameters, such as income, health care expenditures, anticipated income of heirs, and (as in the model above) the size of the consumption floor.

This type of model is not at odds with the difficulty in detecting a correlation between income inequality and the saving rate (Blinder 1975; Schmidt-Hebbel and Serven 2000). Were the income share received by the bottom quintile in 1998 cut from 4.2 percent of aggregate income (U.S. Bureau of the Census 2000, p. 471) to *zero*, with all the dollars appearing in the income share of the top income quintile (boost-

⁴² See Bernheim et al. (2001) and De Nardi (2002). Mulligan (1997) has also criticized this simple "bequest smoothing" model because the data suggest that consumption is no less likely to converge toward the mean among those making intergenerational financial transfers. He proposes an alternative model in which households choose how much to transfer depending on their financial resources and the degree of (endogenous) altruism toward their children.

⁴³ This view of saving behavior can reconcile the high fraction of respondents in surveys who say that they save primarily against future contingencies and the substantial fraction who say that they want to leave a bequest for their children (Laitner and Juster 1996; Horioka and Watanabe 1997; Horioka et al. 2000). For further empirical evidence, see Jianakoplos, Menchik, and Irvine (1996).

⁴⁴ Extreme time preference rates would be required to generate observed low saving rates near retirement; even in buffer stock models of consumption, saving rates revert to traditional life cycle levels as households near retirement (Carroll and Samwick 1997). Engen, Gale, and Uccello (1999) also suggest that low wealth accumulation near retirement may be the consequence of adverse income shocks. However, their model is less likely to explain *median* saving rates for households with low *permanent* income.

ing it from 47.3 to 51.5 percent), the rise in the aggregate personal saving rate predicted by the average values of active saving in figure 1 would be just 0.55 percentage point.⁴⁵ In other words, increased income inequality should tend to increase aggregate saving, but the magnitudes of such changes are likely to be modest and therefore difficult to find in the time-series data.

B. *What Are the Policy Implications?*

Our results suggest that the marginal propensities to save and to consume differ substantially across income groups, suggesting that government policies that redistribute across income groups can have real effects on saving. Although the differential saving effects of typical government transfer programs are not so large as to make a measurable dent in aggregate saving statistics (as illustrated above), our results cast doubt on the assumption of a representative agent in the economy.

In addition, our results that high-income households save a greater fraction of income while working than low-income households imply that a flat rate consumption tax will (on a lifetime basis) be regressive relative to a flat rate income tax. There are two reasons for this. First, even if there were no bequests (so that all households consumed their wealth during their lifetime), households with higher incomes would pay more under an income tax compared to an equal-revenue consumption tax, simply because such households can expect to receive a disproportionate share of interest and dividend income.

Second, one needs to take into account that bequests are not zero and may vary with lifetime earnings. A number of studies have calculated the distributional effects of moving from an income-based tax system to one that taxes just consumption. Fullerton and Rogers (1993), Metcalf (1994), and Altig et al. (2001) use data from Menchik and David (1982) to estimate how bequests differ by lifetime income group. The Menchik and David data, however, exhibited a U-shaped relationship between the fraction of resources bequeathed and lifetime income; bequests were the largest share of lifetime resources for the *lowest* income decile.⁴⁶ This pattern is consistent with the poor (in the lowest income decile) saving, on average, a larger fraction of their lifetime resources than the rich! Not surprisingly, Metcalf (1994) found, using these data, that adjusting

⁴⁵ To compute this change, we multiply 0.042 by the difference between the quintile 5 and quintile 1 saving rates. The direct estimates of the MPS based on the change-in-wealth saving measure from table 9 yield a slightly larger estimate, whereas those based on active saving imply a much smaller estimate.

⁴⁶ This U-shaped pattern may be the consequence of using Wisconsin probate records matched with tax returns; farmers are more likely to show low or zero after-tax income but leave sizable farms.

for bequests added progressivity to a consumption tax relative to an income tax. In contrast, our results show substantially higher saving rates (and wealth-income ratios) among families with higher lifetime incomes (see also Mieszkowski and Palumbo 2002). These results, together with the lack of evidence that the elderly high-income households dissave at a higher rate, suggest that, on average, higher-income households bequeath a larger fraction of lifetime earnings than lower-income households. Because bequests are effectively exempt from a consumption tax (at least for the current generation) but are not exempt from an income tax, our evidence suggests further regressivity in a consumption tax relative to an income tax.

Finally, the results have implications for the “choice versus chance” question first raised by Friedman (1953) and more recently by Venti and Wise (1998). Is the considerable variation in accumulated wealth the consequence of choice (preferences and tastes) or chance? Venti and Wise argue that most of the variation in wealth *within* lifetime income groups is due to saving decisions, or choice. Our finding that differences in saving behavior *across* income groups are also an important source of the overall variation in wealth of the U.S. population suggests a diminished role for choice. For example, wealth accumulation because of government or private policies that differentially affect saving (such as asset-based means testing or the availability of 401(k) plans) cannot be readily attributed to tastes or preferences. Nor can a lucky career outcome such as joining Microsoft in 1984, and a subsequently higher rate of saving, be attributed entirely to choice.

In sum, much remains to be learned about household saving behavior, especially that of elderly households and that of the very top of the income distribution. Still, we believe that our work has established one fact: The rich do, indeed, save more.

References

- Abowd, John M., and David Card. 1989. “On the Covariance Structure of Earnings and Hours Changes.” *Econometrica* 57 (March): 411–45.
- Alessie, Rob, Annamaria Lusardi, and Trea Aldershof. 1997. “Income and Wealth over the Life Cycle: Evidence from Panel Data.” *Rev. Income and Wealth* 43 (March): 1–32.
- Alessie, Rob, Annamaria Lusardi, and Arie Kapteyn. 1999. “Saving after Retirement: Evidence from Three Different Surveys.” *Labour Econ.* 6 (June): 277–310.
- Altig, David, Alan J. Auerbach, Laurence J. Kotlikoff, Kent A. Smetters, and Jan Walliser. 2001. “Simulating Fundamental Tax Reform in the United States.” *A.E.R.* 91 (June): 574–95.
- Altonji, Joseph G., and Ernesto Villanueva. 2003. “The Marginal Propensity to Spend on Adult Children.” Working Paper no. 9811 (July). Cambridge, Mass.: NBER.

- Atkeson, Andrew, and Masao Ogaki. 1996. "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data." *J. Monetary Econ.* 38 (December): 507–34.
- Attanasio, Orazio P. 1994. "Personal Saving in the United States." In *International Comparisons of Household Saving*, edited by James M. Poterba. Chicago: Univ. Chicago Press (for NBER).
- Attanasio, Orazio P., and Martin Browning. 1995. "Consumption over the Life Cycle and over the Business Cycle." *A.E.R.* 85 (December): 1118–37.
- Attanasio, Orazio P., and Hilary W. Hoynes. 2000. "Differential Mortality and Wealth Accumulation." *J. Human Resources* 35 (Winter): 1–29.
- Auerbach, Alan J., and Laurence J. Kotlikoff. 1987. *Dynamic Fiscal Policy*. Cambridge: Cambridge Univ. Press.
- Bar-Ilan, Avner. 1995. "On the Proportionality and Homogeneity of Consumption and Income." *Canadian J. Econ.* 28, no. 4b (November): 1153–60.
- Becker, Gary S., and Casey B. Mulligan. 1997. "The Endogenous Determination of Time Preference." *Q.J.E.* 112 (August): 729–58.
- Becker, Gary S., and Nigel Tomes. 1986. "Human Capital and the Rise and Fall of Families." *J. Labor Econ.* 4, no. 3, pt. 2 (July): S1–S39.
- Bernheim, B. Douglas, and John Karl Scholz. 1993. "Private Saving and Public Policy." In *Tax Policy and the Economy*, vol. 7, edited by James M. Poterba. Cambridge, Mass.: MIT Press (for NBER).
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among U.S. Households?" *A.E.R.* 91 (September): 832–57.
- Blinder, Alan S. 1975. "Distribution Effects and the Aggregate Consumption Function." *J.P.E.* 83 (June): 447–75.
- Bosworth, Barry, Gary Burtless, and John Sabelhaus. 1991. "The Decline in Saving: Evidence from Household Surveys." *Brookings Papers Econ. Activity*, no. 1, pp. 183–241.
- Branch, E. Raphael. 1994. "The Consumer Expenditure Survey: A Comparative Analysis." *Monthly Labor Rev.* 117 (December): 47–55.
- Browning, Martin, and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *J. Econ. Literature* 34 (December): 1797–1855.
- Cameron, Stephen, and Christopher Taber. 2000. "Borrowing Constraints and the Returns to Schooling." Working Paper no. 7761 (June). Cambridge, Mass.: NBER.
- Carroll, Christopher D. 2000. "Why Do the Rich Save So Much?" In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, edited by Joel B. Slemrod. Cambridge, Mass.: Harvard Univ. Press.
- Carroll, Christopher D., and Andrew A. Samwick. 1997. "The Nature of Precautionary Wealth." *J. Monetary Econ.* 40 (September): 41–71.
- Caselli, Francesco, and Jaume Ventura. 2000. "A Representative Consumer Theory of Distribution." *A.E.R.* 90 (September): 909–26.
- Coronado, Julia Lynn, Don Fullerton, and Thomas Glass. 2000. "The Progressivity of Social Security." Working Paper no. 7520 (February). Cambridge, Mass.: NBER.
- Covinsky, Kenneth E., et al. 1994. "The Impact of Serious Illness on Patients' Families." *J. American Medical Assoc.* 272 (December 21): 1839–44.
- Crystal, Steven, et al. 2000. "Out-of-Pocket Health Care Costs among Older Americans." *J. Gerontology*, ser. B, *Psychological Sci. and Soc. Sci.* 55 (January): S51–S62.

- De Nardi, Mariacristina. 2002. "Wealth Inequality and Intergenerational Links." Staff Report no. 314 (November). Minneapolis: Fed. Reserve Bank, Res. Dept.
- Diamond, Peter A., and Jerry A. Hausman. 1984. "Individual Retirement and Savings Behavior." *J. Public Econ.* 23 (February–March): 81–114.
- Duesenberry, James S. 1949. *Income, Saving, and the Theory of Consumer Behavior*. Cambridge, Mass.: Harvard Univ. Press.
- Dynan, Karen E. 1994. "Relative Wage Changes and Estimates of the Rate of Time Preference." Manuscript. Washington, D.C.: Fed. Reserve Bd.
- . 2002. "The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer." *A.E.R. Papers and Proc.* 92 (May): 274–78.
- Elmendorf, Douglas. 1996. "The Effect of Interest-Rate Changes on Household Saving and Consumption: A Survey." Finance and Econ. Discussion Ser., Working Paper no. 1996-27. Washington, D.C.: Fed. Reserve Bd.
- Engen, Eric M., William G. Gale, and Cori E. Uccello. 1999. "The Adequacy of Household Saving." *Brookings Papers Econ. Activity*, no. 2, pp. 65–165.
- Evans, Michael K. 1969. *Macroeconomic Activity: Theory, Forecasting, and Control*. New York: Harper and Row.
- Evans, William N., and Edward Montgomery. 1995. "Education and Health: Where There's Smoke There's an Instrument." Manuscript. College Park: Univ. Maryland.
- Feinstein, Jonathan S., and Chih-Chin Ho. 2000. "Elderly Asset Management and Health: An Empirical Analysis." Working Paper no. 7814 (July). Cambridge, Mass.: NBER.
- Feldstein, Martin, and Andrew A. Samwick. 1992. "Social Security Rules and Marginal Tax Rates." *Nat. Tax J.* 45 (March): 1–22.
- Fisher, Irving. 1930. *The Theory of Interest, as Determined by Impatience to Spend Income and Opportunity to Invest It*. New York: Macmillan.
- French, Eric, and John Baily Jones. 2003. "On the Distribution and Dynamics of Health Care Costs." Manuscript. Chicago: Fed. Reserve Bank.
- Friedman, Milton. 1953. "Choice, Chance, and the Personal Distribution of Income." *J.P.E.* 61 (August): 277–90.
- . 1957. *A Theory of the Consumption Function*. Princeton, N.J.: Princeton Univ. Press (for NBER).
- Friend, Irwin, and Irving B. Kravis. 1957. "Consumption Patterns and Permanent Income." *A.E.R. Papers and Proc.* 47 (May): 536–55.
- Fullerton, Don, and Diane Lim Rogers. 1993. *Who Bears the Lifetime Tax Burden?* Washington, D.C.: Brookings Inst.
- Gale, William G., and Samara R. Potter. 2002. "An Economic Evaluation of the Economic Growth and Tax Relief Reconciliation Act of 2001." *Nat. Tax J.* 55 (March): 133–86.
- Gentry, William M., and R. Glenn Hubbard. 2000. "Entrepreneurship and Household Saving." Working Paper no. 7894 (September). Cambridge, Mass.: NBER.
- Greenhouse, Steven. 1992. "Fairness and a Flat Tax." *New York Times* (March 26), p. 1.
- Gruber, Jonathan, and Aaron Yelowitz. 1999. "Public Health Insurance and Private Savings." *J.P.E.* 107, no. 6, pt. 1 (December): 1249–74.
- Gustman, Alan L., and Thomas L. Steinmeier. 1989. "An Analysis of Pension Benefit Formulas, Pension Wealth, and Incentives from Pensions." In *Research in Labor Economics*, vol. 10, edited by Ronald G. Ehrenberg. Greenwich, Conn.: JAI.
- . 1997. "Effects of Pensions on Savings: Analysis with Data from the Health

- and Retirement Study." Report. Washington, D.C.: U.S. Dept. Labor, Pension and Welfare Benefits Admin.
- . 1999. "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study." *Carnegie-Rochester Conf. Ser. Public Policy* 50 (June): 271–324.
- . 2000. "Social Security Benefits of Immigrants and U.S. Born." In *Issues in the Economics of Immigration*, edited by George J. Borjas. Chicago: Univ. Chicago Press (for NBER).
- . 2001. "How Effective Is Redistribution under the Social Security Benefit Formula?" *J. Public Econ.* 82 (October): 1–28.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence." *J.P.E.* 86 (December): 971–87.
- Hamermesh, Daniel S. 1984. "Consumption during Retirement: The Missing Link in the Life Cycle." *Rev. Econ. and Statis.* 66 (February): 1–7.
- Harris, Christopher, and David Laibson. 2001. "Dynamic Choices of Hyperbolic Consumers." *Econometrica* 69 (July): 935–57.
- Hicks, John R. 1950. *A Contribution to the Theory of the Trade Cycle*. London: Oxford Univ. Press.
- Hildebrand, Vincent. 2001. "Wealth Accumulation of US Households: What Do We Learn from the SIPP Data?" Research Paper no. 41 (June). Hamilton, Ont.: McMaster Univ., Res. Program Soc. and Econ. Dimensions Aging Population.
- Horioka, Charles Yuji, Hideki Fujisaki, Wako Watanabe, and Takatsuga Kouno. 2000. "Are Americans More Altruistic than the Japanese? A U.S.-Japan Comparison of Saving and Bequest Motives." Working Paper no. 7463 (January). Cambridge, Mass.: NBER.
- Horioka, Charles Yuji, and Wako Watanabe. 1997. "Why Do People Save? A Micro-Analysis of Motives for Household Saving in Japan." *Econ. J.* 107 (May): 537–52.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conf. Ser. Public Policy* 40 (June): 59–125.
- . 1995. "Precautionary Saving and Social Insurance." *J.P.E.* 103 (April): 360–99.
- Huggett, Mark, and Gustavo Ventura. 2000. "Understanding Why High Income Households Save More than Low Income Households." *J. Monetary Econ.* 45 (April): 361–97.
- Hurd, Michael D. 1987. "Savings of the Elderly and Desired Bequests." *A.E.R.* 77 (June): 298–312.
- . 1999. "Portfolio Holdings by the Elderly." Manuscript. Santa Monica, Calif.: Rand Corp.
- Hurd, Michael D., and James P. Smith. 2001. "Anticipated and Actual Bequests." In *Themes in the Economics of Aging*, edited by David A. Wise. Chicago: Univ. Chicago Press (for NBER).
- Hurd, Michael D., and David A. Wise. 1989. "The Wealth and Poverty of Widows: Assets before and after the Husband's Death." In *The Economics of Aging*, edited by David A. Wise. Chicago: Univ. Chicago Press (for NBER).
- Hurst, Erik, and Annamaria Lusardi. 2004. "Liquidity Constraints, Wealth Accumulation, and Entrepreneurship." *J.P.E.* 112 (April): 319–47.
- Jianakoplos, Nancy A., Paul L. Menchik, and F. Owen Irvine. 1996. "Saving Behavior of Older Households: Rate-of-Return, Precautionary and Inheritance Effects." *Econ. Letters* 50 (January): 111–20.

- Juster, F. Thomas, Joseph Lupton, James P. Smith, and Frank Stafford. 2001. "The Decline in Household Saving and the Wealth Effect." Manuscript. Ann Arbor: Univ. Michigan.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. New York: Harcourt, Brace.
- Laibson, David I. 1997. "Golden Eggs and Hyperbolic Discounting." *Q.J.E.* 112 (May): 443–77.
- Laibson, David I., Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *Brookings Papers Econ. Activity*, no. 1, pp. 91–172.
- Laitner, John, and F. Thomas Juster. 1996. "New Evidence on Altruism: A Study of TIAA-CREF Retirees." *A.E.R.* 86 (September): 893–908.
- Lawrance, Emily C. 1991. "Poverty and the Rate of Time Preference: Evidence from Panel Data." *J.P.E.* 99 (February): 54–77.
- Leeper, Eric M., and Christopher A. Sims. 1994. "Toward a Modern Macroeconomic Model Usable for Policy Analysis." In *NBER Macroeconomics Annual 1994*, edited by Stanley Fischer and Julio J. Rotemberg. Cambridge, Mass.: MIT Press.
- Liebman, Jeffrey B. 2002. "Redistribution in the Current U.S. Social Security System." In *The Distributional Aspects of Social Security and Social Security Reform*, edited by Martin Feldstein and Jeffrey B. Liebman. Chicago: Univ. Chicago Press (for NBER).
- Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." In *The Phillips Curve and Labor Markets*, edited by Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conf. Ser. Public Policy, vol. 1. Amsterdam: North-Holland.
- Lusardi, Annamaria. 1999. "Information, Expectations, and Savings for Retirement." In *Behavioral Dimensions of Retirement Economics*, edited by Henry J. Aaron. Washington, D.C.: Brookings Inst. Press.
- MaCurdy, Thomas E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *J. Econometrics* 18 (January): 83–114.
- Mayer, Thomas. 1966. "The Propensity to Consume Permanent Income." *A.E.R.* 56 (December): 1158–77.
- . 1972. *Permanent Income, Wealth, and Consumption: A Critique of the Permanent Income Theory, the Life-Cycle Hypothesis, and Related Theories*. Berkeley: Univ. California Press.
- Menchik, Paul L., and Martin David. 1982. "The Incidence of a Lifetime Consumption Tax." *Nat. Tax J.* 35 (June): 189–203.
- Metcalfe, Gilbert E. 1994. "Life Cycle versus Annual Perspectives on the Incidence of a Value Added Tax." In *Tax Policy and the Economy*, vol. 8, edited by James M. Poterba. Cambridge, Mass.: MIT Press (for NBER).
- Mieszkowski, Peter, and Michael G. Palumbo. 2002. "Distributive Analysis of Fundamental Tax Reform." In *United States Tax Reform in the 21st Century*, edited by Gene R. Zodrow and Peter Mieszkowski. Cambridge: Cambridge Univ. Press.
- Modigliani, Franco, and Albert Ando. 1960. "The 'Permanent Income' and 'Life Cycle' Hypothesis of Saving Behavior: Comparison and Tests." In *Proceedings of the Conference on Consumption and Saving*, vol. 2, edited by Irwin Friend and Robert Jones. Philadelphia: Univ. Pennsylvania, Wharton School Finance and Commerce.
- Mulligan, Casey B. 1997. *Parental Priorities and Economic Inequality*. Chicago: Univ. Chicago Press.
- Nelson, Julie A. 1994a. "Consumer Expenditure Surveys, 1980–89: Interview

- Surveys, for Household-Level Analysis." Computer file. Washington, D.C.: U.S. Dept. Labor, Bur. Labor Statist.
- . 1994*b*. "On Testing for Full Insurance Using Consumer Expenditure Survey Data." *J.P.E.* 102 (April): 384–94.
- Ogaki, Masao, Jonathan D. Ostry, and Carmen M. Reinhart. 1996. "Saving Behavior in Low- and Middle-Income Developing Countries: A Comparison." *Internat. Monetary Fund Staff Papers* 43 (March): 38–71.
- Passell, Peter. 1995. "Spending It; the Tax Code Heads into the Operating Room." *New York Times* (September 5), sec. 3, p. 1.
- Pigou, Arthur C. 1951. "Professor Duesenberry on Income and Savings." *Econ. J.* 61 (December): 883–85.
- Powers, Elizabeth T. 1998. "Does Means-Testing Welfare Discourage Saving? Evidence from a Change in AFDC Policy in the United States." *J. Public Econ.* 68 (April): 33–53.
- Quadri, Vincenzo. 1999. "The Importance of Entrepreneurship for Wealth Concentration and Mobility." *Rev. Income and Wealth* 45 (March): 1–19.
- . 2000. "Entrepreneurship, Saving, and Social Mobility." *Rev. Econ. Dynamics* 3 (January): 1–40.
- Sabelhaus, John. 1993. "What Is the Distributional Burden of Taxing Consumption?" *Nat. Tax J.* 46 (September): 331–44.
- Samwick, Andrew A. 1998. "Discount Rate Heterogeneity and Social Security Reform." *J. Development Econ.* 57 (October): 117–46.
- Schmidt-Hebbel, Klaus, and Luis Servén. 2000. "Does Income Inequality Raise Aggregate Saving?" *J. Development Econ.* 61 (April): 417–46.
- Shapiro, Matthew D., and Joel B. Slemrod. 2003. "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys." In *Tax Policy and the Economy*, vol. 17, edited by James M. Poterba. Cambridge, Mass.: MIT Press (for NBER).
- Skinner, Jonathan. 1985. "The Effect of Increased Longevity on Capital Accumulation." *A.E.R.* 75 (December): 1143–50.
- . 1987. "A Superior Measure of Consumption from the Panel Study of Income Dynamics." *Econ. Letters* 23 (2): 213–16.
- Smith, James P. 1999*a*. "Healthy Bodies and Thick Wallets: The Dual Relationship between Health and Economic Status." *J. Econ. Perspectives* 13 (Spring): 145–66.
- . 1999*b*. "Inheritances and Bequests." In *Wealth, Work, and Health: Innovations in Measurement in the Social Sciences: Essays in Honor of F. Thomas Juster*, edited by James P. Smith and Robert J. Willis. Ann Arbor: Univ. Michigan Press.
- Stoker, Thomas M. 1986. "Simple Tests of Distributional Effects on Macroeconomic Equations." *J.P.E.* 94 (August): 763–95.
- Thaler, Richard H. 1994. "Psychology and Savings Policies." *A.E.R. Papers and Proc.* 84 (May): 186–92.
- U.S. Bureau of the Census. 2000. *Statistical Abstract of the United States 2000*. Washington, D.C.: Government Printing Office.
- Venti, Steven F., and David A. Wise. 1998. "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *A.E.R. Papers and Proc.* 88 (May): 185–91.
- . 1999. "Lifetime Earnings, Saving Choices, and Wealth at Retirement." In *Wealth, Work, and Health: Innovations in Measurement in the Social Sciences: Essays in Honor of F. Thomas Juster*, edited by James P. Smith and Robert J. Willis. Ann Arbor: Univ. Michigan Press.

- Vickrey, William. 1947. "Resource Distribution Patterns and the Classification of Families." In *Studies in Income and Wealth*, vol. 10. New York: NBER.
- Yitzhaki, Shlomo. 1987. "The Relation between Return and Income." *Q.J.E.* 102 (February): 77-95.
- Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *J.P.E.* 97 (April): 305-46.
- Zellner, Arnold. 1960. "Tests of Some Basic Propositions in the Theory of Consumption." *A.E.R. Papers and Proc.* 50 (May): 565-73.