WHAT ACCOUNTS FOR THE VARIATION IN RETIREMENT WEALTH AMONG U.S. HOUSEHOLDS?

B. Douglas Bernheim
Jonathan Skinner
Steven Weinberg

Working Paper **6227**

NBER WORKING PAPER SERIES

WHAT ACCOUNTS FOR THE VARIATION IN RETIREMENT WEALTH AMONG U.S. HOUSEHOLDS?

B. Douglas Bernheim
Jonathan Skinner
Steven Weinberg

Working Paper 6227
http://www.nber.org/papers/w6227

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1997

We are grateful for very helpful comments from Martin Browning, Richard Blundell, Karen Dynan, Leora Friedberg, Al Gustman, Annamaria Lusardi, and Steven Venti, and seminar participants at the University of Georgia, Northwestern University, M.I.T., the University of California, Berkeley, and the NBER Summer Institute. This work has been supported by a grant from the National Science Foundation to the National Bureau of Economic Research. This paper is part of NBER's research programs in Aging and Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

What Accounts for the Variation in Retirement
Wealth Among U.S. Households?
B. Douglas Bernheim, Jonathan Skinner
and Steven Weinberg
NBER Working Paper No. 6227
October 1997
JEL Nos. E2, D1
Aging and Public Economics

ABSTRACT

Household survey data consistently depict large variations in saving and wealth, even among households with similar socio-economic characteristics. Within the context of the life cycle hypothesis, families with identical lifetime resources might choose to accumulate different levels of wealth for a variety of reasons, including variation in time preference rates, risk tolerance, exposure to uncertainty, relative tastes for work and leisure at advanced ages, income replacement rates, and so forth. These factors can be divided into a small number of classes, each with a distinctive implication concerning the relation between accumulated wealth and the shape of the consumption profile. By examining this relation empirically, one can test for the presence or absence of these particular explanations for differences in wealth. Using the Panel Study of Income Dynamics and the Consumer Expenditure Survey, we find very little support for life cycle models that rely on the above factors to explain wealth variation. The data are, however, consistent with "rule of thumb" or "mental accounting" theories of wealth accumulation.

B. Douglas Bernheim
Department of Economics
Stanford University
Stanford, CA 94305
and NBER
bernheim@leland.stanford.edu

Jonathan Skinner
Department of Economics
Dartmouth College
Hanover, NH 03755
and NBER
jon.skinner@dartmouth.edu

Steven Weinberg
Department of Economics
Stanford University
Stanford, CA 94305
wsteve@leland.stanford.edu

I. Introduction

Household survey data consistently depict large variations in saving and wealth, even among households with similar socio-economic characteristics (see e.g. Bernheim, Lemke, and Scholz, 1997, Venti and Wise, 1996). Since one would hardly expect all individuals to have identical tastes and experiences, the mere existence of this variation is not at all surprising. However, the interpretation of this variation is of paramount importance. If one takes the view that saving reflects rational, farsighted optimization, then low savers are simply expressing their preferences for current consumption over future consumption -- one cannot coherently claim that they are saving "too little" given the after-tax rate of return, any more than one can assert that people would be better off if they spent more time listening to classical music (e.g., Lazear, 1994). However, if one takes the view that households are shortsighted, irrational, prone to regret, or heavily influenced by psychological motives, then the adequacy of saving among various population subgroups emerges as an important and potentially well-posed empirical question (see Bernheim, 1995).

There are, of course, a variety of factors that could in principle account for the observed variation in wealth within the context of models with rational, farsighted optimization. Households may differ with respect to patience (as represented by the rate of pure time preference), risk tolerance, exposure to uncertainty, health status, perceived life expectancies, relative tastes for work and leisure at advanced ages, levels of work-related expenses, complementarities between consumption and leisure activities, lifetime earnings, income replacement rates, and so forth. These factors provide substantial scope for explaining observed behavior.

In this paper, we investigate the empirical validity of various explanation for the observed variation in savings for retirement by studying data on wealth and consumption drawn from the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX). Our investigation proceeds from the central observation that, for a wide range of models, there are clear implications relating the variation in accumulated wealth to *variations in the shape of the consumption profile*. By examining this relation empirically, one can test for the presence or absence of particular factors.

As we argue in section II, factors that could in principle account for variations in accumulated wealth (holding constant earnings and pension income) can be divided into three categories. Factors in the first category give rise to systematic correlations between wealth and the *rate of change of consumption*, either before or after retirement (or both). For example, a direct implication of the standard Euler equation approach is that households with high time preference rates experience a lower (or negative) rate of change in consumption compared to their more patient counterparts. Factors in the second category give rise to systematic correlations between wealth and *one-time changes in consumption at retirement.* For example, with variation in work-related expenses, those with higher expenses will accumulate less wealth and also experience larger declines in measured consumption at retirement. Likewise, if some individuals are unexpectedly forced to retire earlier than expected, they will both retire with less wealth (because their accumulation phase is terminated prematurely), and experience sharper declines in consumption at retirement (because they revise their expectations about lifetime resources). Factors in the third category give rise to systematic correlations between accumulated wealth and the *level of consumption*, but not between accumulated wealth and changes in consumption. Households with strong bequest motives, for example, would tend to experience a lower level of consumption throughout the life cycle.

The first central finding of this paper is that there is no discernable relation in the data between accumulated wealth and rates of change in consumption prior to retirement, or after retirement. In the six years prior to retirement, there is virtually no difference in the growth rate of consumption among households with different relative wealth holdings. Nor is there discernable systematic differences by wealth in the growth path of consumption in the years following retirement. This strong finding suggests that a wide range of factors (including possible variation in pure rates of time preference) fail even to provide contributory explanations for the observed variations in accumulated wealth.

Our second central finding is the existence of a correlation between accumulated wealth and declines in consumption at retirement. There are several plausible explanations for why consumption might drop at retirement, such as the termination of work-related expenses or the substitution of leisure or "home production" for market expenditures. However, we do not find empirical support for these explanations; the pure impact of retirement on work-related expenses is trivial, while expenditures on food consumed away from home drops by only a modest amount more (and in some cases by less) than expenditures on food at home. Nor is there evidence that this relationship is the consequence of unanticipated shocks leading both to early retirement and to declines in consumption; the drop in consumption persists even when the unpredictable components of retirement are removed through a two-stage regression analysis.

Taken together, our first two findings imply that a broad range of factors operating within models of rational, farsighted, optimizing agents are collectively incapable of accounting for joint patterns of wealth and consumption (holding constant the lifecycle income profile). Of course, life cycle models imply -- and the data confirm -- that the overall variation in wealth is partly attributable to variations in lifetime earnings. However, life cycle models also imply that wealth should vary systematically with the *shape* of the earnings (plus social security and pensions) profile. In particular, one of the basic tenets of the life cycle hypothesis is that households use accumulated savings to smooth consumption over periods in which there are anticipated changes in income. This is inconsistent with our third central finding -- that the decline in consumption at retirement is highly correlated with the household's income replacement ratio. This pattern is present even when one removes the effects of unanticipated shocks affecting the timing of retirement.

Overall, our results pose a significant challenge to the validity of standard life cycle models. Instead, they appear to suggest that *on average* individuals who arrive at retirement with few resources experience a "surprise" -- they take stock of their finances only to discover that their resources are insufficient to maintain their accustomed standards of living (e.g. because pension income is less than expected, or because they recognize that savings will go less far than they had hoped), and they revise their expectations downward in light of this realization.

This paper is closely related to a number of existing studies that examine the behavior of consumption at and around retirement. Hamermesh (1984) and Mariger (1987) find pronounced declines in consumption as households move into retirement. In addition, Hamermesh (1984) calculates that the *level* of consumption typically exceeds annuitized income, and in this sense does not appear to be sustainable -- a conclusion that is disputed by Kotlikoff, Spivak, and Summers (1982) who focus on average (or permanent) income over a broader time horizon instead of income and consumption near retirement. Using Canadian data, Robb and Burbridge (1989) estimated a much larger (and discrete) drop of consumption at retirement for blue-collar workers relative to white-collar workers. Hausman and Paquette (1987) link the decline in (non-medical) consumption at retirement to unexpected and involuntary job loss, often resulting from health problems. Banks, Blundell, and Tanner (1996) track consumption and earnings of a given British age cohort through retirement years, and document a sharp drop in consumption at age 65 that is difficult to explain with reference to conventional economic factors. In an earlier version of the paper, Banks and Blundell (1993) find that while earnings dropped 39 percent between age 61 and 71, consumption dropped by 35 percent. That is, the sharp decline in average earnings for this cohort was tracked almost exactly by an equivalent decline in average consumption. While these authors attribute some of the decline in consumption to the fall in work-related expenses, there remains an unexplained drop that they attribute to a negative "surprise" at retirement.

II. Theoretical Preliminaries

The logic of the household's budget constraint dictates that variations in wealth held at retirement must correspond to variations in consumption profiles, non-asset income profiles, or both. By studying the relations between these variations, it is possible to shed light on the validity of theories that purport to account for household saving behavior.

A. Variation in consumption profiles

For households with similar earnings and pension income profiles,[1] those reaching retirement with less wealth will have consumed more prior to retirement, and will consume or bequeath less in total after retirement. By itself, the budget constraint does not tie down the characteristics of the consumption profile more precisely than this. Specific models of behavior imply that the consumption profile accommodates the budget constraint in one (or more) of three ways. First, for those with lower wealth at retirement, consumption may decline gradually either before retirement, after retirement, or throughout the life cycle. Second, consumption may decline discontinuously at

[1]Among other things, similarity of income profiles implies that the households in question retire at the same ages.

retirement, and this discontinuity may be larger for those with less accumulated wealth. Third, those with less accumulated wealth at retirement may simply bequeath less, consuming more throughout their lives.[2] By studying the relations between accumulated wealth and consumption profiles, one can therefore identify the factors that contribute (as well as the factors that do not contribute) to the observed variation in wealth.

Factors affecting the slope of the consumption profile

Variation in accumulated wealth at retirement could in principle be attributable to any factor that produces variation in the slope of the consumption profile: subject to the qualifications discussed below, rising consumption profiles correspond to high accumulation, while falling consumption profiles correspond to low accumulation. The most obvious factors of this kind include variations in patience (the pure rate of time preference), longevity, exposure to uncertainty, and risk tolerance.

To illustrate the roles of these factors, consider a time-separable utility function of the form

$$U_t = U(C_t) + E_t\left\{\sum_{s=t+1}^{\infty} \lambda_{t,s} U(C_s)\right\}, \tag{1}$$

[2]Purely as a matter of logic, there are, of course, other possibilities. For example, consumption may decline discontinuously at some point after retirement, and this point may occur at a more advanced age for those who reach retirement with greater wealth. One could produce this pattern in a model with heterogeneity in finite, deterministic lifespans. However, the assumption of a deterministic lifespan is very unattractive. More generally, variation in survival probabilities gives rise to variations in the slope of the consumption profile, which is an example of the first pattern.

with

$$\lambda_{t,s} = \left(\frac{1}{1+\delta}\right)^{s-t} \prod_{k=t+1}^{s} (1-\pi_k), \tag{2}$$

where C_i is consumption in period i, E_t is the expectations operator (conditional on information available at time t), δ is the standard (constant) pure rate of time preference, and π_k is the probability of dying before period k, conditional upon surviving to period k-1. Assume for the moment that income is potentially uncertain and independently distributed across periods, but that the age of retirement is fixed.

Maximization of (1) subject to a resource constraint yields the following first order condition:

$$U'(C_t) = \alpha_t E_t\{U'(C_{t+1})\} \tag{3}$$

where

$$\alpha_t = \frac{(1+r)(1-\pi_{t+1})}{1+\delta} \tag{4}$$

Taking a second-order Taylor-series expansion of (3) yields the familiar Euler equation

$$\frac{E_t(C_{t+1}) - C_t}{C_t} \approx \gamma\left(1 - \frac{1}{\alpha_t}\right) + \frac{\psi}{2}\sigma_{t+1}^2 \tag{5}$$

where γ is the intertemporal elasticity of substitution, ψ reflects the household's precautionary inclinations, and σ_{t+1} is the standard deviation of consumption in period t+1, so that the final term captures "precautionary saving" (see, e.g., Deaton, 1991; Dynan, 1993).

Equation (5) tells us that the slope of the consumption profile depends on a number of factors, including the rate of interest, the pure rate of time preference, perceived survival probabilities, and precaution. Variation in any of these factors can produce variation in the shape of the consumption profile, and associated variation in wealth.

To illustrate, consider first the effects of variation in "patience" (the pure rate of time preference). If the elasticity of substitution is close to zero (Leontief preferences), then variation in δ does not alter the consumption trajectory, and therefore cannot account for observed variation in wealth accumulation. Although early studies placed γ close to zero (see, e.g., Hall, 1988), more recent work finds evidence of a small positive elasticity (see, e.g., Attanasio and Weber, 1993, 1995). Consequently, one would expect the slope of the consumption profile to vary with δ.

To see this correlation explicitly, consider the familiar case where utility is given by $U(C) = C^\gamma/\gamma$, income is deterministic, and π_t is zero until the date of death. Figure 1a shows the consumption paths for two households in a simple simulation model.[3] The interest rate is 3 percent, but the impatient households have a value for δ of 3.5 percent, while for patient households, $\delta = 2.5$ percent. Since both patient and impatient

---

[3]In this certainty model, individuals live from age 21 to 85, earnings are as given in Hubbard, Skinner, and Zeldes (1994) for high school dropouts, and there are no liquidity constraints.

Figure 1a: Age-Consumption Profiles for High- and Low-Time-Preference Households

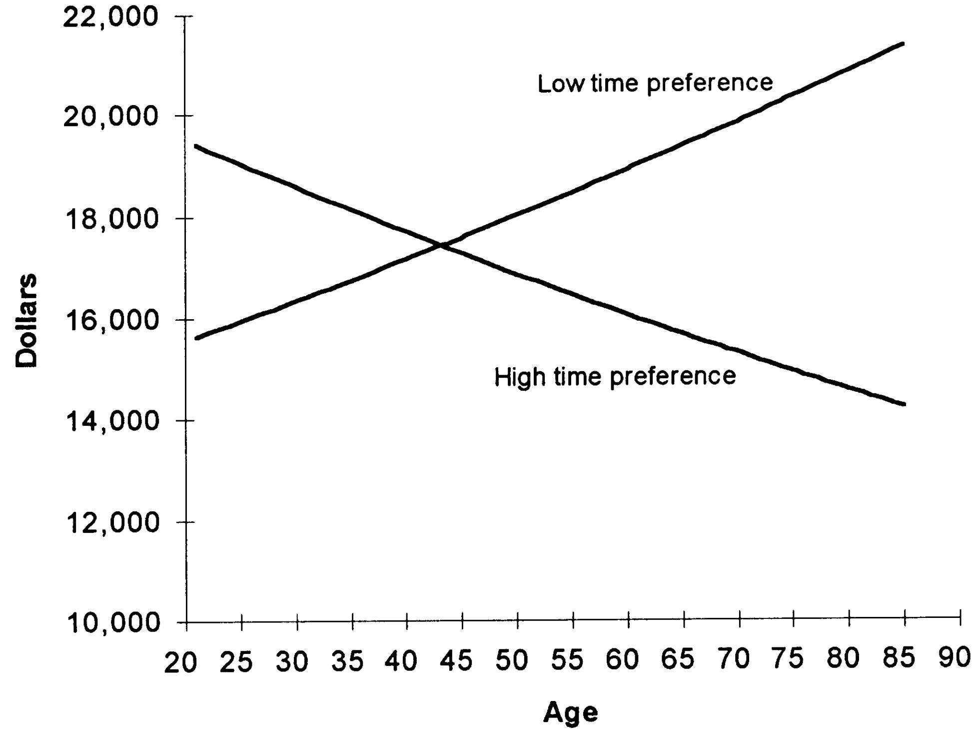

Figure 1b: Age-Wealth Profiles for High- and Low-Time-Preference Households

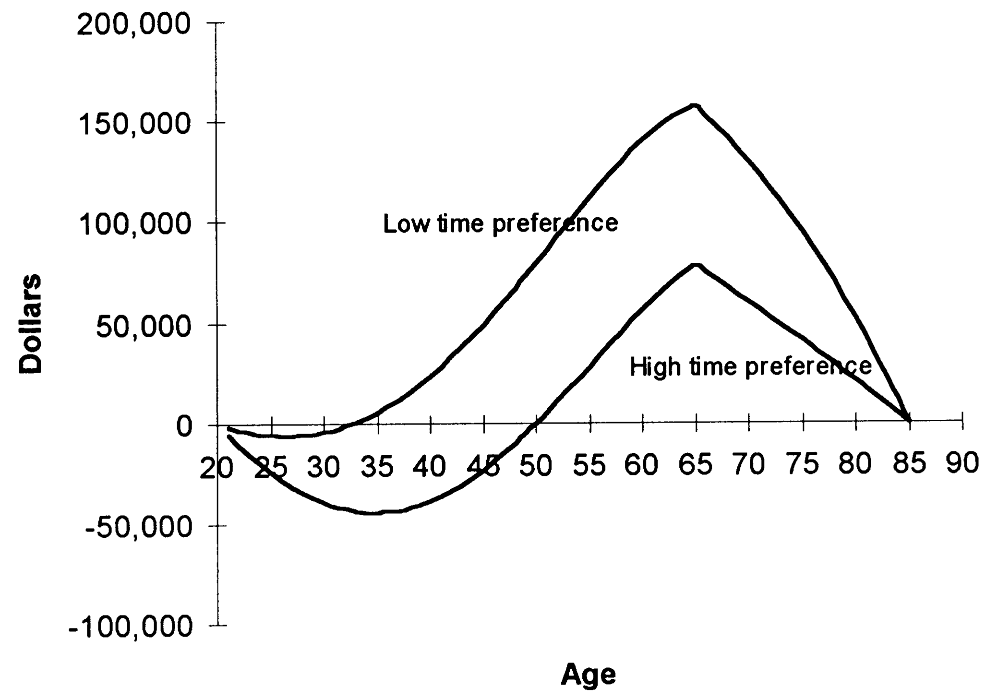

households satisfy the budget constraint, these variations in δ correspond to *rotations* of the consumption profile, with (all else equal) impatient households consuming more than patient households early in life, and patient households consuming more than impatient households later in life. Consequently, patient households accumulate more wealth for retirement, as shown in Figure 1b. This produces a positive correlation between the slope of the consumption profile and accumulated wealth at retirement.

The logic of this argument must be modified somewhat in the presence of liquidity constraints. In the absence of income uncertainty, sufficient impatience drives a household to save nothing, or to exhaust all pre-existing wealth. As a result, consumption tracks income, and the slope of the consumption profile is dictated mechanically by the slope of the income profile. Likewise, with income uncertainty, impatient individuals may act as "buffer stock" savers in the sense of Deaton (1991) and Carroll (1997), with consumption again tracking income.

While consideration of liquidity constraints disrupts the clean prediction of the basic model (a positive correlation between consumption growth and assets accumulated at retirement), it does not fundamentally alter our ability to determine whether variations in wealth are, at least to some extent, attributable to variations in the pure rate of time preference. In particular, among individuals with high levels of wealth, liquidity constraints and buffer stock saving are relatively unimportant. Consequently, if differences in the rate of time preference contribute to observed differences in wealth, then the predicted correlation should be present among households with significant wealth relative to earnings (say the top half of the population distribution). The slope of the consumption profile may be higher or lower for households at the bottom of the wealth distribution, depending on the slope of the income profile; however, except by accident, the slope of the consumption profile should also vary as one moves from the middle of the wealth distribution (where the consumption profile is governed by the Euler equation) to the lower end of the wealth distribution (where the consumption profile is governed by the income profile).

As is clear from equation (5), variations in perceived longevity (survival probabilities) are essentially equivalent to variations in the pure rate of time discount, and therefore produce the same correlations between the slope of the consumption profile and accumulated wealth at retirement. Notably, we have allowed survival probabilities to vary with age. This allows us to underscore the fact that variations in wealth at retirement may be correlated with variations in the slope of the consumption profile at some -- but not all -- stages of the life cycle. Imagine, for example, that two individuals expect to have the same conditional survival probabilities through retirement, but that one correctly anticipates more rapid increases in mortality probabilities at advanced ages. In that case, the one with greater expected longevity will accumulate greater resources. While there will be no observable differences between the slopes of their consumption profiles before retirement, significant differences should emerge after retirement. Specifically, the individual with less accumulated wealth at retirement will experience a more rapid decline in consumption (conditional on survival) after retirement.

Equation (5) also implies that the slope of the consumption profile is sensitive to income uncertainty and the household's tastes for precaution. Greater uncertainty and/or precaution rotates the consumption profile: a higher growth rate of consumption is accompanied by precautionary saving, which leads (for the same income realizations) to greater accumulated wealth at retirement. As with survival probabilities, income uncertainty may differ across households at some ages, but not at others. Imagine, for example, that pre-retirement earnings are uncertain, but that post-retirement pension income is not. Then households with greater earnings uncertainty will experience higher rates of consumption growth prior to retirement, and reach retirement with greater wealth, again implying a correlation between consumption growth and wealth. Subsequent to retirement, consumption growth rates will equalize.

Thus, variations in wealth accumulation that result from variations in the pure rate of time preference, perceived longevity, income uncertainty, and/or tastes for precaution should manifest themselves through a positive correlation between accumulated wealth at retirement and the growth rate of consumption either before retirement, after retirement, or both. The absence of such correlations would be inconsistent with the hypothesis that these factors contribute to the observed variation in wealth.

Thus far, we have focused exclusively on factors that arise in standard life cycle planning models. Similar factors are also present in other models. There are, for example, a number of models proposed in recent years that posit dynamic inconsistencies, such as the Laibson (1997) model of hyperbolic discounting, the Thaler and Shefrin

(1981) "doer-planner" model, or the Akerlof (1991) model of procrastination. Laibson's model is obtained by inserting an additional parameter, β (normally thought to be less than unity), measuring the extent to which individuals (inconsistently) discount all future utility relative to today's utility,[4] into the utility function of equation (1):

$$U_t = U(C_t) + \beta E_t \left\{ \sum_{s=t+1}^{\infty} \lambda_{t,s} U(C_s) \right\}, \tag{6}$$

With this formulation, the intertemporal consumption path is determined as the equilibrium of a game played by successive incarnations of the decision-maker. In this game, the decision-maker may or may not have access to savings technologies that fully or partially alleviate the dynamic inconsistencies by binding the hands of its future incarnations.

As Laibson shows, there is a tendency for dynamically inconsistent planners to save too little, and to exhibit what appear to be high rates of time preference. Thus, variations in β, like variations in δ, tend to induce rotations of the consumption profile along with corresponding changes in asset accumulation. To the extent differences in time consistency are important determinants of the observed variation in wealth, one would therefore expect to observe the predicted positive correlation between consumption

---

[4]At time t+1, the discount factor β discounts future utility beginning at time t+2. Hence the marginal rate of substitution between consumption at time t+1 versus consumption at time t+2 is different from the perspective of time t and from t+1.

growth rates and accumulated wealth at retirement.

If one steps outside the framework of rational, farsighted optimization, then it is possible to imagine other explanations for variations in wealth that would manifest themselves through correlations between accumulated wealth and the slope of the consumption profile. Imagine, for example, that households follow simple rules of thumb for saving before retirement; after retirement, they gradually learn whether they have saved enough to comfortably sustain their accustomed standard of living. If different households follow different rules of thumb, then (all else equal) they will arrive at retirement with different levels of accumulated wealth. Those with little wealth may initially attempt to sustain their pre-retirement standard of living, but will soon observe rapid depletion of their resources, leading them to reduce consumption, perhaps in a succession of small steps. Those with moderate wealth at retirement might sustain accustomed levels of consumption for a longer period of time before becoming aware of their plight and making accommodating adjustments. Those with substantial wealth at retirement might never confront the need to economize, and indeed might choose to increase consumption after watching their resources grow. Thus, under this view, variations in pre-retirement rules of thumb produce a positive short-term (say 3-4 year) correlation between accumulated wealth at retirement, and the growth of consumption after retirement.

Factors affecting the change in consumption at retirement

A number of factors are capable of producing downward, discontinuous jumps in consumption at retirement. By the logic of the budget constraint, the existence of a sharper decline at retirement generally implies higher consumption, and therefore less wealth accumulation, before retirement. Figure 2 illustrates the impact of such a drop on life cycle wealth holdings. The figure is based on a simulation model similar to that used on Figure 1, but with a time preference rate of 3 percent; hence the benchmark case without any discontinuous jump at retirement shows, in Figure 2a, a flat consumption profile throughout the life cycle. We also include a hypothetical consumption profile satisfying the same lifetime budget constraint, but with an (arbitrarily) inflated level of consumption before retirement, and therefore a substantial drop in consumption at retirement. The household with the consumption drop at retirement holds less wealth throughout the life cycle, as shown in Figure 2b. To the extent such factors account for the observed variation in wealth, one should therefore observe a negative correlation between the absolute size of the discontinuity and accumulated wealth at retirement.

One factor capable of producing this pattern is variation in work-related expenses. In simple terms, one can think of work-related expenses as items that should ideally be netted against income, but which are mistakenly counted as consumption. If households smooth non-work-related consumption, then the cessation of work-related expenses at retirement should produce a discontinuous drop in observed consumption. Comparing households with different levels of work-related expenses (all else equal), those with high

Figure 2a: Age-Consumption Profile With (and Without) A Discrete Shift in Consumption at Retirement

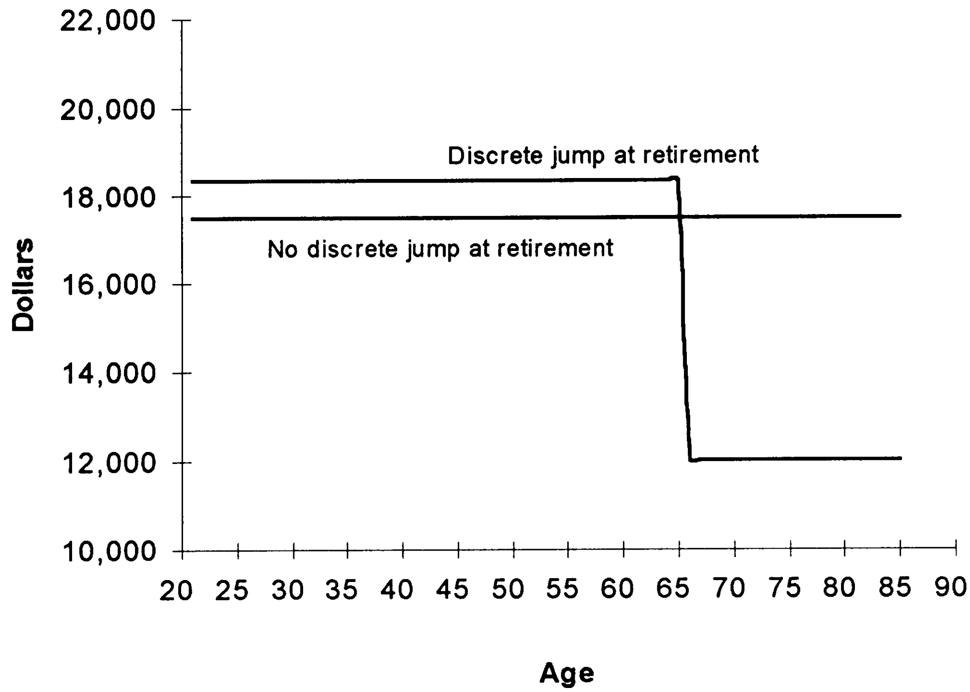

Figure 2b: Age-Wealth Profile With (and Without) a Discrete Shift in Consumption at Retirement

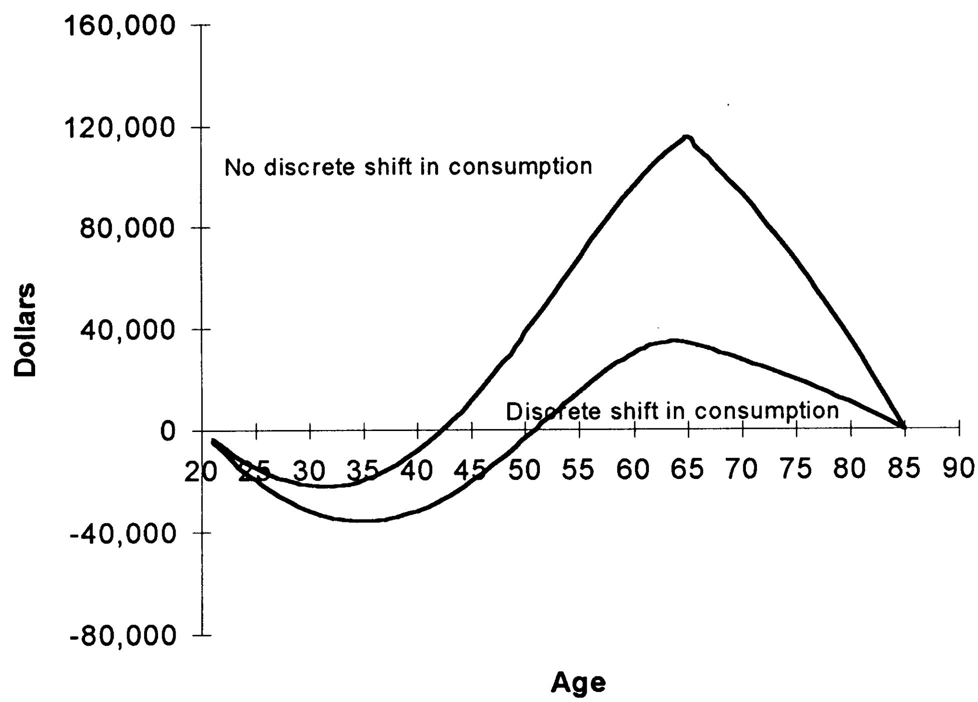

expenses will have higher measured consumption and lower saving before retirement, lower accumulated wealth at retirement, and lower consumption after retirement. If variation in work-related expenses accounts for part of the observed variation in wealth, one should therefore observe a negative correlation between the absolute size of the discontinuity *in expenditure categories that are specifically work related*, and the accumulation of wealth at retirement. If there is a negative correlation between the size of the consumption discontinuity and accumulated wealth, but if this correlation is similar for work-related and non-work-related expenses, then this particular factor would not be supported as an important source of the variation in observed wealth.

Similar comments apply to the possibility that there may be variation in the extent to which leisure acts as a substitute for or complement to the consumption of goods and services. As is well-known, substitutability between consumption and leisure can in theory produce a decline in spending at retirement (Ghez and Becker, 1975, Baxter and Jerman, 1994, Banks, Blundell, and Tanner, 1996, Butler, 1996). For example, retired individuals may spend their retirement leisure time in "household production," mowing one's own lawn rather than hiring somebody to do it, or cooking one's own meals rather than eating out. Likewise, complementarities between consumption and leisure can generate increases in spending at retirement, as when retirees seek to spend leisure time traveling or playing golf.

Variation in the degree of substitutability and/or complementarity across the population translates directly into variation in the size (and possibly the direction) of the discontinuity in consumption at retirement. Individuals who view leisure as a powerful substitute for consumption expenditures may consume more and save less prior to retirement, e.g. counting on their inclination to engage in home production after retirement. Conversely, individuals who think of leisure as a complement to marketed goods and services may consume less and save more prior to retirement in order to finance costly leisure activities during their "golden" years. If variation in substitutabilities and complementarities account for part of the observed variation in wealth, one should therefore observe a negative correlation between the absolute size of the discontinuity *in expenditure categories that are complements to leisure*, and the accumulation of wealth at retirement. If there is a negative correlation between the size of the consumption discontinuity and accumulated wealth, but if this correlation is similar for complements to and substitutes for leisure, then this particular factor would not be supported as an important source of the variation in observed wealth.

Discontinuities in consumption at retirement could also arise for liquidity constrained households, whose consumption profiles simply track income.[5] Since pensions and social security typically replace significantly less than 100 percent of pre-retirement earnings, low savers would experience sharp drops in consumption at retirement, while high savers would smooth consumption across retirement. This would

[5]For example, Hubbard, Skinner, and Zeldes (1995) showed that asset-based means testing (for Medicaid, SSI, or AFDC) in combination with uncertainty about earnings or medical expenses can lead to such discontinuities. However, even households with very high time preference rates (e.g., "buffer stock" savers) should be accumulating wealth just prior to a discontinuous drop in income at retirement, so as to smooth consumption.

again produce a negative correlation between retirement wealth and the size of the consumption discontinuity. However, this factor would be incapable of explaining the existence of a similar correlation among the subset of households with non-negligible retirement savings.

Consumption discontinuities (and the associated variation in wealth) may also be attributable to variation in plans concerning the timing of retirement. All else equal, individuals who plan to retire at advanced ages will tend to save and accumulate less than those who plan to retire early. If plans are always realized, this cannot account for the observed variation in wealth among those who retire at any given age. Suppose, however, that retirement is sometimes unexpected and involuntary, perhaps resulting from a sudden deterioration of health or loss of job (Diamond and Hausman, 1984, Hausman and Paquette, 1987). Then, even among those who actually retire at the same age, there may be substantial variation in the *planned* age of retirement.

To illustrate the implications of this observation, consider three individuals, A, B, and C, who are identical in all but the following respect: while working, A planned to retire at age 65, B planned to retire at age 70, and C planned to retire at age 75. Despite this divergence of intentions, all actually retire at age 65 (B and C experience unexpected and involuntary job loss). Standard considerations imply that, in planning for a longer retirement period, A will have saved more than B, who will in turn have saved more than C, by age 65. Since retirement at age 65 comes as no surprise to A, this individual continues on the same consumption trajectory as before retirement; abstracting from the considerations discussed earlier in this section, A would experience no consumption discontinuity at retirement. In contrast, B is surprised by the events that trigger retirement at age 65. From B's perspective, retirement is "bad news" in that it signals the loss of five years worth of future earnings. B adjusts to this news by reducing consumption discontinuously at retirement. C is similar to B, except that the "news" is more severe, and results in an even greater reduction in consumption.

Thus, if variation in retirement plans accounts for the observed variation in wealth among those retiring at a given age, one should again observe a negative correlation between the absolute size of the consumption discontinuity at retirement, and the level of accumulated wealth at retirement. Unlike work-related expenses or substitutabilities and complementarities, this factor is capable of explaining the potential finding that this negative correlation exists for all categories of consumption, rather than for categories that are specifically related to work and/or leisure. However, under this view, if one can statistically remove the effects of "news" associated with unexpected retirement, then the correlation between the consumption discontinuity and wealth at retirement should disappear. Banks, Blundell, and Tanner (1996) focus on changes in consumption around age 65 when many British workers retire. Thus they implicitly remove the idiosyncratic news from retirement events, and implicitly estimate the reduced form correlation between the exogenous factors leading to retirement (being 65) and changes in consumption. We follow a similar approach below in a two-stage model that estimates how predicted changes in retirement (once again removing the potentially endogenous news) affects consumption behavior.

If one again steps outside the framework of rational, farsighted optimization, then it is possible to imagine other explanations for variations in wealth that would manifest themselves through correlations between accumulated wealth and the magnitude of the consumption discontinuity at retirement. Suppose, for example, that individuals engage in a variety of heuristic and quasi-rational strategies to determine their saving prior to retirement (or simply procrastinate, as in Akerlof, 1991). Once they reach retirement, they take stock of their financial situations, and adjust their living standards to accommodate their resources. In that case, the adequacy of savings at retirement is "news": those with bad news (inadequate savings) may decrease their consumption, while those with good news (excessive savings) may increase it. Note that behavior of this sort would imply correlations between wealth and the size of discontinuities in all consumption categories. Moreover, these correlations would persist even if the "news" associated with the age of retirement was removed statistically (in contrast to the other sources of variation in wealth discussed earlier in this section).

Yet another possibility is suggested by theories of "mental accounting" (Thaler and Shefrin, 1981, Thaler, 1994). Within the context of these theories, current income is generally regarded as more "spendable" than assets, particularly those held in retirement accounts or converted into annuities. Consequently, consumption may decline at retirement for the simple reason that an important category of spendable resources (current earnings) disappears. Suppose that individuals differ in the extent to which they exercise self-discipline over spendable funds. Then those with little self-discipline will accumulate little wealth for retirement, and will experience substantial declines in consumption with the cessation of spendable earnings, while those with greater self-discipline will accumulate greater wealth for retirement, and experience smaller declines in consumption with the cessation of spendable earnings. Once again, the source of variation in wealth manifests itself through a negative correlation between the size of the consumption discontinuity and accumulated wealth; this correlation should cut across consumption categories, and should persist even when one removes the "news" associated with unexpected retirement.

Factors affecting the overall level of consumption

Variation in wealth at retirement could be attributable to factors that raise or lower the overall level of consumption throughout the life cycle. The logic of the budget constraint dictates that this variation in expenditures would have to be absorbed by bequests (Dynan, Skinner, and Zeldes, 1997). Figure 3a illustrates two simulated consumption profiles (again with the time preference rate equal to the interest rate), with the lower consumption stream being for the household who seeks to bequeath 5% (in present value terms) of their lifetime wealth. The associated wealth profile for the household with the bequest motive is consistently above that for the household without an operative bequest motive, as shown in Figure 3b. Thus variation in wealth accumulation

Figure 3a: Age-Consumption Profiles With (and Without) a Bequest Motive

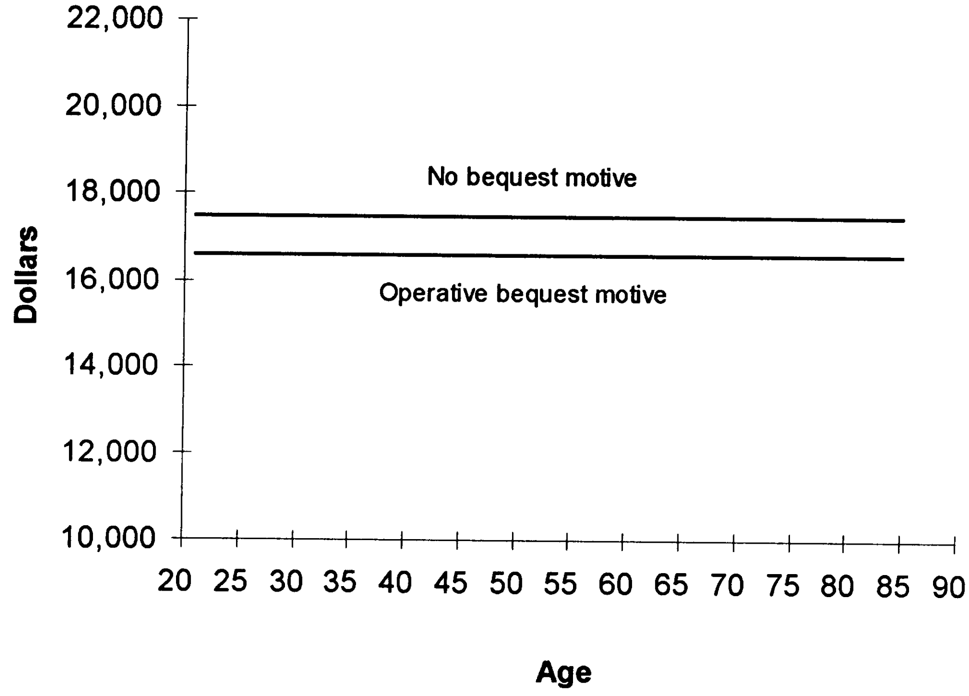

Figure 3b: Age-Wealth Profiles With (and Without) a Bequest Motive

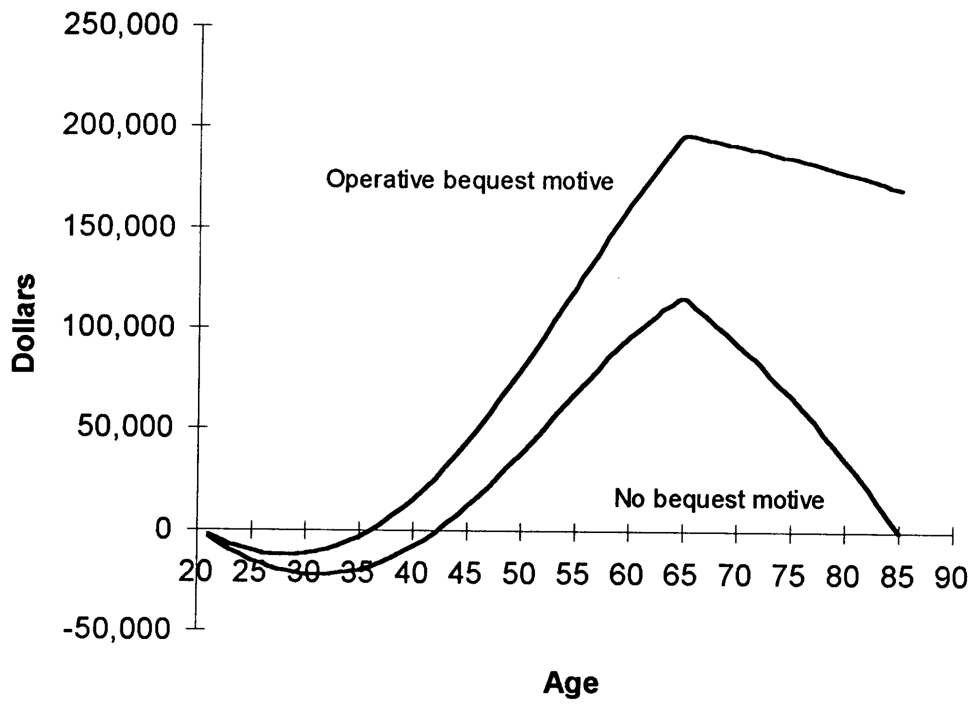

might occur if, for example, there is significant variation in the strength of bequest motives across the population (for supporting empirical evidence see Bernheim, 1991, and Laitner and Juster, 1996). Note that theories of this kind have no implications concerning the likely correlations between accumulated wealth and changes in consumption, either before, at, or after retirement.

B. Variation in income profiles

In the preceding subsection, we have focused our attention on factors that might explain variations in wealth for households with similar earnings and pension income profiles. Naturally, variation in these profiles may also contribute to the overall variation in wealth at retirement.

Survey data generally show that wealth at any age rises significantly with proxies for lifetime resources, such as household earnings. While this pattern certainly helps to account for the fact that some households accumulate more wealth than others, it fails to discriminate between any interesting behavioral hypotheses, and therefore sheds very little light on the underlying determinants of household saving.

Variations in wealth at retirement may also result from variations in the *shape* of the income profile, and the nature of this relation has the potential to shed considerably more light on the validity of competing behavioral models. One widely understood implication of the life cycle hypothesis is that household saving should vary inversely with earnings replacement rates (households anticipating larger declines in non-asset income at retirement should save more). This observation forms the basis for the literature examining the relation between savings and retirement benefits (mainly social security and private pensions). Since private saving and retirement benefits may not be perfect substitutes, estimated rates of "crowding out" shed little light on the validity of the life cycle hypothesis. However, irrespective of the degree of substitutability between these classes of resources, one of the central tenets of the life cycle hypothesis is that households use accumulated savings to smooth consumption over periods in which there are anticipated changes in income. Consequently, the life cycle hypothesis would be undermined by a finding that, among those with non-negligible retirement savings, lower earnings replacement ratios tend to be associated with larger consumption discontinuities at retirement.

The life cycle model might be reconciled with the hypothesized finding if there was a significant unanticipated component of post-retirement non-asset income. Yet this income consists primarily of social security and private pension benefits. Information concerning the size of these benefits is readily available to households approaching retirement. Consequently, any indication that households are "surprised" by post-retirement non-asset income casts doubt on the view that these same households are engaged in the kind of deliberate retirement planning envisioned in theory.

In principle, the hypothesized finding could also be attributable to sample selection problems. The mere existence of a correlation between pension eligibility and tastes for saving would be insufficient to explain the existence of a relation between the earnings replacement rate and the consumption discontinuity at retirement. However, sample selection issues could obscure our findings in more subtle ways. In particular, employers may offer less generous pensions (relative to earnings) to workers with higher work-related expenses, in anticipation of a decline in their total expenditures at retirement. While this would explain the hypothesized correlation between income the income discontinuity and the consumption discontinuity, it would also imply that this relation should be much stronger among categories of consumption that are work-related, than among those categories that are not work-related.

III. Empirical Analysis

In this section, we use household survey data to examine the empirical relations between wealth at retirement, consumption profiles, and (non-asset) income profiles. For the reasons discussed in section II, these relations shed considerable light on the underlying determinants of (and sources of variation in) wealth accumulation.

We present our empirical analysis in three stages. First, we analyze patterns of total consumption around retirement to determine whether the slope of the consumption profile, the discontinuity in consumption at retirement, or both are systematically related to accumulated wealth and/or non-asset income profiles. Second, we examine consumption on a disaggregated basis to determine whether different patterns emerge for goods that are complements to work (i.e. work-related expenses) or leisure. Third, we reexamine the patterns of interest using statistical techniques to remove the "news" associated with the timing of retirement.

A. The path of consumption

Specification

As we noted in section II, a convenient approach to discriminating among competing models is to compare *growth rates* of consumption, as well as *discontinuities* in consumption at retirement, across households with differing levels of wealth accumulation and rates of earnings replacement. Our general strategy is to estimate a function of the following form:

$$\ln(C_{it}) = \mu_i + f(t,X_i) + Z_{it}\Gamma + \varepsilon_{it} \qquad \textbf{(7)}$$

where $\ln(C_{it})$ is the log of consumption for household i at time t relative to retirement (i.e, t = -2 means two years prior to retirement), μ_i is the household-specific effect (which proxies for all factors that would cause permanent differences in the level of consumption across households, such as permanent income and age), Z_{it} is a vector of household characteristics that can change through time (e.g. health status and household composition), Γ is a parameter vector, and ϵ_{it} is a random disturbance. Most importantly, $f(\bullet)$ is some flexible functional form that allows the shape of the consumption profile before and after retirement, as well as discontinuities at retirement, to vary with the vector X_i, which includes measures of the household's retirement wealth and earnings replacement rates. By the logic of section II, estimation of $f(\bullet)$ allows us to test for the presence of various factors that could in principle explain the variation in personal savings.

In the text of this paper, we report results based on specifications of $f(\bullet)$ that are linear in t both before and after retirement, and that allow for a discontinuous jump in consumption at retirement. We also allow the slope of the consumption profile to differ during the pre-retirement and post-retirement periods, and to vary with the household's income replacement rate and retirement wealth. A number of the figures included in subsequent sections are based on an alternative, more flexible (and much less parsimonious) nonparametric specification (presented in the appendix), wherein t enters through a collection of dummy variables. While it is possible in principle to define separate dummy variables for each value of t, in practice this requires the estimation of a very large number of parameters (since each dummy is interacted with each of the earnings replacement and retirement savings variables). We therefore use dummies based on two-year intervals (i.e. t = -5 or -6, t = -3 or -4, etc.).

Data

The data sample consists of the set of all households surveyed in the Panel Study of Income Dynamics (PSID) with a transition to retirement between the years 1978 and 1990. We define non-retired households to be those with at least one member (head or spouse) working more than 1500 years annually. We define a household to be retired if no member works more than 500 hours annually in the current year, or in any subsequent year for which data are available. Naturally, some individuals made the transition from non-retired status to retired status over the course of several years, during which time one or more member worked part-time (between 1500 and 500 hours).[6] We restricted the sample to households with transition periods of less than five years.[7] For the remaining observations, the variable t is set equal to -1 in the last year in which the household is non-retired, and equal to +1 in the first year in which the household is retired;[8] transition period data are excluded from our analysis. In the course of our empirical investigation, we controlled for lengthy transitions by including a dummy variable to identify households that spent more than two years (but fewer than four years) in transition to retirement; however, the coefficients of this variable were rarely significant and generally not large.

The total sample used in this analysis generally includes in excess of 3500 observations on 430 households, with the specific number depending on the regression specification. The samples are generally unbalanced, since later-retiring households are

---

[6]Households may have shifted from part time back to full time as well during this transition period. See Gustman and Steinmeier, 1984.

[7]39 percent of households made the transition to retirement in one year, 73 percent took no more than two years, 80 percent took no more than three years, and 85 percent took no more than four years. Thus, the restriction excludes roughly 15 percent of the potential sample. Of these, roughly half made the transition in five or six years.

[8]Note that for someone who made the transition from full time to part time work in 1982, and from part time to fully retired in 1984, the year prior to retirement would be 1981 and the year after retirement would be 1984.

observed for fewer years after retirement, and more years before retirement, than early-retiring households. This does not, however, appear to drive our findings, as similar results are obtained when the regressions are estimated using subsamples consisting of balanced panels.

Unfortunately, the PSID does not contain ideal data on consumption. While many past researchers have used food consumption to proxy for total consumption, Skinner's (1987) analysis of the Consumer Expenditure Survey (CEX) indicates that superior measures are available. Specifically, by using additional information on consumption reported in both the PSID and the CEX, such as the composition of food expenditures, utility payments, value of the house, and car ownership, one can increase the predictive power of the PSID consumption index three-fold.[9] We use this approach below with a more restricted set of consumption indicators -- food at home, food away from home (excluding meals at work or school), and the imputed or actual rental value of one's residence (utilities and autos were reported only sporadically) -- although we also report regression results for each component separately.[10] In certain years, the PSID did not collect information on food consumption, and we were forced to exclude these

---

[9]In other words, using the Consumer Expenditure Survey data, the R^2 of a regression of total consumption on total food consumption was 0.26; with the additional components of consumption, the R^2 rose to 0.78. This is not surprising given the importance of rental or owner-occupied housing expenses in a typical budget.

[10]In a well-specified demand structure for consumption goods, relative prices might be expected to affect the consumption weights. Palumbo (1997) found that accounting for price changes during retirement did not much affect his results (although his focus was somewhat different from ours).

observations from the sample.[11]

As mentioned above, equation (7) allows the shape of the consumption profile to depend upon the household's earnings replacement rate and retirement savings, which are summarized by the vector X_i. To allow for functional flexibility, we use dummy variables that indicate the household's position in the sample distribution of each variable.[12] Specifically, we divide our sample into four equally sized quartiles based on the ratio of income post-retirement income (total pension, social security, transfer, and earned income for years t = 1 through t = 3) to pre-retirement non-asset income (for years t = 1 through t = 3).[13] All measures of income are after-tax, where the household's tax rate is determined by taking the ratio of federal taxes paid by the head and spouse to the total income of the head and spouse. It is important to emphasize that we define earnings replacement in terms of ratios, so that, for example, the fourth (or highest) quartile includes both high income households with generous post-retirement compensation packages, as well as low income households with high social security replacement rates.

Similarly, we also divide our sample into four equally sized quartiles based on the ratio of wealth in the year prior to retirement to average pre-retirement (year t = 1 through

---

[11]We also dropped 28 observations for which the respondent reported zero food consumption, either away or at home.

[12]For the results reported in the text, we use variables that measure the household's position in the sample distribution, rather than in the (weighted) population distribution. Similar results are obtained when our variables are defined in terms of population distributions.

[13]Note that we exclude asset income from the numerator and denominator of the ratio.

t = 3) non-asset after-tax income. Recognizing that there is some disagreement in the literature on retirement saving concerning the appropriate measure of wealth (compare, for example, Poterba, Venti, and Wise, 1996, and Engen, Gale, and Scholz, 1996), we estimate separate specifications for financial wealth and total wealth; generally, the results are quite similar.

Unfortunately, since the PSID collected comprehensive information on the components of wealth only in 1984 and 1989, we typically do not directly observe wealth in the year prior to retirement. In such cases, we extrapolate retirement wealth by applying the intertemporal budget constraint. Specifically, using observed wealth in either 1984 or 1989 (whichever is closer) along with our estimates of consumption and measures of money income, we backcast (or project) wealth inductively according to the equation

$$W_{t-1} = \frac{W_t - Y_{t-1} + C_{t-1}}{1+r} \tag{8}$$

where r, the interest rate, is assumed in calculations to be equal to 3 percent. We have also estimated (but do not report) specifications based on the following alternative approach: first, we estimated regressions explaining wealth as a non-linear function of age, retirement, and other characteristics; second, we used the coefficients for age and retirement from these regressions to adjust the observed wealth of each household (in either 1984 or 1989) to the year prior to retirement. This approach yielded similar results.

Equation (7) also permits the consumption profile to depend on a vector of other time-varying household characteristics, Z_{it}. For the specifications reported here, this vector includes demographic variables that are likely to affect consumption changes, such as family size, disability status and gender of the household head, and marital status. Recall that the household-specific constant term μ_i in equation (7) absorbs factors that remain constant during the sample period. The coefficient on female headship, for example, is thus identified by the effect on consumption of the death of a husband. Twenty seven percent of household heads report disabilities, 17 percent are female, and 68 percent are married.

Additional summary statistics for the sample appear in Table 1. There is substantial variation in income replacement rates, wealth ratios, and the level of wealth. These numbers are suggestive of the wide heterogeneity in retirement preparation among the retirement-aged population. The extent to which income profiles differ across the population is illustrated in Figure 4, which traces the evolution of income through retirement for households in each of the four income replacement quartiles, controlling for the household's retirement wealth quartile.[14] After-tax income is normalized relative to income in the first and second years prior to retirement (-1&-2 in Figure 4). Note that there is little difference in the shape of the income profile prior to retirement across income replacement quartiles. Naturally, the normalized profiles for these four groups

[14]Figure 4 is based on a regression of log after-tax-income on the same explanatory variables as in equation 7.

Figure 4: Change in Income at Retirement, by Quartile

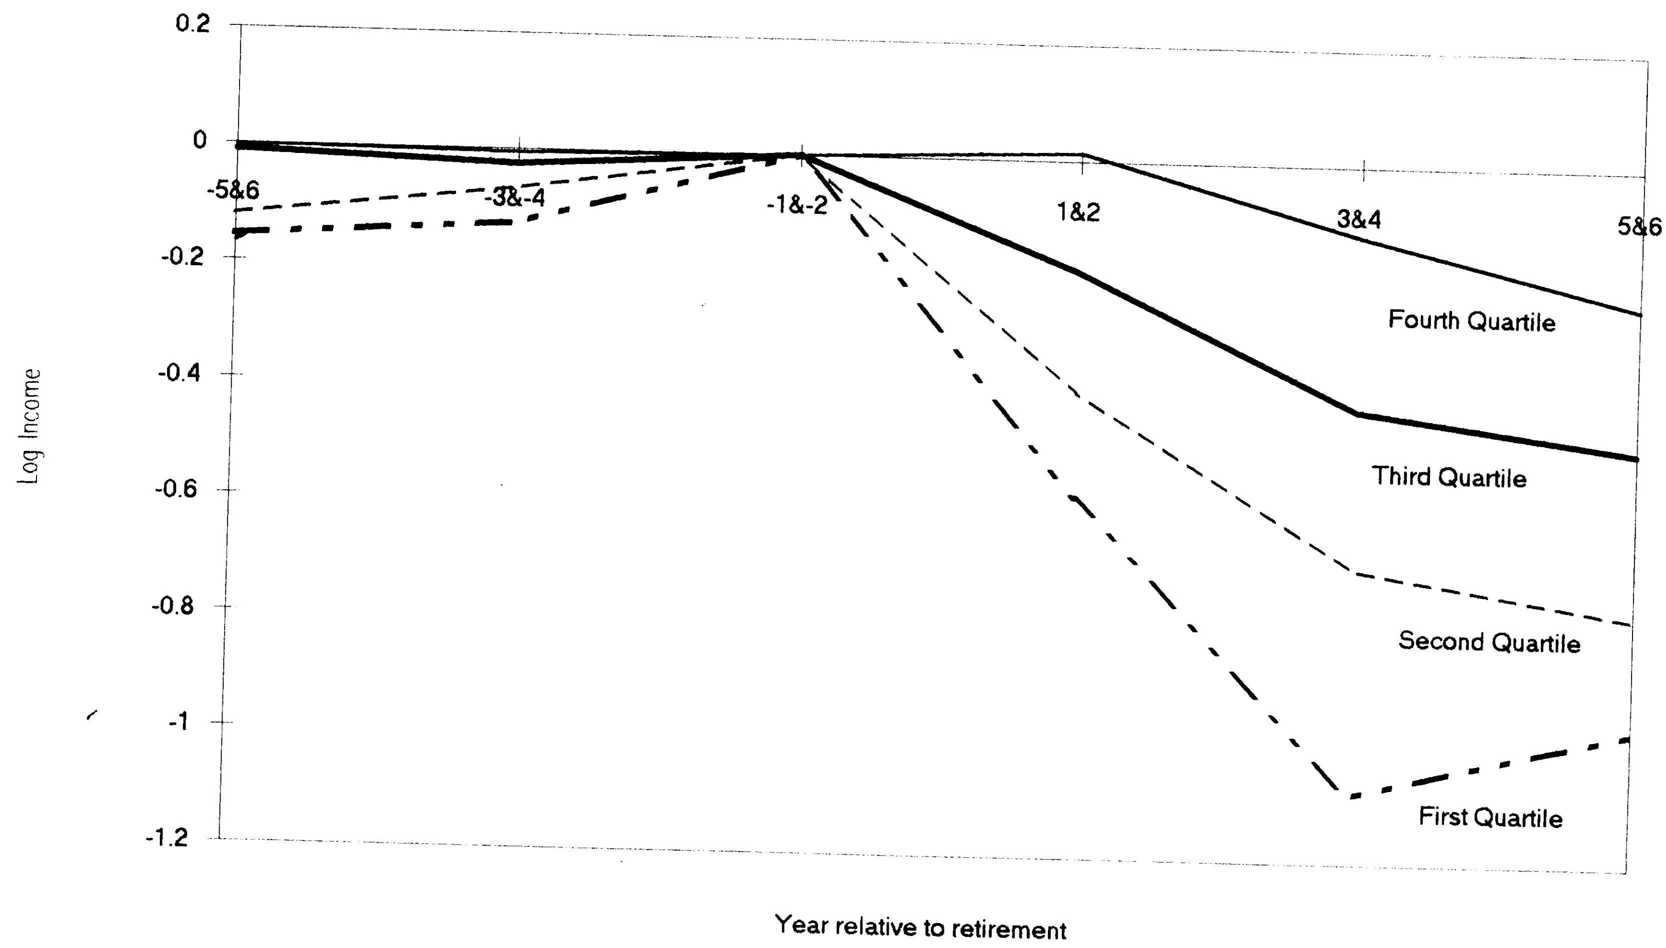

diverge significantly post-retirement, with the top income replacement quartile showing only a small change in income, while income drops dramatically (by more than 60 percent) for households in the lowest quartile. Thus, for a large subset of the population, there is a dramatic decline in after-tax income at retirement.

Further descriptive statistics are provided in table 2. For the purpose of this table, we have limited our definition of wealth to financial assets; however, virtually identical patterns are observed when a more comprehensive measure of wealth is used. Table 2a shows the joint distribution of the sample over the wealth ratio quartiles and income replacement quartiles. Note that a substantial fraction of the population saves little despite the fact that income declines significantly after retirement. Indeed, the distribution of the sample is surprisingly uniform. There is no evidence that the sample is concentrated along the Southwest/Northeast diagonal of the table -- those in the high income replacement quartiles are no more likely (or less likely) to be in the high wealth quartiles. This is not what one would expect to find if wealth varies across the population in part because households save more when their income replacement rates are lower. There may, of course, be a variety of subtle explanations for this pattern; taken by itself, it does not justify strong inferences concerning behavior.

Table 2b shows average retirement ages for each of the wealth ratio and income replacement quartiles. Differences in retirement age across these groups are small, and no systematic relations are apparent. This does not lend much credence to the view that the observed variation in wealth is in part attributable to unexpected developments affecting the timing of retirement. In particular, though early retirees are most likely to have entered retirement unexpectedly, they do not appear to have accumulated less for retirement, relative to their incomes.

Table 2c displays average pre-retirement income for each of the wealth ratio and income replacement quartiles. It is readily apparent that higher income households save more relative to their incomes. However, these same households also tend to have lower income replacement rates. In part, this is no doubt a reflection of the progressivity of social security.

Results

Tables 3a and 3b present estimates of equation (7), based on measures of financial wealth and total wealth, respectively. In each table, the first row describes the shape of the consumption profile for households in the lowest income replacement and wealth ratio quartiles (henceforth, the "benchmark" group). There is virtually no growth rate in consumption prior to retirement (-0.007 in table 3a, -0.005 in table 3b), a substantial drop downward at retirement (-0.322 in table 3a, -0.309 in table 3b) and a somewhat more pronounced annual decline in consumption post-retirement (-0.030 in table 3a, -0.040 in table 3b).

The second section of each table contains coefficients that describe how the consumption profiles of individuals in the second, third, and fourth wealth ratio quartiles differ from those of households in the benchmark group. Notice, in table 3a (which uses financial wealth), that there are virtually no difference across the wealth quartiles in the growth rates of consumption either before retirement or after retirement. Tests for equality of consumption growth across all wealth ratio quartiles before retirement and after retirement, as well as the joint test of both hypotheses (P-value for F(6,3432) of 0.78), fail to reject even at low levels of confidence.

Results based on total wealth (table 3b) also exhibit negligible differences in consumption growth rates by retirement savings, except within the second wealth ratio quartile. Because of this exception, tests for equality of coefficients before and after retirement are somewhat marginal. Nevertheless, there is no evidence of any *systematic* relation that would point to the presence of one or more of the factors discussed in the first portion of section II.A. The consumption growth rate of the second wealth ratio quartile is abnormally low relative to the benchmark group before retirement, but abnormally high after retirement. The first pattern suggests a lower pure rate of time preference (or similar parameter) than the benchmark group, while the second suggests the opposite. In addition, there is no evidence of any *systematic* relation between retirement savings and consumption growth either before or after retirement -- the second wealth quartile is simply an outlier. Finally, unlike other patterns that we emphasize in this paper, the anomalous behavior of the second wealth quartile is not robust with respect to variations in the method of estimation.[15]

[15]For example, when wealth ratio and income replacement quartiles are defined using population weights rather than sample frequencies, the anomalous findings for the second wealth quartile disappear. Specifically,

Thus, there is no indication in the data that any of the factors discussed in the first portion of section II.A (those that should manifest themselves through systematic differences in the slopes of consumption profiles before and/or after retirement) explain any portion of the variation in actual wealth.

In contrast, the results in tables 3a and 3b indicate that there are large differences across wealth ratio quartiles in the size of the consumption discontinuity at retirement. Generally, a higher wealth ratio is associated with a smaller decline in consumption. We find very little decline in consumption for households in the highest wealth ratio quartile (-0.040 in table 3a, and -0.074 in table 3b). The associated coefficients are individually significant at high levels of confidence, and one rejects the hypothesis of equal discontinuities across wealth ratio quartiles at the 99.9 percent level of confidence in both regressions. These findings suggest that the factors discussed in the second portion of section II.A (those that should manifest themselves through systematic differences in the size of the consumption discontinuity) may contribute significantly to the variation in actual wealth; however, further investigation is required before one can draw inferences concerning any particular factor.

Figure 5 depicts normalized consumption paths for households in each of the four financial wealth ratio quartiles (assuming that the household falls into the first income replacement quartile), as well as for households in the top wealth ratio and income

---

the P-values for the test of equal consumption growth rates across wealth ratio quartiles reported in table 3b change to 0.65 (from 0.06) pre-retirement, and to 0.92 (from 0.11) post-retirement. Otherwise, the results change very little.

Figure 5: Change in Consumption at Retirement, by Wealth Quartile

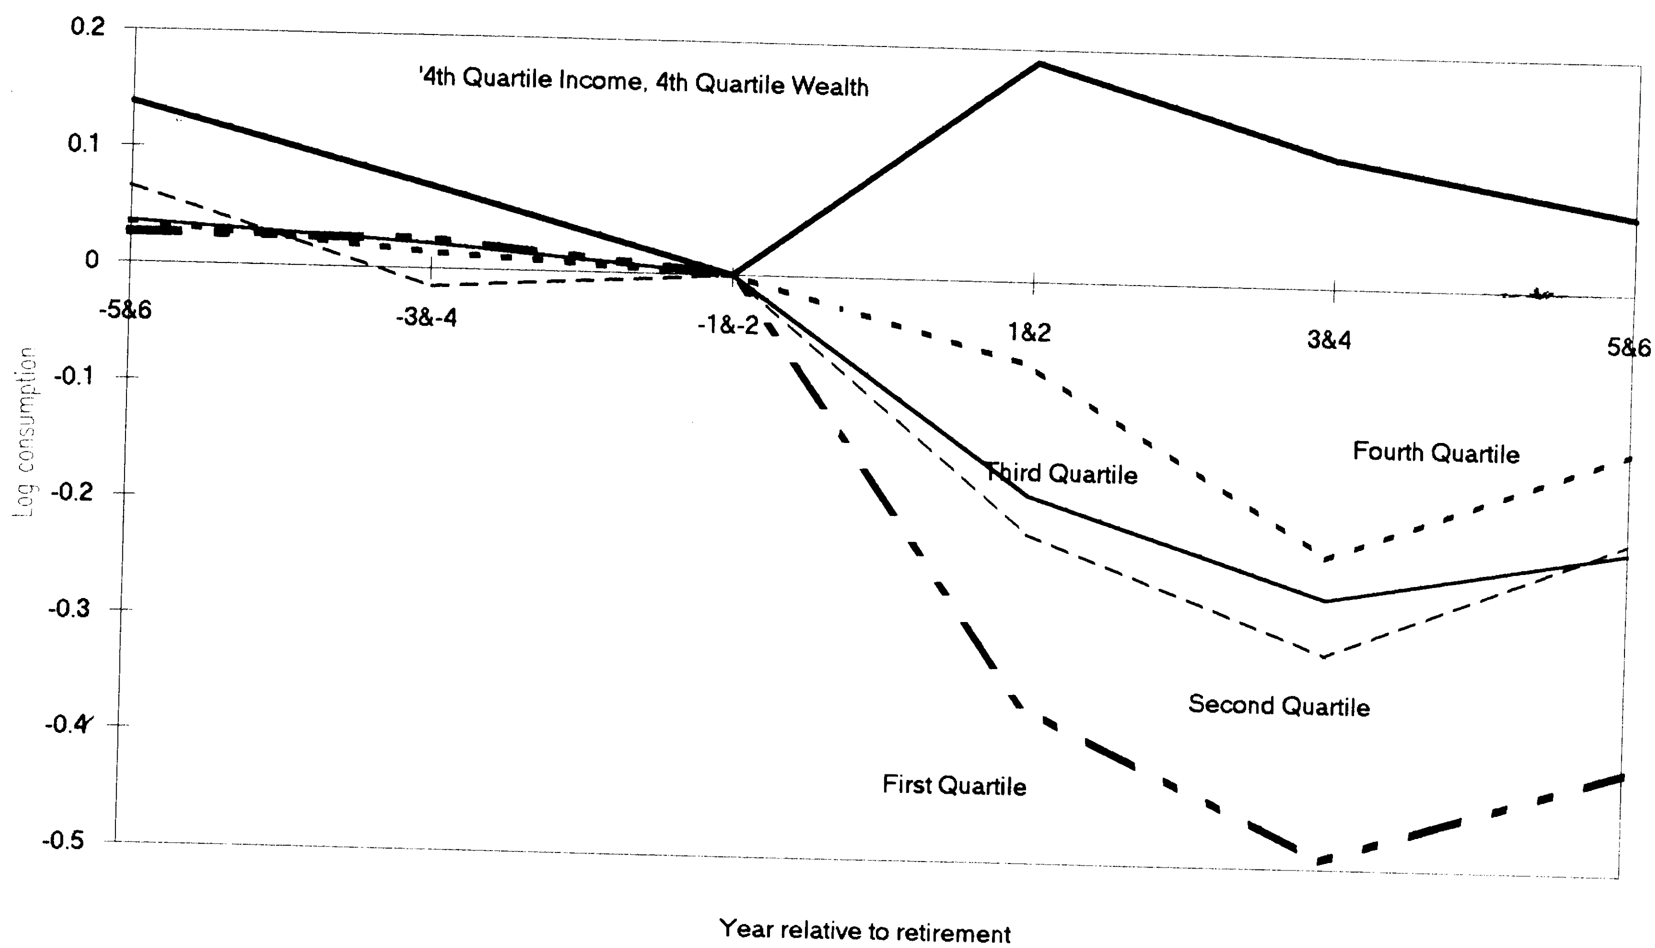

replacement quartiles.[16] This figure is based on the more flexible, less parsimonious, non-parametric specification discussed above (see Appendix Table A.1 for the associated parameter estimates). For each group, the figure normalizes consumption relative to its level in the first and second years before retirement. Once again, one sees the absence of systematic differences across wealth ratio quartiles in consumption growth rates either before or after retirement. Indeed, one cannot reject the joint hypothesis that consumption profiles are flat for every wealth quartile both before and after retirement (P-value for F(12, 3414) of 0.73). However, the figure also exhibits large differences in consumption discontinuities at the time of retirement.[17]

Differences in the shape of consumption profiles across income replacement quartiles are also intriguing. In both tables 3a and 3b, we see large and highly significant differences across these groups in the size of the consumption discontinuity at retirement. Higher income replacement rates are associated with smaller declines in consumption at retirement, and one can reject the hypothesis of an equal discontinuity across income replacement quartiles at the 99.9 percent level of confidence in both regressions.[18] At

---

[16]Similar patterns and results are obtained using total wealth. This is true for all of the specifications reported in the appendix; to conserve space, we only include results for financial wealth.

[17]To look for longer-term effects of income and wealth on consumption patterns, we have extended the sample to include data up to 10 years past retirement; these results are shown in Appendix Table A.4. There is some evidence that even those in the high wealth ratio and income replacement quartiles experience a drop in consumption 6-10 years past retirement. However, these results must be treated cautiously given the greater imbalance of the panel.

[18]Notably, these differences remain even when one restricts attention to households that actually engage in non-trivial saving (i.e. those in the top three wealth ratio quartiles).

first, this might appear reconcilable with models of consumption smoothing, since the regression also controls for the household's retirement wealth ratio, thereby effectively ruling out smoothing through the endogenous adjustment of personal saving. However, recall from table 2a that there is virtually no correlation between a household's positions in the wealth ratio and income replacement distributions. Consequently, even when wealth ratio variables are excluded from the regressions in tables 3a and 3b, one still observes significantly larger discontinuities in consumption for households in lower income replacement quartiles (we exclude these estimates to conserve space). As discussed in section II.B, this relation could also be attributable, at least in principle, to sample selection problems. However, as we have noted, it is possible to test for the most natural version of this problem by examining the composition of consumption, which we do in section III.B, below.

Tables 3a and 3b also indicate the presence of some differences across income replacement quartiles in the slope of the consumption profile, both before and after retirement. Generally, prior to retirement, consumption falls more rapidly for households with higher income replacement ratios. This finding is conceivably attributable to the existence of systematic differences in preferences between households with and without pensions. However, the sign of the correlation is difficult to rationalize in this way. The usual argument (e.g. Gale, 1997) is that individuals who discount the future relatively little will tend to save more, and to self-select into jobs with pension coverage. But if this were true, one would expect to observe a *positive* correlation between income replacement rates and the slope of the consumption profile, rather than a negative one. The negative correlation is also puzzling in light of the relationship between income replacement rates and the consumption discontinuity at retirement. Were households reasonably farsighted and concerned about consumption smoothing, one would expect to observe downward adjustments in consumption prior to retirement for those who are likely to find themselves least able to maintain their standards of living after retirement. The data show exactly the opposite pattern.

With respect to the slope of the consumption profile after retirement, the third income replacement quartile differs significantly from the benchmark (first quartile) at the 90 percent level of confidence, but the second and fourth quartiles do not. Thus, while there are some statistically significant differences between quartiles, there is no systematic pattern. We suspect that the coefficient for the third quartile is an anomaly, perhaps driven by one or more outliers; indeed, unlike most of our other results, it is not robust with respect to other natural estimation procedures and variable definitions.

Figure 6 depicts consumption paths for households in each of the four income replacement quartiles (assuming that the household falls into the first wealth ratio quartile), as well as for households in the top wealth ratio and income replacement quartiles. Like figure 5, this figure is based on our more flexible, less parsimonious, non-parametric specification (Appendix Table A.1). As before, we normalize consumption relative to its level in the first and second years before retirement. The figure clearly shows much steeper declines in consumption at retirement for households with lower

Figure 6: Change in Consumption at Retirement, by Income Quartile

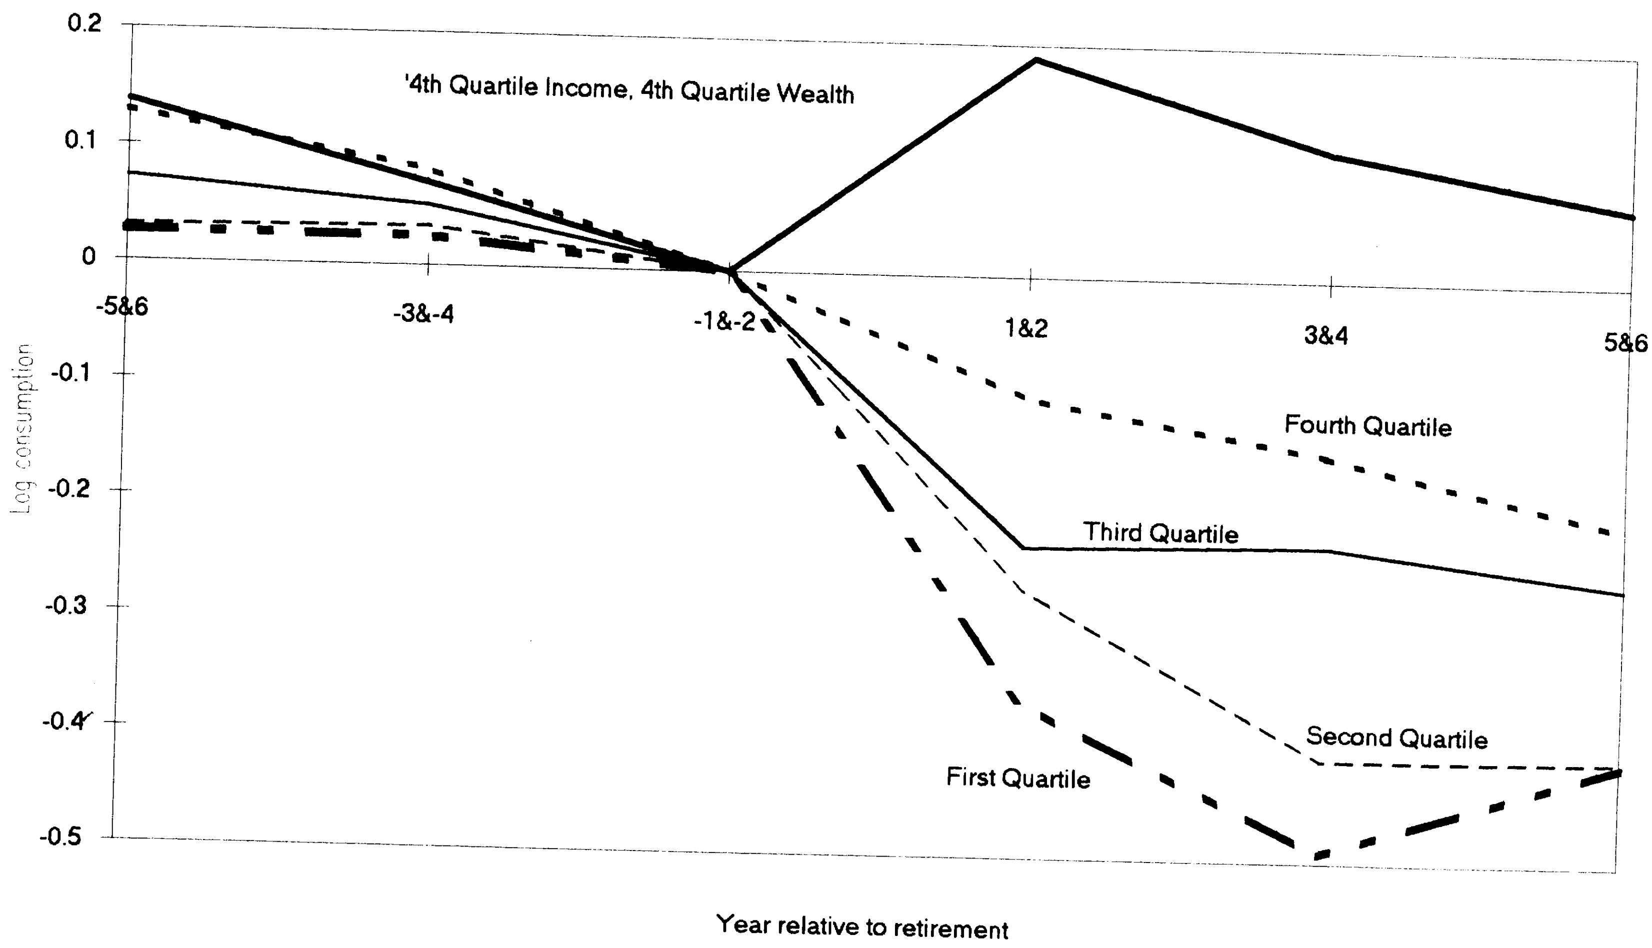

income replacement rates.

B. The composition of consumption

As noted in section II, there are a number of factors that could in principle account for the existence of a consumption discontinuity at retirement, as well as for correlations between the size of this discontinuity and variables such as wealth or income replacement rates. In most cases, theory also implies that these patterns should be present (or at least present to a much greater extent) for particular kinds of expenditure categories (i.e. work-related expenses and substitutes for leisure). Thus, one can distinguish between theories at a more refined level by examining disaggregated measures of consumption.

Although the PSID is not ideally suited for this task, one can disaggregate somewhat by analyzing expenditure patterns separately for food consumed at home and away from home (where the latter category excludes meals consumed at work or school).[19] If, as one would suspect, home cooking is more highly complementary with leisure than are restaurant meals, then, according to various theories (see section II), the patterns noted above should be considerably more pronounced for food consumed away from home.[20]

---

[19]The exclusion of meals at work is notable, since it reduces the likelihood that the observed discontinuity of consumption at retirement is attributable to the disappearance of work-related expenses. Indeed, since retiring works must replace meals at work with meals at or away from home, theory implies that the disappearance of work-related expenses should actually produce an *upward* discontinuity in food consumption, as measured in the PSID. The fact that we find the opposite is therefore particularly telling.

[20]We also estimate consumption profiles for housing expenditures; results appear in Appendix table A.4. Qualitative patterns are similar, although the magnitude of these effects are considerable smaller.

Figure 7 (based on the regression in appendix table A.2) plots normalized consumption trajectories for food consumed away from home. The patterns are broadly similar (although estimated with less statistical precision) to those observed using data on total consumption. There is little or no evidence of differences in consumption growth rates, but there are large and significant differences in the consumption discontinuity at retirement, depending on wealth and replacement income quartiles.

Figure 8 (based on the regression in appendix table A.3) plots normalized consumption trajectories for food consumed at home. The observed trajectories of food consumption at home are generally similar to the trajectories for total consumption and food away from home. Food consumption at home drops most dramatically among those in the lowest income and wealth quartiles. The (log) decline in home food consumption for this group is in fact larger in magnitude than the decline in food consumption away from home. These patterns are difficult to reconcile with the hypothesis that consumption declines discontinuously at retirement because of substitution between market expenditures and leisure.[21]

Since the PSID collected data on relatively few expenditure categories, it is important to determine whether our results are more broadly representative. We therefore

---

[21]There are of course some limitations of this test. It could be that households are switching from high-cost prepared food purchased at the supermarket to less expensive basic foods prepared at home. However, the degree of substitution for low-cost items would have to be quite extreme (and among just the low wealth and income quartiles) to generate the pattern seen in the data. Martin Browning also noted that these results appear to contradict the notion that food at home should be more income inelastic than food away from home. Tabulations of the ratio of food away divided by total food consumption in fact shows some evidence of increasing post-retirement.

Figure 7: Changes At Retirement of Food Away From Home

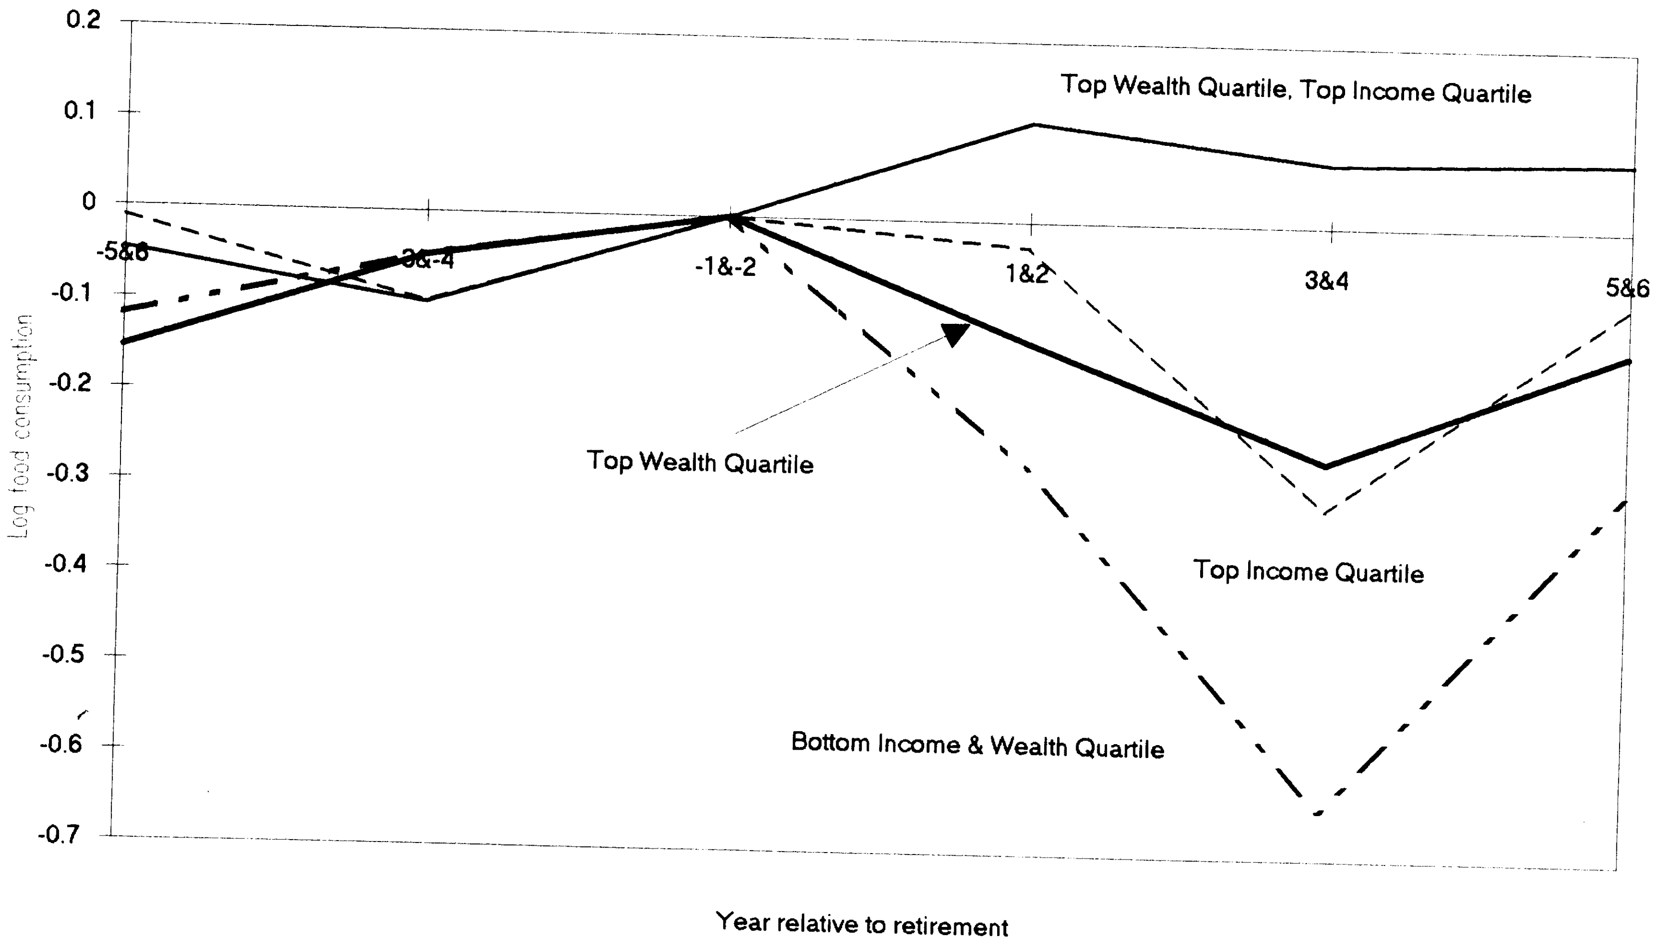

Figure 8: Changes At Retirement of Food at Home

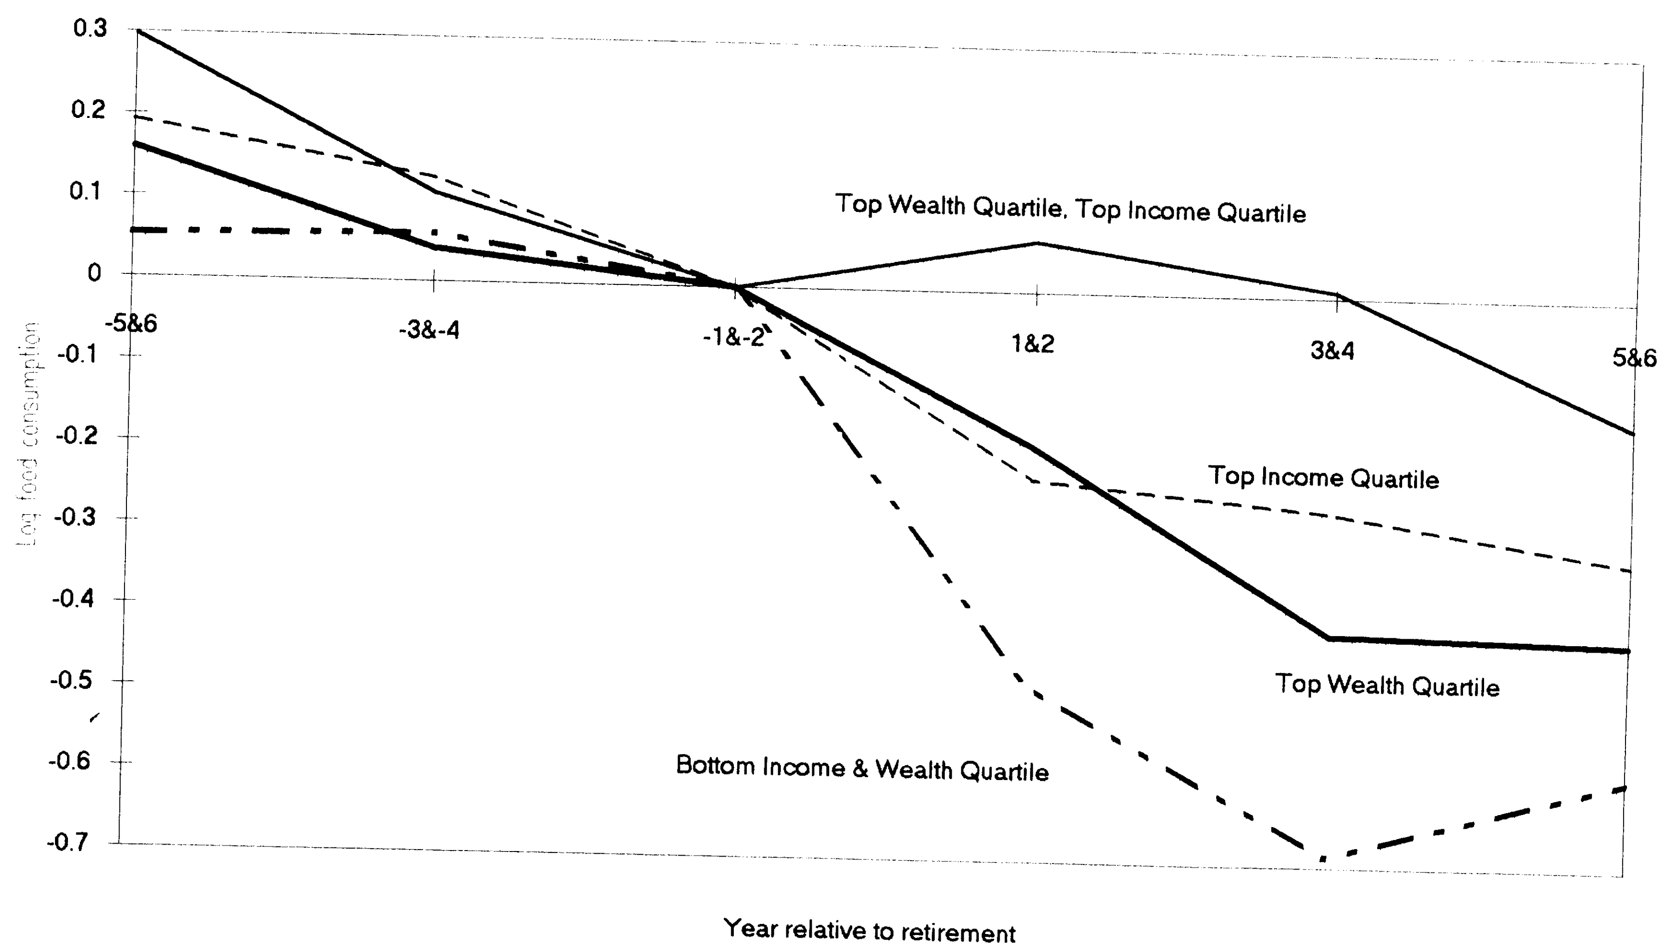

extend our investigation by analyzing Nelson's (1994) extract of the 1980-89 Consumer Expenditure Survey (CEX), which contains merged and annualized quarterly consumption data. We restrict our attention to 1982 through 1989 because of concerns about poor data quality in the start-up years. Observations are dropped if the household's data are missing for any quarter, or if the appropriate fields indicate that income data are of poor quality.

The CEX data are, for practical purposes, a series of cross-sections. The survey does not provide a measure of each household's wealth at retirement, or of changes in income after retirement. We therefore use the conditional commodity demand approach (see e.g. Browning and Meghir, 1991) to study the effect of retirement on relative budget shares. In effect, our object is to determine whether labor market status (here, retirement) has independent predictive power for the budget shares of particular expenditure categories (work- related expenses and leisure complements), controlling for prices and total expenditures. Since we are not interested in estimating the complete budget system, we adopt a simplified version of the Browning-Meghir specification, wherein we simply include dummy variables to allow for variation in prices across years. We would expect that individuals with higher-than-average tastes for work might also have higher-than-average tastes for goods complementary with work; thus we would expect our OLS estimates to be upper bounds on the "true" impact of retirement on commodities such as transportation expenses, adult clothing, and fuel.[22]

Table 4 contains average budget shares and parameter estimates for a simple conditional budget share model. The model explains relative expenditures on several categories of goods that are plausibly complements to working (adult clothing, transportation, fuel, and food away from home), as well as one non-work-related expenditure category that is presumably a complement to leisure (food consumed at home). The model includes dummy variables for age and year (not reported) as well as the demographic variables listed in the table.[23]

A negative partial correlation between retirement and work-related expenses would not by itself indicate that such factors can explain the *variation* in wealth across the population. To conclude that lower wealth households save less because they *anticipate* larger declines in work-related expenses, one would have to find that households in lower wealth quartiles experience *larger* declines in work-related expenses at retirement. To investigate this possibility, we interact the retirement dummy variable with wealth-to-income quartile dummies constructed from the limited wealth data (on

---

[22]Browning and Meghir (1991) used asset income and education as instruments for labor market participation to avoid endogeneity issues of this kind. However, this would be problematic in the current context, since we are also concerned with the separate effect of asset income and education on budget shares near retirement.

[23]We also estimated a model closer in spirit the Almost Ideal Demand System (see e.g. Deaton and Muellbauer, 1980, pp. 75-84), with interaction terms between all of the explanatory variables and the log of total consumption expenditures. Results were similar for the interacted model and the simple model when the simple model included just the retirement dummy variable. However, the model coefficients were less stable when the retirement variable, the wealth quartiles, and the log of consumption expenditures were all interacted.

stocks, bonds, saving bonds, and checking accounts) contained in the Consumer Expenditure Survey.[24] More than one-third of the sample has a missing value for at least one wealth category; all missing wealth values were set to zero. Thus we recognize that these wealth quartiles are likely measured with error.

As is clear from the coefficients in the second through fifth rows of Table 4, retirement is indeed associated with lower budget shares for goods that are complementary to work. While we recognize the perils of interpreting these coefficients as causal, the results suggest that retirement is associated with a reduction in adult clothing purchases of (on average) 0.4 percent of the total consumption budget, a sizeable fraction of the 2.5 percent budget share for all adult clothing.[25] Similarly, retirement is associated with an (average) 2 percentage point fall in total transportation expenditures; of that, the drop in fuel consumption is about 0.7 percentage point. Retirement is also associated with a nearly one percentage point increase in the budget share for food at home, which is consistent with the notion that "home production" substitutes for marketed goods.[26]

---

[24]The wealth quartile dummy variables indicate the household's position within the wealth-to-income distribution after adjusting for differences in retirement status, age, and marital status.

[25]Because each quartile comprises exactly one-quarter of the sample, one can take the arithmetic average of the four coefficients for the retirement interactions to obtain an "average" retirement effect.

[26]While the estimated changes in relative budget shares for food consumed at home and away from home are apparently inconsistent with findings based on panel data from the PSID, it is important to remember that, due to sample selection issues, the use of a cross-section tends to bias upward the estimated reduction in complements to work.

Focusing on *differences* in the impact of retirement on budget shares across wealth groups, there is relatively little variation in the decline of spending on adult clothing (-0.25 percent for the high wealth group and -0.45 for the low wealth group), or on fuel (-0.67 for the high wealth group and -0.70 for the low wealth group). Furthermore, the transportation share declines by *more* among the top wealth group than among the bottom wealth group, which is clearly inconsistent with the hypothesis that wealth varies across households due to the anticipation of differential declines in work-related transportation expenses. Smaller declines in work-related expenses among households with greater wealth are observed only for food consumed away from home. Taken together, these results provide little support for the view that work-related expenses can account for variations in wealth holdings across households.

The final column of table 4 does indicate that food consumed at home increases after retirement, and that this increase is larger for individuals with less wealth. Within the first wealth quartile, there is a 1.9 percentage point increase in the budget share for food consumed at home, compared to a 0.3 percentage point decline for the fourth wealth quartile. This finding is consistent with the view that the variation in wealth is in part attributable to rationally anticipated differences in the extent to which a household is able (or willing) to substitute home production for marketed commodities: those with greater inclinations to rely on home production do not need to save as much. However, the magnitude of the effect is not sufficiently large to explain much of the variation in wealth.

C. Removing the "news" associated with retirement

For the reasons discussed in section II, the findings of sections III.A and III.B are difficult to reconcile with many explanations for the observed variation in wealth based on standard models of life-cycle optimization. This evidence presented so far does not, however, exclude the possibility that the variation in retirement wealth and the associated variation in the size of the consumption discontinuity at retirement are both attributable, at least in part, to unexpected events that affect the timing of retirement.

To examine this possibility, we investigate the manner in which consumption responds to *predictable* events that affect the probability of retirement. In particular, to identify the effect on consumption of predictable retirement, we exploit the fact that the retirement hazard function varies sharply with age (spiking at ages 62 and 65). Our approach is similar in spirit to that of Banks, Blundell, and Tanner (1996), who studied a similar issue using British data.

Formally, we derive empirical consumption profiles through a two step procedure.[27] First, we estimate simple probit models explaining retirement status as a function of education, family size, gender of household head, and marital status as

---

[27]In both stages, we include observations during the transition between full work and retirement; otherwise, our sample is the same as in section III.A.

independent variables.[28] Separate models are estimated for each integer age from 54 through 70, and in each instance we augment the data set to include all households irrespective of whether they retired between 1978 and 1990.[29] The predicted probability of retirement is then introduced in a second-stage consumption regression similar to equation (7), except that $f(\bullet)$ is now specified as a function of the household's retirement probability, rather than of t. Naturally, this fitted probability is interacted with the dummy variables for wealth ratio and income replacement. We bootstrap the entire two-step procedure 1000 times to obtain standard errors.

Table 5 presents estimates of the second stage. The baseline impact of predicted retirement in columns A and B of Table 5 are large in magnitude (-0.74 and -0.78 for estimates based on financial wealth and total wealth, respectively) and highly significant. Furthermore, many of the interactions with wealth ratio and income quartile variables are also large in magnitude and statistically significant. Consequently, the central patterns identified in previous sections are still very much apparent even when we remove the effects of unpredictable events affecting the timing of retirement.

It is of course possible that the probability of retirement may be acting as a proxy for other age-related factors. To evaluate this possibility, we add age to the second stage

---

[28]Neither occupation nor industry appeared to exert a strong effect and so were excluded from the final specification.

[29]An alternative approach would be to estimate a hazard model of retirement, and then use the estimated hazard function to calculate a survival function. The advantage of this approach is that, fixing other characteristics, the unconditional probability of retirement is always higher at older ages. However, in the limit, the two approaches should yield the same predicted probability of retirement.

regression (thereby relying implicitly on the highly non-linear relationship between retirement probabilities and age to identify the effects of interest). Results appear in columns C and D of Table 5. The estimated baseline retirement effect drops (to -0.27 and -0.31, respectively), but the implied effects are still substantial. More importantly, the coefficients of the interaction terms (which indicate the manner in which the consumption discontinuity varies with the household's wealth ratio and income replacement rate) are not much affected by the inclusion of the age variable.

V. Conclusion

In this paper, we have found that consumption growth rates near retirement do not vary systematically with retirement wealth. Thus there is no indication that heterogeneity in pure rates of time preference, subjective survival probabilities, income uncertainty, or tastes for precaution -- all of which should manifest themselves through systematic differences in consumption growth rates -- play a role in determining the distribution of retirement savings. We have found a pronounced discontinuity in consumption at retirement, with the size of the discontinuity negatively correlated with retirement savings and income replacement rates. However, none of these phenomena are confined to work-related expenses or leisure substitutes. The empirical evidence therefore casts doubt on theories that rely on differences in relative tastes for leisure, home production, or work-related expenses to explain the variation in wealth at retirement.

Likewise, differences in retirement wealth for households with similar income and pension profiles (including retirement ages) could in principle result from heterogeneity in planned retirement ages, provided that realized retirement ages are affected to some degree by unanticipated events. While this factor is also consistent with the negative correlation between the consumption discontinuity and retirement savings, it cannot account for the fact that this same pattern remains readily apparent even when we remove the effects of unpredictable events that affect the timing of retirement. Moreover, while variation in the strength of bequest motives may contribute to differences in wealth accumulation, it fails to explain the strong negative correlation between retirement savings and the magnitude of the consumption discontinuity at retirement.

We are also unable to attribute differences in retirement accumulation to variation in the shape of the income/pension profile Households with lower income replacement rates have larger consumption discontinuities at retirement. Contrary to the central tenets of life cycle theory, there is little evidence that households use savings to smooth the effects on consumption of predictable income discontinuities.

These findings are difficult to interpret in the context of the life cycle model. While it may be possible to formulate some model of rational life cycle planning that would account for our findings, in our view, the empirical patterns in this paper are more easily explained if one steps outside the framework of rational, farsighted optimization. If, for example, households follow heuristic rules of thumb to determine saving prior to retirement, and if they take stock of their financial situation and make adjustments at retirement (so that the adequacy of saving is "news"), then one would expect to observe the patterns documented in this paper. A similar conclusion would follow from a theory of "mental accounting," in which individuals differ in the extent to which they can exercise self-discipline over the urge to spend current income.

References

Akerlof, George A., "Procrastination and Obedience," *American Economic Review* 81(2) (May 1991): 1-19.

Attanasio, Orazio, and Guglielmo Weber, "Consumption Growth, the Interest Rate, and Aggregation," Review of Economic Studies 60 (July 1993): 631-49.

_______ and ________, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," Journal of Political Economy 103(6) (December 1995): 1121-1157.

Banks, James, Blundell, Richard, and Sara Tanner, "Is There a Retirement-Savings Puzzle?" mimeo, Institute for Fiscal Studies (July 1996).

Banks, James, and Blundell, Richard, "Is There a Retirement-Savings Puzzle?" mimeo, Institute for Fiscal Studies (1993).

Baxter, Marianne, and Urban Jermann (1995), "Household Production and the Excess Sensitivity of Consumption to Current Income," mimeo, University of Virginia.

Bernheim, B. Douglas, "How Strong Are Bequest Motives? Evidence Based on Estimates of the Demand for Life Insurance and Annuities," *Journal of Political Economy* 99 (5) (October 1991): 899-927.

Bernheim, B. Douglas, "Do Households Appreciate their Financial Vulnerabilities? An Analysis of Actions, Perceptions, and Public Policy," Tax Policy and Economic Growth, Washington DC: American Council for Capital Formation (1995): pp. 1-30.

Bernheim, B. Douglas, Robert Lemke, and John Karl Scholz, "U.S. Household Saving in the 1980s: Evidence from the Surveys of Consumer Finances," mimeo, Stanford University (1997).

Browning, Martin, and Costas Meghir, "The Effects of Male and Female Labor Supply on Commodity Demands," Econometrica 59(4) (July, 1991): 925-51.

Butler, Monika, "A Neoclassical Exploration of Consumption and Income Profiles over the Life-Cycle," mimeo, University of Virginia (November 1996).

Carroll, Chris, "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," Quarterly Journal of Economics 92(1) (February 1997): 1-56.

Congressional Budget Office, Baby Boomers in Retirement: An Early Perspective. Washington DC: U.S. Government Printing Office (1993).

Deaton, Angus (1991) "Saving And Liquidity Constraints," Econometrica, v59(5), 1221-1248.

Deaton, Angus, and John Muellbauer, Economics and Consumer Behavior (New York: Cambridge University Press, 1980).

Diamond, Peter A., and Jerry A. Hausman, "The Retirement and Unemployment Behavior of Older Men," in Henry A. Aaron and Gary Burtless (eds.) Retirement and Economic Behavior Washington DC: The Brookings Institution (1984).

Dynan, Karen, "How Prudent are Consumers?" Journal of Political Economy 101(6) (December 1993): 1104-1113.

Dynan, Karen, Jonathan Skinner, and Stephen Zeldes, "Do the Rich Save More?" mimeo, Columbia University (1997).

Engen, Eric M., William G. Gale, and John Karl Scholz, "The Illusory Effects of Saving Incentives on Saving," *Journal of Economic Perspectives* 10(4) (Fall 1996), 113-138.

Gale, William G., "The Effects of Pensions on Household Wealth: A Re-evaluation of Theory and Evidence," mimeo, Brookings Institution (1997).

Gustman, Alan L., and Thomas Steinmeier, "Partial Retirement and the Analysis of Retirement Behavior," Industrial and Labor Relations Review 37(3) (April 1984): 403-415.

Ghez, Gilbert, and Gary Becker, The Allocation of Time and Goods Over the Life Cycle. New York: Columbia University Press (1975).

Hall, Robert, "Intertemporal Substitution in Consumption," Journal of Political Economy 96(2) (April 1988): 339-357.

Hamermesh, Daniel S., "Consumption During Retirement: The Missing Link in the Life Cycle," Review of Economics and Statistics 66(1) (February 1984): 1-7..

Hausman, Jerry A., and Lynn Paquette, "Involuntary Early Retirement and Consumption," in Gary Burtless (ed.) Work, Health, and Income Among the Elderly (Washington DC: The Brookings Institution, 1987).

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* 40 (June 1994).

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," Journal of Political Economy, Volume 103, No. 2, April 1995, 360-399.

Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers, "The Adequacy of Saving," American Economic Review 72(5) (December 1982): 1056-1069.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," Quarterly Journal of Economics 112 (May 1997): pp. 443-478.

Laitner, John, and F. Thomas Juster, "New Evidence on Altruism: A Study of TIAA-CREF Retirees," American Economic Review , September 1996, 86(4), 893-908.

Lazear, Edward P., "Some Thoughts on Saving," in D. A. Wise (ed.), Studies in the Economics of Aging Chicago: University of Chicago Press and NBER (1994).

Mariger, Randall P., "A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results," Econometrica 55(3) (May 1987): 533-557.

Nelson, Julie A., Consumer Expenditure Surveys, 1980-89: Interview Surveys, for Household-Level Analysis (computer file, 1994), Washington, DC: United States Department of Labor, Bureau of Labor Statistics (producer), Ann Arbor, MI: Inter-university Consortium for Political and Social Research (distributor).

Palumbo, Michael (1996), "Health Uncertainty and Precautionary Saving Near the End of the Life Cycle," working paper.

Poterba, James M., Steven F. Venti, and David A. Wise, "How Retirement Saving Programs Increase Saving" *Journal of Economic Perspectives* 10(4) (Fall 1996), 91-112.

Robb, A.L., and J.B. Burbidge, "Consumption, Income, and Retirement," Canadian Journal of Economics 22(3) (August 1989): 522-542.

Skinner, Jonathan, "A Superior Measure of Consumption from the Panel Study of Income Dynamics," Economic Letters (1987).

Thaler, Richard H., "Psychology and Savings Policies," *American Economic Review* 84(2) (May 1994): 186-192.

Thaler, Richard H., and Hersh M. Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy* 89(2) (April 1981): 392-406.

Venti, Steven, and David Wise, "Lifetime Income, Saving Choices, and Wealth at Retirement," mimeo, Dartmouth College (1996).

Table 1 : Summary Statistics

Variable Names	Mean	Median	Standard Deviation	Minimum	Maximum
Age	62.1	62	5.4	43	80
Income Replacement Rate	0.64	.60	0.32	.024	2.87
Financial Wealth-income ratio	4.31	2.53	7.43	-6.10	72.09
Total Wealth-income ratio	5.96	4.07	8.51	-5.95	84.49
Log Consumption (1984 dollars)	9.57	9.63	5.19	5.45	11.50
Family Size (does not include spouse)	1.76	1	1.28	1	14

Table 2a: Proportional Distribution of Sample by Income Replacement and Wealth Quartile

	Income Q1	Income Q2	Income Q3	Income Q4	Total
Wealth Q1	6.2	5.4	6.9	6.6	25.1
Wealth Q2	4.5	9.2	5.0	6.1	24.9
Wealth Q3	6.6	5.2	7.8	5.4	25.1
Wealth Q4	7.8	5.2	5.0	6.9	24.9
Total	25.1	25.1	24.6	25.1	100.0

Table 2b: Average Age at Retirement by Income Replacement and Wealth Quartile

	Income Q1	Income Q2	Income Q3	Income Q4	Total
Wealth Q1	59.6	60.1	60.6	60.7	60.3
Wealth Q2	59.7	61.6	61.2	61.1	61.0
Wealth Q3	60.1	61.6	62.0	60.0	61.0
Wealth Q4	60.0	61.0	60.5	61.8	60.8
Total	59.9	61.2	61.1	61.0	60.8

Table 2c: Median Pre-Retirement Income by Income Replacement Rate and Wealth Quartile

	Income Q1	Income Q2	Income Q3	Income Q4	Total
Wealth Q1	25,142	23,058	27,512	23,032	24,781
Wealth Q2	24,817	27,214	30,517	30,744	28,315
Wealth Q3	38,993	33,421	32,513	28,584	33,560
Wealth Q4	39,560	34,666	34,148	24,602	33,330
Total	33,240	29,147	31,046	26,558	29,993

Table 3a: Consumption Growth Pre- and Post-Retirement (Financial Wealth Ratios)

	Pre-retirement log consumption growth	Retirement "jump" in log consumption	Post-retirement log consumption growth
Benchmark (Quartile 1 for Income and Wealth Ratios)	-0.007 (0.7)	-0.322 (5.7)	-0.030 (2.5)
Financial Wealth Ratio Quartile			
Quartile 2	-0.015	-0.165**	-0.015
Quartile 3	-0.007	-0.150**	-0.023
Quartile 4	-0.008	-0.040**	-0.030
P-value for Equality of Coefficients F(3,3432)	0.76	0.00	0.58
Income Replacement Quartile			
Quartile 2	-0.007	-0.228	-0.039
Quartile 3	-0.020	-0.186*	-0.008
Quartile 4	-0.030*	-0.019**	-0.030
P-value for Equality of Coefficients F(3,3432)	0.06	0.00	0.07

N = 3881, 422 households, each with individual constant terms. R^2=.76. Dependent variable is log of consumption. Absolute value of t-statistics (for benchmark) in parentheses; * denotes hypothesis that the quartile growth rate (or jump) differs from the benchmark is rejected at the 5% level in a two-tailed test; ** denotes that the hypothesis is rejected at the 1% level. Additional variables: family size (coeff. = 0.04, t-stat. = 4.7); marital status (coeff. = 0.06, t-stat.= 3.4); disability (coeff. = -0.03, t-stat. = 2.3); female widower (coeff. = -0.15, t-stat. = 4.2); and a dummy variable for whether the household was working part-time for 3-4 years prior to full retirement (coeff. = -0.03, t-stat. = 1.2).

Table 3b: Consumption Growth Pre- and Post-Retirement (Total Wealth Ratios)

	Pre-retirement log consumption growth	Retirement "jump" in log consumption	Post-retirement log consumption growth
Benchmark (Quartile 1 for Income and Wealth Ratios)	-0.005 (0.5)	-0.309 (5.4)	-0.040 (3.3)
Wealth Ratio Quartile			
Quartile 2	-0.026*	-0.180*	-0.011*
Quartile 3	-0.001	-0.121**	-0.029
Quartile 4	-0.005	-0.074**	-0.024
P-value for Equality of Coefficients F(3,3432)	0.06	0.00	0.11
Income Replacement Quartile			
Quartile 2	-0.005	-0.213	-0.046
Quartile 3	-0.020	-0.183*	-0.016
Quartile 4	-0.029*	-0.002**	-0.039
P-value for Equality of Coefficients F(3,3432)	0.04	0.00	0.08

N = 3881, 422 households, each with individual constant terms. R^2=.76. Dependent variable is log of consumption. Absolute value of t-statistics (for benchmark) in parentheses; * denotes hypothesis that the quartile growth rate (or jump) differs from the benchmark is rejected at the 5% level in a two-tailed test; ** denotes that the hypothesis is rejected at the 1% level. Additional variables: family size (coeff. = 0.04, t-stat. = 4.8); marital status (coeff. = 0.06, t-stat.= 3.5); disability (coeff. = -0.03, t-stat. = 2.2); female widower (coeff. = -0.15, t-stat. = 4.2); and a dummy variable for whether the household was working part-time for 3-4 years prior to full retirement (coeff. = -0.02, t-stat. = 0.8).

Table 4: Conditional Budget Share Regressions, Consumer Expenditure Survey, 1982-89

Dependent Variable	Adult Clothing	Transpor-tation	Fuel	Food Away from Home	Food at Home
Average Budget Share (Percent)	2.50	14.85	4.39	4.22	14.63
Retired (yes = 1) x Wealth Quartile 1	-0.448 (6.1)	-1.665 (4.6)	-0.704 (6.4)	-0.982 (7.0)	1.893 (9.4)
Retired (yes = 1) x Wealth Quartile 2	-0.539 (6.5)	-1.634 (4.0)	-0.930 (7.5)	-1.306 (8.2)	1.367 (6.0)
Retired (yes = 1) x Wealth Quartile 3	-0.227 (1.9)	-1.506 (2.6)	-0.487 (2.8)	-0.534 (2.4)	0.040 (0.1)
Retired (yes = 1) x Wealth Quartile 4	-0.243 (3.0)	-2.827 (7.1)	-0.665 (5.5)	-0.229 (1.5)	-0.301 (1.4)
Married	0.418 (5.5)	-2.232 (5.9)	0.464 (4.1)	-3.144 (21.6)	3.797 (18.2)
Single Female Head	1.110 (14.5)	-2.771 (7.3)	-1.369 (12.0)	-3.300 (22.6)	-0.734 (3.5)
Any Children ≤ 18 yrs? (yes = 1)	-0.204 (1.5)	1.315 (2.0)	0.389 (2.0)	-0.649 (2.5)	1.685 (4.6)
Number of Children	-0.229 (3.2)	-0.822 (2.3)	0.084 (0.8)	-0.347 (2.5)	1.783 (9.1)
Log (Total Expenditures)	0.843 (19.0)	8.455 (39.4)	-0.475 (7.1)	1.341 (16.2)	-8.094 (66.45)
R^2	.12	.26	.15	.12	.37

Notes: Five-year dummy variables for age and individual year dummy variables included in all regressions. The sample size is 10,260. All dependent variables are percentages of total expenditures.

Table 5: Second Stage Consumption Regressions

Variable Name	A Financial Wealth	B Total Wealth	C Financial Wealth	D Total Wealth
Probability of Retirement	-0.738 (7.7)	-0.779 (7.9)	-0.274 (2.8)	-0.312 (3.1)
Wealth Quartile 2 x P(Ret)	0.020 (0.3)	0.051 (0.7)	0.073 (0.9)	0.099 (1.3)
Wealth Quartile 3 x P(Ret)	0.160 (1.7)	0.256 (2.7)	0.212 (2.2)	0.284 (3.0)
Wealth Quartile 4 x P(Ret)	0.177 (1.7)	0.266 (2.5)	0.195 (1.9)	0.272 (2.6)
Income Quartile 2 x P(Ret)	0.178 (1.6)	0.174 (1.6)	0.157 (1.5)	0.151 (1.4)
Income Quartile 3 x P(Ret)	0.351 (3.5)	0.334 (3.4)	0.331 (3.4)	0.312 (3.2)
Income Quartile 4 x P(Ret)	0.318 (3.0)	0.293 (2.9)	0.298 (2.9)	0.277 (2.8)
Family Size	0.078 (7.9)	0.078 (8.0)	0.074 (7.6)	0.074 (7.5)
Marital Status	0.176 (9.7)	0.174 (9.7)	0.130 (6.9)	0.129 (6.9)
Female Head	-0.153 (3.4)	-0.158 (3.5)	-0.143 (3.2)	-0.148 (3.3)
Age			-0.024 (7.7)	-0.024 (7.6)
Disability	-0.049 (3.4)	-0.049 (3.3)	-0.040 (2.9)	-0.040 (2.7)
R^2	.77	.77	.77	.78

N = 4742. Dependent variable is log of consumption. Household-specific effects included. First stage probits are run separately for each age group; all z-statistics (in parentheses) are based on bootstrapped standard errors.

Table A.1: Regression for Total Consumption

		Interaction terms with year relative to retirement				
		-5 and -6	-3 and -2	+1 and +2	+3 and +4	+5 and +6
Benchmark (1st Quartile, Wealth and Income)		0.026 (0.6)	0.024 (0.6)	-0.365 (8.6)	-0.493 (10.7)	-0.413 (8.1)
Increment - 2nd Wealth Quartile		0.040 (0.9)	-0.042 (1.0)	0.149 (3.4)	0.178 (3.7)	0.198 (3.9)
Increment - 3rd Wealth Quartile		0.011 (0.3)	-0.003 (0.1)	0.182 (4.1)	0.226 (4.8)	0.189 (3.6)
Increment - 4th Wealth Quartile		0.010 (0.2)	-0.011 (0.3)	0.294 (6.6)	0.261 (5.4)	0.276 (5.3)
Increment - 2nd Income Quartile		0.005 (0.1)	0.010 (0.2)	0.097 (2.1)	0.081 (1.7)	0.003 (0.1)
Increment - 3rd Income Quartile		0.047 (1.1)	0.028 (0.7)	0.133 (3.0)	0.266 (5.5)	0.153 (2.8)
Increment - 4th Income Quartile		0.104 (2.4)	0.057 (1.4)	0.262 (5.9)	0.343 (6.7)	0.203 (3.8)
Part-time work		-0.054 (1.3)	-0.069 (1.7)	-0.075 (1.7)	-0.079 (1.5)	-0.068 (1.3)
Family size	0.038 (4.8)					
Female Head (widowhood)	-0.146 (4.2)					
Married	0.063 (3.4)					
Disabled	-0.034 (2.3)					

Sample size = 3881. Estimates correct for household-specific effects. $R^2 = .76$. Dependent variable is log total consumption. Absolute value of t-statistics in parentheses. Shaded area shows benchmark coefficients for bottom wealth and income quartiles. The "increment" for each wealth ratio and income replacement quartile measures log consumption *relative* to the benchmark group.

Table A.2: Regression for Food Consumption Away from Home

		Interaction terms with year relative to retirement				
		-5 and -6	-3 and -2	+1 and +2	+3 and +4	+5 and +6
Benchmark (1st Quartile, Wealth and Income)		-0.120 (1.0)	-0.048` (0.4)	-0.272 (2.1)	-0.645 (4.7)	-0.294 (1.9)
Increment - 2nd Wealth Quartile		-0.072 (0.6)	-0.026 (0.2)	0.034 (0.3)	0.242 (1.7)	0.010 (0.1)
Increment - 3rd Wealth Quartile		0.074 (0.6)	0.015 (0.1)	0.170 (1.3)	0.541 (4.1)	0.410 (2.7)
Increment - 4th Wealth Quartile		-0.035 (0.3)	-0.001 (0.0)	0.140 (1.1)	0.383 (2.8)	0.160 (1.1)
Increment - 2nd Income Quartile		0.066 (0.6)	0.004 (0.0)	0.027 (0.2)	0.077 (0.6)	0.016 (0.1)
Increment - 3rd Income Quartile		0.166 (1.5)	0.029 (0.3)	0.029 (0.2)	0.332 (2.6)	-0.077 (0.5)
Increment - 4th Income Quartile		0.109 (0.9)	-0.052 (0.5)	0.244 (2.0)	0.332 (2.4)	0.211 (1.4)
Part-time work		0.041 (0.4)	-0.117 (1.0)	-0.036 (0.3)	0.056 (0.4)	-0.071 (0.4)
Family size	-0.019 (0.8)					
Female Head (widowhood)	-0.214 (2.2)					
Married	-0.016 (0.3)					
Disabled	0.006 (0.2)					

Sample size = 3104. Estimates correct for household-specific effects. $R^2 = .56$. Dependent variable is log of food consumption away from home. Absolute value of t-statistics in parentheses. Shaded area shows benchmark coefficients for bottom wealth and income quartiles. The "increment" for each wealth ratio and income replacement quartile measures log consumption *relative* to the benchmark group.

Table A.3: Regression for Food Consumption at Home

		Interaction terms with year relative to retirement				
		-5 and -6	-3 and -2	+1 and +2	+3 and +4	+5 and +6
Benchmark (1st Quartile, Wealth and Income)		0.054 (1.0)	0.059 (1.1)	-0.484 (8.6)	-0.689 (11.4)	-0.593 (8.6)
Increment - 2nd Wealth Quartile		0.115 (2.00)	-0.004 (0.1)	0.185 (3.2)	0.158 (2.5)	0.086 (1.3)
Increment - 3rd Wealth Quartile		0.014 (0.2)	-0.039 (0.7)	0.270 (4.7)	0.227 (3.7)	0.108 (1.6)
Increment - 4th Wealth Quartile		0.105 (1.9)	-0.019 (0.4)	0.294 (5.0)	0.274 (4.3)	0.169 (2.5)
Increment - 2nd Income Quartile		0.058 (1.0)	0.003 (0.1)	0.081 (1.4)	0.181 (2.8)	0.086 (1.2)
Increment - 3rd Income Quartile		0.119 (2.1)	0.074 (1.3)	0.215 (3.6)	0.373 (5.9)	0.265 (3.7)
Increment - 4th Income Quartile		0.139 (2.5)	0.068 (1.2)	0.253 (4.3)	0.424 (6.3)	0.268 (3.8)
Part-time work		-0.042 (0.8)	-0.072 (1.4)	0.002 (0.0)	-0.002 (0.0)	-0.060 (0.8)
Family size	0.083 (7.9)					
Female Head	-0.193 (4.2)					
Married	0.109 (4.5)					
Disabled	-0.034 (1.8)					

Sample size = 3874. Estimates correct for household-specific effects. $R^2 = .62$. Dependent variable is log food consumption at home. Absolute value of t-statistics in parentheses. Shaded area shows benchmark coefficients for bottom wealth and income quartiles. The "increment" for each wealth ratio and income replacement quartile measures log consumption *relative* to the benchmark group.

Table A.4: Additional Estimation Results: Extended Panel and Housing Consumption

	Sample Extended to 10 Years After Retirement			Housing Consumption (rental payments or imputed housing flows)		
	Bottom Wealth and Income Quartile	Top Wealth Quartile	Top Income Quartile	Bottom Wealth and Income Quartile	Top Wealth Quartile	Top Income Quartile
5&6 yrs pre-retirement	0.018 (0.4)	0.032	0.124*	0.021 (0.4)	0.042	0.088
3&4 yrs pre-retirement	0.020 (0.5)	0.011	0.079	0.078 (1.3)	0.095	0.088
1&2 yrs pre-retirement are the excluded dummy variables						
1&2 yrs post-retirement	-0.361 (8.2)	-0.069**	-0.102**	-0.204 (3.3)	0.010**	0.059**
3&4 yrs post-retirement	-0.488 (10.4)	-0.229**	-0.147**	-0.293 (4.4)	-0.161	0.061**
5&6 yrs post-retirement	-0.410 (7.9)	-0.139**	-0.198**	-0.175 (2.3)	-0.050	0.003*
7&8 yrs post-retirement	-0.403 (6.6)	-0.255*	-0.185**			
9&10 yrs post-retirement	-0.417 (5.7)	-0.248*	-0.241			

Notes: * denotes coefficient is different from the benchmark Quartile 1 wealth and income coefficient at the .05 significance level in a two-tailed test; ** denotes difference from the benchmark at the .01 significance level.