ELSEVIER

Journal of Economic Behavior & Organization
Vol. 36 (1998) 59–83

JOURNAL OF
Economic Behavior
& Organization

Are assets fungible? Testing the behavioral theory of life-cycle savings

Laurence Levin*

Cornerstone Research Suite 250, 1000 El Camino Real, Menlo Park, CA 94025, USA

Received 7 February 1994; received in revised form 14 January 1998; accepted 11 February 1998

Abstract

This paper is an empirical investigation of the behavioral life-cycle savings model. This model posits that self-control problems causes individuals to depart substantially from rational behavior. I show that this model can explain how the consumption of individuals at or near retirement vary with changes in different types of financial assets. Specifically, consumption spending is sensitive to changes in income and in liquid assets, but not very sensitive to changes in the value of other types of assets such as houses and social security (even though the value of non-liquid assets is relatively large for most of the households in the sample). In general, the evidence presented here favors the Behavioral Life-Cycle Model over the conventional life-cycle model even when liquidity constraints are introduced.

JEL classification: D91: Intertemporal consumer choice; Life-cycle models and savings; E21: Consumption; Savings

Keywords: Savings; Life-cycle; Behavioral models

1. Introduction

According to the simplest versions of the life-cycle theory, a person's consumption should depend only on the present value of his wealth. This theory has many implications, the most studied being that individuals should engage in consumption smoothing. However, this implication has been rejected by most empirical studies.[1] Often the failure

* Corresponding author. Tel.: +1 415 239 8224; e-mail: 1mlevin@usa.net

[1] For e.g., see papers by Flavin (1981), Mankiw (1981), Hall and Mishkin (1982), Hansen and Singleton (1983), Courant et al. (1986), Carroll and Summers (1989) and Campbell and Mankiw (1989). An excellent summary of empirical work on the life-cycle hypothesis can be found in Thaler (1990).

PII S0167-2681(98)00070-5

of the life-cycle model is attributed to the presence of liquidity constraints for significant subsets of the population (see Hayashi (1985) or Zeldes (1989) for examples). This paper examines whether these anomalous results can be better explained by an alternative model of savings behavior that takes into account the problem of self-control. This model is called the behavioral life-cycle model, which was formulated by Shefrin and Thaler (1988).[2]

To compare these two models, I examine how consumption varies with both the level of wealth and the form that wealth takes. In the simplest life-cycle models, consumption should only depend on an individual's total wealth, in other words, assets should be fungible; receiving $ 100 today should have as much effect on an individual's spending as receiving $ 100+interest next year, or a $ 100 appreciation of the value of his house.[3] The behavioral life-cycle model predicts that assets should not be fungible, implying that an individual's consumption decisions will be affected by asset composition as well as total wealth.

The behavioral life-cycle model as developed by Shefrin and Thaler is a simple model of self-control based on three ideas. First, individuals are tempted to spend all their resources on current consumption instead of saving for the future. Second, individuals who save, overcome this self-control problem by investing in a variety of assets that have different levels of temptation associated with them. A historic example of this type of asset was Christmas clubs which were savings plans (usually with low or non-existent interest rates) that did not allow withdrawal unto December. Individuals invested in them to insure they had enough money saved for Christmas presents. Thus, individuals create mental accounts for their different assets causing their marginal propensities to consume from those assets to vary with the level of temptation associated with each one. Third, setting up these mental accounts implies that individuals engage in 'framing'; a person's consumption spending not only depends on total wealth but also depends on how that wealth is allocated among assets with differing levels of 'temptation.' For example, individuals are more willing to spend assets they have labeled current income than those they have labeled wealth or those that they expect in the future. Essentially, the behavioral life-cycle posits that there are psychological as well as financial transaction costs associated with spending from different types of assets.

This paper contains the first formal empirical tests of the behavioral life-cycle model using a large panel dataset such as those that have been used to study other models of savings. Previous evidence for the behavioral life-cycle (collected in Thaler, 1992) has either come from small surveys of college or MBA students or been garnered from anomalies found in other studies that were not designed to test the behavioral life-cycle model. Although this type of evidence is valuable, it cannot be definitive because of small sample size or possible biases in studies that were not designed to study behavioral decision making. Testing the behavioral life-cycle model necessitates taking the assertions of that model and creating empirical tests that distinguish between it and

---

[2] This theory can be interpreted as a formalization of an older literature which had rejected some of the implications of permanent income/life-cycle models but found that wealth was still an important element in making consumption decisions. This literature is summarized in Mayer (1972).

[3] Liquidity constraints and other extensions of the life-cycle model may also cause fungibility to be rejected. See the discussion in Section 2.

conventional models of life-cycle savings. It also requires building an econometric model that controls for the biases inherent in examining consumption data.

Another novel feature of this paper is that many goods are used to study consumption. In previous empirical studies of consumption only food spending is used because it is the only category of consumption included in the PSID and other commonly used datasets. In this paper, I use the Retirement History Survey which contains information on how much is spent on ten different goods. This creates an opportunity to investigate how spending patterns differ between goods.

In the next section, I derive four testable differences between the behavioral and conventional life-cycle models. The first two tests are straightforward comparisons of the marginal propensity to consume (MPC) out of different types of assets. The third test examines whether liquidity constraints are causing the observed differences in assets' MPC and is derived from differences between the two models on the effect of liquidity constraints. Essentially this test examines whether liquidity constraints occur because of an inability to liquify certain assets or because of large psychological or financial transaction costs in doing so. The final test is an examination of whether transaction costs seem to be financial or psychological in nature. This test is derived by looking for evidence of framing in consumption. For a conventional life-cycle model individual, the propensity to consume a particular good out of any asset will depend only on the good's wealth elasticity and possible financial transaction costs. However, if an individual acts according to the behavioral life-cycle model, then the form wealth takes affects the type of wealth the individual consumes. In other words, the psychological cost of using an asset to finance the consumption of a particular good is a function of that good's attributes. In this case the propensity to consume a particular good out of a particular asset will depend on the attributes of both the good and the asset.[4]

In Section 3, the data and the statistical methods are discussed. Section 4 contains the empirical results, and Section 5 is the conclusion.

The results of this paper provide strong evidence against assets being fungible. For many individuals consumption does not seem to be very sensitive to changes in their wealth portfolios. Further, assets' marginal propensities to consume are very different and show strong evidence in favor of the presence of large transaction costs. Finally, the behavior of liquidity constrained individuals is consistent with the expected effects from framing. In general then, this paper finds strong evidence in favor of the behavioral life-cycle model in comparison to the conventional life-cycle model.

2. Modelling consumption expenditures

2.1. Framework

In this section a simple framework is introduced in order to formalize tests of the differences between the behavioral and conventional life-cycle models. To establish the

[4] Although this implication of the behavioral life-cycle theory is not explicitly discussed by Shefrin and Thaler (1988), Shefrin and Thaler have derived this result in an unpublished mimeo, Shefrin and Thaler (1981).

framework, compare the expenditure function of a conventional life-cycle individual (E_{L}) to one who acts in a behavioral manner (E_{B}), both with income (Y) and K different types of assets (A_1-A_K).

In the simplest model, a conventional life-cycle individual's expenditures will only be a function of his total resources. The individual's total resources (W) are the sum of his assets plus the capitalized value of his permanent income which can be written as αY where α is

$$\alpha = \sum_{t=1}^{T} \frac{1 \cdot P_{\mathrm{t}}}{(1+r)^t} \tag{1}$$

where P_{t}=survival probability, r=interest rate, T=time horizon.

Thus the total resources of an individual are: $W = \alpha Y + \Sigma A_i$. Therefore, the conventional life-cycle individual's expenditure function for the gth good can be written as

$$E_{\mathrm{L}g} = E_{\mathrm{L}g}(W) = E_{\mathrm{L}g}\left(\alpha Y + \sum_{k=1}^{K} A_k\right) \tag{2}$$

For a behavioral life-cycle individual, consumption will depend not only on his total resources but also on how those resources are split up between different types of assets. Thus the expenditure function for a behavioral life-cycle individual will be

$$E_{\mathrm{B}g} = E_{\mathrm{B}g}(Y, A_1, A_2, \ldots, A_k) \tag{3}$$

In this case, there is no single aggregate that an individual uses when making an expenditure decision. From these expenditure functions one can easily derive the marginal propensity to consume (MPC) out of a given asset. For a conventional life-cycle saver the marginal propensity to consume out of any asset (i or j) is

$$\frac{\partial E_{\mathrm{L}g}}{\partial A_i} = \frac{\partial E_{\mathrm{L}g}}{\partial A_j} = \frac{\partial E_{\mathrm{L}g}}{\partial W} \quad \text{and} \quad \frac{\partial E_{\mathrm{L}g}}{\partial Y} = \alpha \frac{\partial E_{\mathrm{L}g}}{\partial W} \tag{4}$$

For income this formulation assumes that income is permanent income. If changes in income were completely transitory then the MPC of income would equal the MPC of wealth. Thus, the value of the MPC from an addition to income relative to the MPC from an addition to wealth will range between the MPC of wealth (if the income change is transitory) and (if the income change is permanent) and will depend on how permanent the change in income is. Note that the MPC from income will never be greater than α! For a behavioral saver the marginal propensity to consume out of an asset will be

$$\frac{\partial E_{\mathrm{B}g}}{\partial A_i} > \frac{\partial E_{\mathrm{B}g}}{\partial A_j} > \cdots > \frac{\partial E_{\mathrm{B}g}}{\partial A_K} \quad \text{and} \quad \frac{\partial E_{\mathrm{B}g}}{\partial Y} > \alpha \frac{\partial E_{\mathrm{B}g}}{A_1} \tag{5}$$

where the assets are arranged in descending orders of temptation or liquidity (A_1 is more tempting than A_2 etc.); further, even if Y is transitory, its high temptation level will cause its MPC to be higher then the MPC of an equivalent change in wealth.

2.2. Test 1 (Comparing the MPC from income to the MPC from wealth)

The first test derives from comparing Eq. (4) to Eq. (5). This test consists of examining the sensitivity of consumption to income relative to wealth. According to behavioral life-cycle theory, spending should be very sensitive to income and much less sensitive to wealth. Thus if the behavioral theory is right, a regression of consumption against income and other assets should find that income is usually significantly related to consumption, while other assets may not be. Further, if the conventional theory holds, the ratio of the MPC of permanent income to the MPC of wealth should equal, the capitalized value of a $ 1 worth of permanent income; if the behavioral theory is correct this ratio will be larger than.

2.3. Test 2 (Testing for the equality of different assets' MPC)

The second test also directly follows from the framework developed above. This test compares the MPCs of different assets. If individuals are conventional life-cycle savers then the MPCs of the different assets should be equal. In other words assets should be perfect substitutes. However, if individuals are behavioral life-cycle savers then the MPC of different assets will not be the same and should, instead, be a function of the level of temptation associated with the particular asset. Thus, this test consists of examining whether or not the MPCs are equal.

2.4. Test 3 (Testing for liquidity constraints)

Obviously there are explanations other than the behavioral life-cycle hypothesis that predict Tests 1 and 2 would fail. The addition of liquidity constraints to the simplest convention would, by themselves, cause assets to not be fungible. If individuals cannot borrow against some of their assets then even the conventional life-cycle model will predict that less liquid assets will have lower marginal propensities to consume. The third test compares the MPC of different asset types for constrained and unconstrained individuals. The idea behind this test is that in a conventional life-cycle model liquidity constraints are externally imposed by market imperfections which limit the ability to borrow against future income whereas in the behavioral life-cycle model, liquidity constraints may also be internally imposed by the consumer. The testable difference between internal and external liquidity constraints is that if the constraint is external, the value of an illiquid asset will affect consumption as long as other more liquid assets are still being held (since an increase in this asset just adds to general wealth) but will cease to affect consumption when those more liquid assets are depleted, since then the liquidity constraint will be binding. The opposite result holds when the liquidity constraint is internally imposed. As long as more tempting assets have not been exhausted, the value of the less tempting asset will not affect the decision on how much to consume. Instead a less tempting asset will be consumed only after more tempting assets are exhausted.

To formalize this test, assume there are two assets denoted A_1 and A_2. A_1 is a liquid asset (for example, the value of a checking account) while A_2 is an asset that is less liquid (for example, a house) that will eventually be liquidated. For a conventional life-cycle individual, the value of A_2 will only affect consumption when the liquidity constraint is not binding. This occurs when the value of A_1 is large. If the value of A_1 is small the individual is liquidity constrained causing an increase in A_2 to have no effect on consumption since there is no way the individual can access it. In our example, if a person who is not liquidity constrained (a lot of money in his bank account) finds that his house has doubled in price then his consumption will increase. However, if he is liquidity constrained then the increase in the value of his house will not change his consumption expenditures because he cannot borrow against it.

Formally then his consumption expenditures will follow this pattern

$$\left.\frac{\partial E_{\mathrm{L}}}{\partial A_2}\right|_{A_1 \gg 0} > 0 \quad \text{and} \quad \left.\frac{\partial E_{\mathrm{L}}}{\partial A_2}\right|_{A_1 \approx 0} = 0 \tag{6}$$

Two implications of the behavioral life-cycle model predict that the effects of liquidity constraints will be completely different from those described in Eq. (6). First, as discussed earlier, the liquidity constraints of a behavioral life-cycle individual are internally imposed in addition to being externally forced upon him by imperfections in the capital market. In this case a less liquid or tempting asset (A_2) will only affect his consumption when other more tempting assets (A_1) are depleted and he needs to use it. For example, an increase in the value of a house only affects consumption when another less tempting asset is exhausted. If there is plenty of liquid wealth available, a change in the value of the house will not alter the individual's consumption decision because the value of the house does not enter into his decision making process. However, if the person is strapped for cash then he will want to dip into those less liquid assets. Since the constrained individual is (at least implicitly) using those assets, the components of his decision making process will now include them. In this case, a change in the value of these assets will affect his consumption pattern. Second, if the behavioral life-cycle model is true then individuals who adopt the rules of behavior suggested by it will be more successful at saving than those who do not. In this case the individuals who are least tempted to use illiquid assets to support consumption will be the individuals with relatively larger savings. Thus even after controlling for an individual's initial resources (through a fixed effect or other econometric controls-see below) we will find a differential response to an exogenous change in a particular asset's value between the group that looks liquidity constrained and the group that does not. Formally, both these reasons imply that a behavioral life-cycle individual's consumption will follow this pattern

$$\left.\frac{\partial E_{\mathrm{L}}}{\partial A_2}\right|_{A_1 \gg 0} = 0 \quad \text{and} \quad \left.\frac{\partial E_{\mathrm{L}}}{\partial A_2}\right|_{A_1 \approx 0} > 0 \tag{7}$$

Thus the conventional life-cycle theory and the behavioral life-cycle theory have strongly contrasting and testable implications about the effects of liquidity constraints on consumption behavior.

2.5. Test 4 (Testing for psychological versus financial transaction costs)

Testing for framing in consumption is complicated. Proponents of the conventional life-cycle hypothesis maintain that differences in various assets' marginal propensities to consume may result from those assets having differential risk characteristics, or tax rates[5] or other features that cause a one dollar increase in the value of a house to have the same effect on spending as, say a 50 cent increase in some other asset. The most important objection to Test 3 (see Skinner, 1993) is that the existence of fixed monetary costs of transforming illiquid assets into consumption may also cause Test 3 to be rejected. Thus Test 3 is actually a test of whether there are costs to transforming assets but does not test directly whether those costs are financial or psychological in nature.

However there is a testable difference between the psychological and the financial theory of liquidity constraints. Note that if the financial explanation of transaction costs is correct then there still is a single aggregate of resources (W) upon which an individual bases all his consumption decisions on. In this case how an asset maps into W depends on its characteristics alone, not the characteristics of the good being consumed. Let $\Phi_k(A_k)$ be the function that transforms the value of asset A_k into W. W is now a sum of all the (transformed) assets and a conventional life-cycle individual's expenditures on a particular good (g) will be

$$E_{\mathrm{L}g} = E_{\mathrm{L}g}(W) = E_{\mathrm{L}g}\left(\sum_{k=1}^{K} \Phi_k(A_k)\right) \tag{8}$$

Now suppose we have two different goods (g and h) and two different assets (A_1 and A_2). In this case the individual's wealth will be: $W = \Phi_1(A_1) + \Phi_2(A_2)$ and his expenditures on two the goods will be

$$E_{\mathrm{L}g}(W) = E_{\mathrm{L}g}(\Phi_1(A_1) + \Phi_2(A_2)) \quad \text{and} \quad E_{\mathrm{L}h}(W) = E_{\mathrm{L}h}(\Phi_1(A_1) + \Phi_2(A_2))$$

If we take the derivatives of the expenditure functions of one of the goods (g) with respect to one of the assets (k) we have

$$\frac{\partial E_{\mathrm{L}g}}{\partial A_k} = \frac{\partial E_{\mathrm{L}g}}{\partial W} \cdot \frac{\partial W}{\partial A_k} = \frac{\partial E_{\mathrm{L}g}}{\partial W} \cdot \Phi_k'(A_k) \tag{9}$$

The $\partial E_{\mathrm{L}g}/\partial W$ is the marginal propensity of a conventional life-cycle individual to consume a good for a given change in wealth while $\partial W/\partial A_k = \Phi_k'(A_k)$ is how the change in a particular asset is transformed into a change in wealth. The marginal propensity to consume a particular good out of wealth is obviously not going to be equal for different goods, since its value will depend on the good's wealth elasticity. However, the second term does not depend on what particular good is chosen since its value just depends on how the asset is transformed into wealth and should be the same for all goods. If a conventional life-cycle individual consumes two goods (g and h), then the ratios of the

[5] For example, consider a individual with 2 assets, A_1 and A_2. If A_2 is taxed at a rate of the individual's total resources will be $W = A_1 + (1-\tau)A_2$ and the propensity to consume out of those two assets on a good (G) will be: $\partial G/\partial A_1 = \partial G/\partial W$ and $\partial G/\partial A_2 = (1-\tau)\partial G/\partial W$.

marginal propensities to consume out of the two assets will be equal since

$$\frac{\partial E_{\mathrm{L}g}/\partial A_1}{\partial E_{\mathrm{L}g}/\partial A_2} = \frac{\partial E_{\mathrm{L}g}/\partial W \cdot \Phi_1'(A_1)}{\partial E_{\mathrm{L}g}/\partial W \cdot \Phi_2'(A_2)} = \frac{\Phi_1'(A_1)}{\Phi_2'(A_2)} = \frac{\partial E_{\mathrm{L}h}/\partial A_1}{\partial E_{\mathrm{L}h}/\partial A_2} \qquad (10)$$

For a behavioral life-cycle individual, however, the relationship between an expenditure on a certain good and the value of an asset will depend on both the good's and asset's characteristics. In this case the propensity to consume a particular good out of a particular asset should depend on the attributes of both the good and the asset. As an example, compare spending on groceries to spending on vacations. The conventional model, discussed above, would predict that after controlling for the different wealth elasticities of the two goods, a change in the value of a home or a business will affect spending on groceries by as much as it affects spending on vacations. On the other hand, the behavioral model predicts that spending on a good (like groceries) that is purchased in small quantities every day will be less influenced by the value of a particular asset then a good that is purchased infrequently at a higher cost. In other words, the behavioral life-cycle model predicts that when an individual finds his house has increased in value he will be more likely to change his vacation plans than his grocery purchases. For a behavioral life-cycle individual, then, Eq. (10) will not hold; the ratios of the MPCs between assets and spending on different goods given in Eq. (10) will no longer be equal. Thus Eq. (10) suggests a very general test between the two competing hypotheses. If people engage in framing then asset composition will affect both the composition of spending as well as its size.

3. Data and statistical methods used in the study

3.1. Data

The data used in this study is from the Longitudinal Retirement History Survey (RHS). This survey interviewed around 11 000 household heads who were between 57 and 62 in 1969. This survey was repeated every 2 years until 1979, creating six waves. The survey contains detailed information on family status, wealth, social security, income and expenditures.

The RHS has several features that makes it useful in studying the effect of wealth and income on consumption. First, creating a value for future wealth is much simpler for a group near retirement than for a group of younger cohorts. Determining the value of future wealth of a relatively young person is extremely difficult since it will depend on his future wage income as well as his future social security and pension, all of which are extremely difficult to predict. On the other hand, since the individuals in the RHS are either retired or near retirement, their social security wealth and pension wealth are known or can be predicted with a much larger degree of certainty. Further, future wage income is a much lower portion of future wealth for the elderly, making exact modelling relatively less important. The second advantage of using a sample of the elderly is that liquidity constraints are relatively unimportant for them. The inability to borrow against future labor income may cause younger individuals to be seriously liquidity constrained.

However, the peak earning years of the elderly are behind them. Thus their spending patterns should not be as constrained by their inability to borrow against future income. A third advantage of the RHS is that it is a relatively long panel (10 years), during which most of the households progress from one stage of the life-cycle (working) to another (retirement). Thus we may be able to examine how life-cycle changes affect the relationship between consumption spending and asset type.

To study consumption, I use eight of the ten expenditure categories[6] in the RHS as dependent variables. All expenditure categories are annual expenditures for an entire household. The eight consumption variables are listed below. Unfortunately, not all the consumption questions were asked for all of the waves of the RHS; the dependent variables and their range are listed below:

Dependent Variable	In 1969	In 1971	In 1973	In 1975	In 1977	In 1977
Food Eaten at Home (Groceries)	Yes	Yes	Yes	Yes	Yes	Yes
Food Eaten Out (Food Out)	No*	Yes	Yes	No*	Yes	Yes
Charitable Contributions	Yes	No	No	Yes	Yes	Yes
Dues to Clubs	Yes	Yes	No	Yes	Yes	Yes
Entertainment	Yes	Yes	No	Yes	Yes	Yes
Gifts	Yes	Yes	No	Yes	Yes	Yes
Transportation	No	Yes	Yes	Yes	Yes	Yes
Vacations	Yes	Yes	Yes	Yes	Yes	Yes

* For Food Out, 1969 and 1975 were excluded because the values for those 2 years were not consistent with the values for the other years.

Unfortunately this data cannot be combined into one total consumption category because of missing values. Since there are relatively few observations which have non-zero values for all eight categories, a different equation for each consumption category will be estimated.[7]

Two sets of independent variables are used in this study. The first set is household characteristics. Since the equations are estimated using fixed effects (see ahead), only household characteristics that change over time need to be controlled for. The included characteristic variables are married dummy (Married), retirement dummy (Retire), spouse employment dummy (Spouse Work), changes in health status (Feel Better-Health better than 2 years ago, Feel Worse-Health worse than 2 years ago) and number of people in the household (#_in_House). The Retirement dummy was constructed in two steps. First, those who said they were retired and had wage incomes of less than $ 1000 a year were

[6] The two missing categories are amounts spent on chores and amounts spent on non-vacation trips. Very few households have more than one non-zero observation for either of these categories. Therefore these categories could not be analyzed in the manner discussed in Section 3.3.

[7] I have also excluded single women who were never married during the panel (the RHS does not contain men who were single in 1969). Preliminary investigation has shown that their reported consumption behavior is much different from the behavior of couples and is not explained by changes in their income or any other variable. This may be due to less accurate data (we do not have social security records for their husbands) and other wealth data may be incorrect. Although including these women does not substantially change the results, it does weaken them.

considered retired. Second, for individuals who did not respond to the retirement question or claimed they were not retired, the retirement dummy was constructed by finding the first year that their wage earnings were less than $ 1000 and then stayed less than that in subsequent waves. Using this definition, most individuals were retired by 1977 and all were retired by 1979.

The second set of independent variables consists of the different assets. In this study, I use five types of assets: current income, liquid assets, value of house, future assets and non-liquid assets. Current income (Income) consists of wages, net interest income, dividends, social security, pension income, and business or farm income less payments on debts. Liquid assets (Liquid Wealth) consist of assets that are relatively easy to transform into cash. These include checking and savings accounts, stocks and bonds, and (negatively) the value of any non-secured debt. House value (House) is the net equity in the home (market value of home less the mortgage). Future assets (Future Wealth) include the (actuarially discounted) present value of social security and pension income plus the discounted value of future labor income.[8] Non-liquid assets (Property) consist of everything that remains, including net value of farm, business, and other (non-house) property. All asset values are net of current payments to the household (which are counted as income).

In many studies of consumption, rejection of life-cycle models is attributed to liquidity constraints caused by an inability to borrow against future labor income (for e.g., see Zeldes, 1989 or Deaton, 1989). Although most of the individuals in this study are near retirement, some might still wish to borrow against future social security or pension income and find themselves unable to. To examine whether this is important I have split the sample into constrained and unconstrained subsamples in a manner analogous to Zeldes. I have defined the unconstrained group as those households whose liquid wealth exceeds at least 3 months of household income for at least five of the six periods in the sample, while the constrained group consists of everyone else.

3.2. Means

Table 1 contains the means of the consumption and asset variables used in the different regressions (which are labelled at the top).[9] Since each regression is taken from a different sample, the means of the variables in each of the regression samples are reported. The first row (labelled consumption) is the mean of the dependent variable used in that regression. The means are for the entire sample and are also broken down for the constrained and unconstrained subsamples.

One striking result from the table of means is that the households who are in the Entertainment, Vacation, and Dues regressions have a higher mean for income and wealth then the households in the other spending categories. Mean Income for these three

[8] I assumed that individuals have perfect foresight about their future career path (i.e. know both their future income and when they will retire) and so use their actual earnings to calculate their expected future income. It turns out that including or not including future labor income does not change the results very much, so a more sophisticated treatment of future income is probably not worthwhile (almost two-thirds of the sample were retired by the third wave).

[9] Means of the household characteristic variables are available upon request.

Table 1
Means and number of observations for the wealth and consumption variables

	Group	Groceries	Food out	Charity	Dues	Entertainment	Gifts	Transport	Vacation	All[a]
Dependent variable[*]	Overall	$ 1146	$ 691	$ 171	$ 62	$ 111	$ 165	$ 307	$ 635	$ 3287
	Unconstrained	$ 1117	$ 723	$ 194	$ 68	$ 115	$ 183	$ 314	$ 677	$ 3393
	Constrained	$ 1191	$ 639	$ 125	$ 44	$ 92	$ 128	$ 293	$ 418	$ 2930
# of observations	Overall	2505	2211	2432	835	520	2356	909	594	
	Unconstrained	1548	1360	1598	608	428	1570	597	498	
	Constrained	957	851	834	227	92	786	312	96	
Income	Overall	$ 8602	$ 5514	$ 5826	$ 9393	$ 11 127	$ 7463	$ 7665	$ 12 493	
	Unconstrained	$ 9752	$ 6294	$ 6497	$ 10 075	$ 11 453	$ 8190	$ 8553	$ 12 934	
	Constrained	$ 6740	$ 4268	$ 4539	$ 7565	$ 9609	$ 6012	$ 5964	$ 10 203	
Liquid wealth	Overall	$ 13 483	$ 9104	$ 10 953	$ 19 891	$ 24 603	$ 12 420	$ 13 772	$ 26 201	
	Unconstrained	$ 20 681	$ 14 028	$ 15 680	$ 26 309	$ 28 899	$ 17 553	$ 19 923	$ 30 332	
	Constrained	$ 1838	$ 1234	$ 1896	$ 2702	$ 4618	$ 2167	$ 2001	$ 4771	
Future wealth	Overall	$ 39 923	$ 27 216	$ 25 565	$ 38 301	$ 42 151	$ 33 527	$ 37 297	$ 50 459	
	Unconstrained	$ 43 103	$ 29 627	$ 27307	$ 39600	$ 42 933	$ 35 502	$ 40 213	$ 51 045	
	Constrained	$ 34 779	$ 23 362	$ 22 227	$ 34 824	$ 38 514	$ 29 581	$ 31 718	$ 47 416	
Property	Overall	$ 14 189	$ 8952	$ 9516	$ 17 565	$ 19 159	$ 11 726	$ 11 765	$ 21 855	
	Unconstrained	$ 17 838	$ 11 271	$ 11 280	$ 20 146	$ 20 567	$ 13 739	$ 14 426	$ 23 298	
	Constrained	$ 8286	$ 5246	$ 6137	$ 10 650	$ 12 605	$ 7706	$ 6674	$ 14 370	
Home value	Overall	$ 20 254	$ 13 868	$ 14 694	$ 20 654	$ 22 159	$ 17 689	$ 17 796	$ 24 783	
	Unconstrained	$ 23 180	$ 15 712	$ 16 165	$ 21 883	$ 23 378	$ 19 445	$ 20 084	$ 25 748	
	Constrained	$ 15 520	$ 10 923	$ 11 877	$ 17 362	$ 16 487	$ 14 182	$ 13 417	$ 19 777	

The top row in each box is the mean for the entire sample.
The row titled ‘Constrained’ contains means for the subsample whose liquid wealth is less than 6 months of income.
The row titled ‘Unconstrained’ contains means for the subsample whose liquid wealth is greater than 6 months of income.
[*] Means of Dependent and Independent Variables used in regression for that consumption category.
[a] All is the sum of the means of the consumption categories.

samples is over $ 10 000 while for the other five categories mean income is near $ 8000. A similar result holds for the various wealth variables. The reason for this is that there are fewer observations for these categories than there are for the other ones, and being in this sample is positively correlated with wealth. Although it is obvious why taking vacations correlates strongly with wealth it is somewhat odd that joining a social club (which has an average expenditure of only $ 69) or spending on entertainment should be. This heterogeneity in the samples should be remembered when comparing the results from the eight regressions. Further, even though the combined results will be reported they can only be interpreted as indicative, not definitive.

Another less surprising result is that the constrained group is on average poorer than the unconstrained group. For example, the constrained group's mean income is around 25 percent lower than the unconstrained group's mean income. Even more striking, the constrained group's mean liquid wealth is only 10–15 percent of the unconstrained group's mean liquid wealth. The other wealth variables are also considerably lower for the constrained group then for the unconstrained group.

The column marked ALL contains the sum of the means of the eight consumption variables. The value of $ 3287 is the sum of the eight categories mean spending levels and is between 25–40 percent of income. This provides good evidence that the eight consumption variables are a large proportion of total consumption spending.

3.3. Empirical methods

To consistently estimate different assets' MPCs, five major econometric problems need to be solved. These problems are: Unobserved differences between individuals, autocorrelation of errors between years, measurement error in the wealth data, possible simultaneity bias between spending and savings decisions, and sample selection bias caused by missing observations. Any consistent estimator of the effect of asset composition on consumption must take these problems into account.

Unobserved differences between individuals may result in people with the same level of income having different amounts of savings. These differences may be caused by different tastes for savings, different income paths (individuals whose income is more variable may have more precautionary savings), or differences in luck. The coefficients from a standard regression of assets on consumption may be contaminated by these unobserved differences. This bias occurs because the estimated coefficient of a particular asset will include both the marginal propensity to consume from that asset and the cross-correlation between the unobservable differences, the value of the asset and consumption. To solve this problem, I will use fixed effects. Fixed effects sweep out the unobserved differences between people by subtracting out the individual means from all the variables before estimation.[10] This method is equivalent to using individual dummy terms in all the equations. This technique should eliminate the contamination caused by the unobserved

[10] This method is somewhat unorthodox. Griliches and Hausman (1986) suggest sweeping out the means by first differencing and using a slightly different set of instruments. Since some of the dependent variables are missing in the middle of the sample (usually 1973) first differencing would lead to some excess data loss. Preliminary investigations using first differences on grocery and vacation spending (which span all 6 years) suggest that the results of the Griliches-Hausmann technique and those reported here are very similar.

differences and allow for unbiased estimates of the MPCs. Even though this technique allows the estimates to be conditioned on the individual's initial resources, it does not control for differential responses of these individuals to exogenous changes in the value of those assets. In this case the estimated parameter can be interpreted as the average MPC of a particular asset for a given sample. If the sample changes then the MPC may change with it (see discussion in Section 2 on why samples might have different MPCs).

The problem of auto-correlation of the equations' error term can be dealt with straightforwardly. Instead of assuming that the error term follows an AR(1) or AR(2) process that is constant between periods, I allow the autocorrelations between the different years to be free parameters that are estimated directly. The method[11] for doing this is to treat each year as a separate equation. Then the equations for the different years are estimated jointly by a method similar to seemingly unrelated regressions or three-stage least squares in which the covariances between equations are actually the auto-correlations of the error terms and are estimated directly.

A typical problem with using wealth data is that this data is often riddled with measurement error. The RHS is no exception (see Hurd, 1987). To control for measurement error, I use a technique devised by Griliches and Hausman (1986), which utilizes the panel structure of the RHS. Their technique exploits the multiple observations per individual that panel datasets contain. To instrument a variable contaminated by measurement error one uses the value of that variable from different periods. For example, to instrument liquid wealth in the 1975 equation one can use the value of liquid wealth in 1973 and 1977 as instruments, while to instrument liquid wealth in the 1977 one would use the 1975 and 1979 values. This method is consistent as long as the measurement error does not follow a moving average (MA) or auto-regressive process but there is autocorrelation in the true value (since otherwise the variables would not be good instruments). The assumption of no MA component in the measurement error can be tested for by calculating a Hausman-Specification test comparing the estimates when the adjacent time period values are used as instruments with the estimates when the instruments are drawn from values from periods further away (which would not be affected by a MA error).[12]

In a similar manner, simultaneity bias can also be tested for and corrected. A potential problem for estimating the effect of wealth on consumption is that decisions on how much to consume and how much to save may be made simultaneously; if you decide to save more you will consume less. In this case, using the future values of the wealth variables as instruments will not lead to consistent estimators because a consumption/savings decision in an earlier period will affect the future values of those variables too. This effect can also be tested for by calculating a Hausman test statistic comparing the estimates that use both the future values and past values of the asset values as instruments (and so should be more efficient if measurement error is the major problem) with the estimates that only use assets' past values, which are consistent even when simultaneity

[11] This method is outlined in Hsiao (1987).

[12] Testing for autocorrelation in the measurement error is much harder (if not impossible). However, if autocorrelation of the measurement error was a serious problem then the estimates would still change when instruments of different lengths were used, which might be discerned by a Hausman test.

bias occurs because their values are not affected be current decisions (i.e. their values are predetermined).[13] The test statistics for both measurement error and simultaneity bias will be reported below.

The final problem with the data is sample selection. Many of the observations do not have complete (or even any) information on the consumption categories for all the periods, or claim that the value of that good is zero (which seems implausible for some of the categories, particularly groceries). Although one could handle this problem by making some strong assumption about the error distribution (such as normality) I will instead only use households that have valid non-zero consumption values for all the periods. I am thus assuming that for the households selected by this criteria, the process that puts them into the sample is constant over the time-span and so is a fixed effect. By controlling for these fixed effects, I will control for any sample selection effect.

The method outlined above differs from normal three-stage least squares in two ways. First, the coefficients of the variables are restricted to be the same across all the equations (which are the years). Second, to use three-stage least squares one must assume that all the instruments can be used in all of the equations. Since the instruments are the values of the non-adjacent variables, this method will not work. Instead one must use a generalized method of moments estimator that allows the instruments to vary between equations.

4. Empirical results

4.1. General regression results

The baseline regression results are in Table 2. Table 3 contains results on the differences between the estimated MPCs. The estimates of the MPCs for both the liquidity constrained and non-liquidity constrained subsamples are in Table 4. For all four tables, columns are labelled with the consumption category that is the dependent variable. The column labelled All is the aggregate consumption category and is the sum of the previous eight columns. In this section, I discuss some general results from the regressions. In the sections following, specific tests between the conventional and behavioral life-cycle hypotheses developed earlier are conducted.

The first parts of Tables 2 and 4 contain the R^2 and the number of observations for the different regressions. Table 2 also contains the chi-square statistics on tests of various hypotheses. The second part of Tables 2 and 4 are estimates of the MPCs of the different assets.

The first part of each table reveals very low values of R^2. One reason for this is that by using a fixed effect one takes out the between variation (in a cross-section regression the R^2 is between 0.3 and 0.4) but retains a lot of the variation caused by measurement error in the dependent variable. Luckily, even though the R^2 is small, some effects can still be discerned.

[13] At a minimum, acceptance of the hypothesis of no simultaneity bias implies that even if there is actually some bias, the effect of this bias is small relative to the variance of the estimates.

Table 2
Expenditure function results – regression statistics and wealth coefficients for the combined sample

	Groceries	Food out	Charity	Dues	Entertainment	Gifts	Transport	Vacation	All†
R^2	0.11	0.09	0.06	0.02	0.01	0.03	0.02	0.01	
# of obervations	2505	2211	2432	835	520	2356	909	594	
Simultaneity test χ^2	13.96	26.48**	8.82	18.96	10.06	15.21	5.87	22.86	
Measurement error test χ^2	22.65	28.01**	10.17	11.51	10.63	22.26	16.22	20.73	
Estimates of marginal propensities to consume (MPC) of different assets									
Income	11.47**	1.78	8.82**	1.08**	0.94*	6.68**	10.71**	0.24	41.73
(t-statistic)	(5.50)	(0.71)	(7.25)	(2.62)	(1.74)	(7.94)	(5.26)	(0.13)	
Liquid wealth	−0.04	0.33*	−0.06	0.08*	0.07**	0.46**	0.33**	0.44**	1.61
(t-statistic)	(0.28)	(1.90)	(0.61)	(2.90)	(2.40)	(5.49)	(3.09)	(2.87)	
Future wealth	1.68**	0.12	0.27	−0.11	0.10	0.32**	−0.59	−0.57	1.23
(t-statistic)	(3.14)	(0.18)	(1.14)	(1.35)	(0.59)	(2.01)	(1.63)	(1.18)	
Property	0.12	−0.17	−0.06	−0.05**	0.11**	0.03	0.45**	0.13	0.55
(t-statistic)	(1.00)	(0.73)	(0.70)	(2.33)	(3.07)	(0.40)	(4.03)	(1.05)	
Home value	0.47	−0.05	−0.18	−0.06	−0.04	−0.14	0.24	−0.24	0.01
(t-statistic)	(1.25)	(0.09)	(1.14)	(1.46)	(0.50)	(0.91)	(0.85)	(0.84)	

** Significant at 5 percent level
* Significant at 10 percent level.*
† All is the sum of the income or wealth coefficients for all of the consumption categories.
The χ^2 tests for potential simultaneity or measurement error bias. See text for details.

Table 3
Tests for the equality of different wealth coefficients MPCs

	Groceries	Food Out	Charity	Dues	Entertainment	Gifts	Transport	Vacation	All[a]
Income–Liquid wealth	11.51**	1.46	8.88**	1.00**	0.88	6.21**	10.39**	–0.20	40.12
(*t*-statistic)	(5.45)	(0.57)	(7.27)	(2.40)	(1.62)	(7.28)	(5.04)	(0.11)	
Income–Future wealth	9.79**	1.66	8.55**	1.19**	0.85	6.35**	11.31**	0.81	40.50
(*t*-statistic)	(4.33)	(0.62)	(6.72)	(2.73)	(1.46)	(7.17)	(5.31)	(0.41)	
Income–Property	11.35**	1.96	8.88**	1.13**	0.83	6.65**	10.26**	0.12	41.18
(*t*-statistic)	(5.40)	(0.76)	(7.17)	(2.67)	(1.52)	(7.83)	(4.99)	(0.06)	
Income–Home value	11.00**	1.83	8.99**	1.14**	0.98*	6.81**	10.48**	0.48	41.72
(*t*-statistic)	(5.13)	(0.69)	(7.25)	(2.69)	(1.73)	(7.67)	(4.94)	(0.27)	
Income χ^2	46.31**	3.13	57.06**	20.43**	6.11	74.8**	30.31**	7.86*	–
Liquid wealth–Future wealth	–1.72**	0.21	(0.33)	0.19**	(0.03)	0.14	0.92**	1.00**	0.38
(*t*-statistic)	(3.06)	(0.31)	(1.29)	(2.46)	(0.19)	(0.78)	(2.53)	(2.00)	
Liquid wealth–Property	(0.16)	0.5*	0.00	0.13**	(0.04)	0.44**	(0.12)	0.31	1.06
(*t*-statistic)	(0.93)	(1.66)	(0.01)	(3.60)	(1.02)	(3.57)	(0.83)	(0.83)	
Liquid wealth–Home value	(0.51)	0.37	0.11	0.14**	0.11	0.60**	0.09	0.68**	1.60
(*t*-statistic)	(1.29)	(0.65)	(0.68)	(2.85)	(1.21)	(3.23)	(0.33)	(2.09)	
Liquid wealth χ^2	10.79**	2.94	2.43	17.77**	4.20	16.16**	8.09**	7.85**	–
Future wealth–Property	1.56**	0.29	0.34	–0.06	–0.01	0.30*	–1.05**	–0.69	0.68
(*t*-statistic)	(2.79)	(0.42)	(1.29)	(0.72)	(0.08)	(1.66)	(2.83)	(1.37)	
Future wealth–Home value	1.21*	0.17	0.45	–0.05	0.14	0.46**	–0.83*	–0.33	1.22
(*t*-statistic)	(1.79)	(0.20)	(1.55)	(0.56)	(0.71)	(2.09)	(1.83)	(0.58)	
Property–Home value	–0.35	–0.13	0.11	0.01	0.15*	0.17	0.22	0.36	0.54
(*t*-statistic)	(0.91)	(0.21)	(0.61)	(0.18)	(1.85)	(1.00)	(0.69)	(1.13)	
Other χ^2	8.47**	0.20	2.42	0.52	3.52	4.47	8.08**	2.82	–

** Significant at 5 percent level, * Significant at 10 percent level.
[a] All is the sum of the income or wealth coefficients for all of the consumption categories.
The χ^2 tests the joint hypothesis that the differences between the named coefficient and all the other coefficients are zero.

Table 4
Expenditure function results – regression statistics and wealth coefficients for the constrained and unconstrained subsamples

	Groceries	Food Out	Charity	Dues	Entertainment	Gifts	Transport	Vacation	All[a]
All coefficients χ^2	43.96**	71.66**	22.62	19 624**	733.32**	21.70	146.91**	1237.4**	
Wealth coefficients χ^2	19.88**	17.56**	4.51	778.7**	75.99**	5.57	27.31**	303.3**	
Results for the constrained subsample									
R^2	0.12	0.11	0.33	0.03	0.03	0.02	0.02	0.02	
# of obervations	957	851	834	227	92	786	312	96	
Estimated marginal propensities to consume									
Income	13.69**	−4.46	3.94	0.78	−3.03**	5.48**	7.28*	1.56	25.25
Liquid wealth	−1.86	4.92	−1.75	0.12	0.77	0.43	0.40	−2.01	1.01
Future wealth	3.17**	0.23	−0.10	−0.03	0.38	0.14	−0.39	−1.24	2.16
Property	0.73	0.61	0.13	−0.11	0.25	0.33	1.01**	0.31	3.26
Home value	1.08	−0.57	0.36	−0.06	0.63	0.17	−0.55	4.36*	5.42
Results for the unconstrained subsample									
R^2	0.11	0.08	0.06	0.02	0.01	0.03	0.03	0.01	
# of obervations	1,548	1,360	1,598	608	428	1,570	597	498	
Estimated marginal propensities to consume									
Income	9.67**	3.08	10.35**	0.95**	0.80	6.63**	9.08**	0.04	40.61
Liquid wealth	−0.05	0.18	−0.05	0.09**	0.05*	0.48**	0.29**	0.41**	1.41
Future wealth	0.91	0.02	0.45*	−0.13	0.07	0.42**	−0.35	−0.58	0.80
Property	−0.04	−0.15	−0.13	−0.03	0.11**	−0.04	0.41**	0.11	0.25
Home value	0.11	−0.09	−0.34**	−0.04	−0.04	−0.17	0.15	−0.17	−0.59

** Significant at 5 percent level, * Significant at 10 percent level.
[a] All is the sum of the income or wealth coefficients for all of the consumption categories.
χ^2 statistics are tests for pooling the constrained and unconstrained group.
The row labelled 'All Coefficients χ^2' has tests of the hypothesis that all the coefficients are the same for the constrained group and unconstrained group.
The row labelled 'Wealth Coefficients χ^2' has tests of the hypothesis that the coefficients of the asset and income variables are equal for both groups.

In Table 2, the line titled 'Measurement Error MA χ^2' contains the chi-square statistics for a Hausman test on whether measurement error follows a first order moving average process (the null hypothesis is that measurement error is white noise). This test consists of comparing the estimates when the instruments are one period away (for example, using wealth from 1973 and 1977 to instrument 1975 wealth) which are valid and most efficient if the measurement error is white noise to those when the instruments are two periods away (in same example, using 1971 and 1979 wealth as instruments) which are valid if the measurement error is either white noise or follows a first order moving average process.[14] The line titled 'Simultaneity χ^2' has the same interpretation except that the null hypothesis is no simultaneity bias. In this case estimates using current and past values of the wealth variables are compared to estimates that use only the past values of the wealth variables as instruments. For both tests the null hypothesis is rejected at the 5 percent level only for the Food Out category, and even those results do not change much when the more robust estimator is used (i.e. the coefficients do not change much but the *t*-statistics become lower). In general then, it seems safe to assume that the measurement error in the data does not have a significant moving average component and that the effect of simultaneity bias, if it exists, is small. The χ^2 tests in Table 4 will be explained below.

4.2. Tests of the behavioral life-cycle theory

The estimates for the income and the wealth variables are reported in two parts. The first part (Table 2) consists of the coefficients and the *t*-statistics. Since the income and wealth variables are divided by one-hundred before entering the regression, the coefficients on these variables are the number of cents spending on a particular category increases when the wealth or income variable is increased by one dollar. The second part (Table 3) consists of tests of linear restrictions on whether various wealth and income coefficients are equal. For example, the row labelled 'Income–Liquid Wealth' consists of the differences between Income and Liquid wealth coefficients (with the *t*-statistic for that difference below it). The rows labelled with 'χ^2' are chi-square statistics for the null hypothesis that the sets of coefficients are equal. There are three sets of χ^2 statistics. The first chi-square (labelled 'Income χ^2') tests the null hypothesis that the coefficient on income is equal to the coefficients of the four different wealth variables. The second set (labelled 'Liquid χ^2') tests the hypothesis that the coefficient on liquid wealth is equal to the coefficients of the three remaining wealth variables. The last chi-square (labelled 'Other χ^2') tests whether the three remaining coefficients are equal.

In general the results from the wealth variable are more consistent with behavioral models of consumption than with the simplest life-cycle model. Income has a large and significantly positive effect on spending in six of the eight categories (for entertainment the significance level is only 10 percent) even though the capitalized value of the different income streams is included in the wealth variables (the two exceptions are for the food out category where the coefficient is still large and positive but not statistically significant

[14] Of course one could do a similar test for the presence of any measurement error by comparing the results from using instruments to the results when instruments are not used. In this case the null hypothesis of no measurement error is always strongly rejected.

and vacation spending which is lumpy, perhaps causing individuals to be more willing to dip into assets to pay for it (see ahead)). The effect of receiving an extra dollar of income is to increase spending on all the categories by almost 42 cents. On the other hand, the wealth variables have much less effect on consumption. Of the four wealth variables, Liquid Wealth seems to have the largest effect on consumption. It is positive and statistically significant at the 10 percent level (or better) for six of the consumption categories. However, its effect on consumption (even when significant) is quite small. On the aggregate measure, a dollar increase in wealth increases consumption by less than 2 cents.

4.2.1. Test 1 (Comparing the MPC from income to the MPC from wealth)

To test the behavioral life-cycle theory against the conventional life-cycle theory a direct comparison of the marginal propensities to consume out of different assets is needed. If the conventional life-cycle theory holds then receiving an extra dollar of permanent income should change consumption as much as receiving an extra amount of wealth equal to the present value of that dollar, (i.e. the from Eq. (1)). The ratio of the marginal propensity to consume out of permanent income to the marginal propensity to consume out of wealth will be an estimate of that from Eq. (1). From Table 1 the aggregated value of this ratio is 41.73/1.61=25.9 which means that a $ 1 increase in income has as much effect on consumption as a $ 26 increase in wealth. For this ratio of the marginal propensities to spend out of income and wealth to be consistent with a conventional life-cycle model (i.e. Eq. (2)), a couple of average age in my sample (husband 65, wife 62) would have to (actuarially) discount income by an interest rate of *negative* 2.3 percent,[15] even under the heroic assumption that all the changes in income are permanent (if the changes in income were not permanent then the ratio should be smaller).[16] Thus the ratio of the MPC of income relative to liquid wealth appears to be too large to be consistent with conventional life-cycle theory; it is, however, consistent with the behavioral model.

4.2.2. Test 2 (Testing for the equality of different assets' MPC)

On the other hand, comparing liquid wealth to the other wealth variables demonstrates that individuals are much more willing to spend out of liquid assets than they are out of other forms of wealth. In three of the eight spending categories the difference between the coefficient of liquid wealth and property is positive and statistically significant (when the difference is negative, the *t*-statistics are always small). In the aggregate, the effect of property on aggregate consumption is quite small (around a half a penny per dollar). This is only one-third the size of liquid wealth's effect.

A similar relationship holds between the coefficients on house value and liquid wealth. In fact, the coefficient on housing wealth is negative for six of the eight consumption categories and is *never* statistically different from zero. This implies that changes in

[15] All calculations of implied interest rates adjust for mortality probabilities. Even if the household assumed instead that it was going to survive for 30 years for sure, the implied interest rate would still be very low (about 1 percent).

[16] For example, social security income changes with marital status. Thus changes in social security income could not be considered 'permanent' for both members of the household.

housing wealth have no discernable effect on consumption at all! Thus the results for both housing and property wealth are more consistent with the behavioral life-cycle hypothesis, which asserts that households do not treat different types of assets as being equivalent to each other (i.e. assets do not seem to be fungible, Eq. (5) seems to describe reality better than Eq. (4)).

Interestingly though, the difference between liquid and future wealth is much smaller. While the differences between the two coefficients are positive five out of the eight times, and this (positive) difference is statistically significant three of eight times, the effect of future wealth on aggregate consumption is around 75 percent of the effect of liquid wealth (although the effect is concentrated on grocery spending). Thus, households seem to be more comfortable consuming out of future wealth than out of other forms of non-liquid wealth. One possible reason for this is that future wealth is mainly the value of social security, which increased rapidly in this period. Perhaps an increase in one period increased people's confidence that it would continue to increase in the future.

4.2.3. Test 3 (Testing for liquidity constraints)

A standard explanation of why consumption is sensitive to income and not to wealth is the existence of liquidity constraints. The third test examines whether the effects of liquidity constraints are consistent with a conventional life-cycle model or instead can be better described by a behavioral model. According to conventional life-cycle theory, the marginal propensity to consume should be the same for all assets except for individuals who are liquidity constrained and cannot borrow against an illiquid asset. In that case, the marginal propensity to consume out of the illiquid asset will be zero (Eq. (6) holds). However, for a behavioral life-cycle individual liquidity constraints occur because of the psychological cost of using a less liquid asset. As long as more liquid assets have not been exhausted the propensity to consume out of the less liquid asset will be zero. Only when the individual has to dip into that asset will consumption be sensitive to its value (Eq. (7) holds).

To examine this issue empirically, I divided my sample into liquidity constrained and unconstrained subsamples (defined in Section 3.1). The regression results and MPC estimates for both subsamples are in Table 4. In the box labelled 'All Coefficients χ^2' is a chi-square statistic and associated probability for the hypothesis that the coefficients from the two equations are the same. In the box labelled 'Wealth Coefficients χ^2' is a test that just the Income and Wealth coefficients are the same for both groups. For six of the eight categories (Charity and Gifts are the only exceptions) both hypotheses are rejected. Thus we can be confident that the two groups are behaving differently.

Comparing the results for the unconstrained sample to the constrained sample, the most striking result is that illiquid forms of wealth affect consumption *less* for the unconstrained subsample than they do for the constrained sample (Table 3). For almost all the consumption categories, the propensity to consume out of any form of illiquid wealth is smaller for the unconstrained group than it is for the constrained group. The result of this can be seen most clearly by comparing the aggregate consumption of the two subsamples. For every extra dollar of future wealth, aggregate consumption for the constrained subsample increases by 2.16 cents for a $ 1.00 increase in this asset's value, while for the unconstrained group this increase is only 0.80 cents. The contrasts are even larger for the other two assets. For the unconstrained group, the value of the propensity to consume out of property is one-thirteenth of the value estimated for the constrained group (0.25 cents per dollar for the unconstrained group versus 3.26 cents of extra spending per dollar for the constrained group). The most extreme difference is caused by changes in housing wealth. The unconstrained group's propensity to consume out of housing is generally negative (six of eight times) and has an aggregate value of −0.59 cents! In contrast, the constrained group's marginal propensity to consume out of housing wealth is quite large (5.42 cents out of every dollar of appreciation). Thus the unconstrained group's propensity to consume out of illiquid forms of wealth is significantly lower than the constrained sample – just the reverse of what the external liquidity constraint hypothesis would have predicted but exactly in line with the prediction of the behavioral life-cycle. Individuals do not appear to finance consumption out of less tempting or illiquid assets unless their more liquid assets are exhausted. Thus liquidity constraints seem to be internally not externally imposed.

The observed difference between the constrained and unconstrained groups' marginal propensity to consume out of illiquid forms of wealth cannot be attributed to higher spending patterns of the constrained group. Note that the unconstrained group's marginal propensity to consume out of income is more than 50 percent larger (at over 40 cents on the dollar) than the constrained group's marginal propensity to consume out of income (about 25 cents). This is a very counter-intuitive result since it implies that the consumption of the unconstrained tracks income more than the consumption of the constrained group. This contradictory behavior is hard to explain with the conventional model but may be consistent with the behavioral model. For the poorer constrained group an extra dollar of income might have a higher chance of being spent on a good that is not included in this study (for example, rent, non-food groceries, clothing, utilities, etc).[17] However, the constrained group has already broken the mental barrier against spending out of the illiquid assets. Therefore, their spending will be sensitive to changes in those assets, whereas the unconstrained group's spending will not be very sensitive. Further, the propensity to consume out of wealth may be higher for the constrained group because these are the individuals who have been less successful at saving and so are more tempted by their less liquid assets; in contrast the unconstrained subsample contains relatively successful savers who may have adopted the behavioral rule of not consuming out of less tempting forms of wealth.

A behavioral bequest motive may also help to explain the failure of the unconstrained group to spend out of their illiquid assets. Individuals with a bequest motive may have tried to lock in their bequest by investing it in an illiquid asset (like their house) that they are less tempted to use for their own consumption. Thus owning a large illiquid asset may act as a disciplinary device to insure the desired bequest occurs. Unlike a standard bequest motive, however, this strategy implies that the size of the bequest is not neutral with respect to which asset changes in value. A large increase in home value will be entirely 'saved' for a bequest while an equivalent rise in a more liquid asset will not.

[17] As an example, note that the only good whose marginal propensity to consume out of income is higher for the constrained group is Groceries, which is almost certainly not a luxury good.

Even though the results for both the constrained and unconstrained groups are not consistent with models of conventional liquidity constraints they are consistent with predictions made by the behavioral life-cycle model. Thus liquidity constraints seem not to arise from an external market failure as much as an internal desire not to spend out of certain types of assets unless one has to.

4.2.4. Test 4 (Testing for psychological versus financial transaction costs)

As stated earlier, the results from the liquidity test are consistent with the existence of financial transaction costs as well as psychological transactions costs. The testable difference between these two theories is that if the costs of spending of an asset were only financial then variations in asset holdings would affect the overall level spending but not the composition of that spending. In this case the amount a conventional life-cycle individual spends on a particular good should only depend on his total resources and the wealth elasticity of that good. Therefore, transforming one asset into another should not affect the split of his spending between different goods. On the other hand, Shefrin and Thaler (1981) posit that behavioral life-cycle individuals create separate mental accounts for the different goods they buy analogous to separate mental accounts they have for different assets. In this case, individuals may be more willing to spend a particular asset on one good than on another. For example, an increase in housing equity coupled with an equivalent decrease in another asset may cause vacation spending to rise while grocery spending falls. Observation of this type of behavior then provides direct evidence for the behavioral life-cycle theory.

One potential problem with this strategy is that with the exception of the constrained group, consumption is not very sensitive to changes in any type of wealth. This may make it difficult to find much of a relationship between the values of particular assets and the amount of spending on a good even if the behavioral life-cycle hypothesis is true.

To test for the effect mentioned above one has to compare the ratios of the marginal propensities to consume out of particular assets for different goods (doing this controls for goods having different wealth elasticities – see the discussion of Test 4 in Section 2 for the derivation). Under the conventional life-cycle hypothesis the ratios should be equal, whereas under the behavioral life-cycle hypothesis they may not be.

To conduct this direct test, observations that had values for more than one consumption category were needed. In order to maximize the number of observations, only the consumption values for 1975, 1977, and 1979 were used (for Food Out which did not have valid data for 1975, the 1973 values were used) and 28 separate regressions were run, in which each consumption category was paired with another one. This task was repeated three times, once for the overall sample and then for the constrained and unconstrained subsamples. Using the results from these paired equations, ratios of the wealth coefficients divided by the income coefficient were constructed for each of the categories. Then using a Wald test, the equality of these ratios between the two equations was tested for. For example, consider consumption goods G and H with estimated coefficients on the wealth and income variables G_I, H_I (for income), G_L, H_L (for liquid wealth), G_F, H_F (for future wealth), G_P, H_P (for property wealth), and G_H, H_H (for home wealth), the test statistic is a χ^2 for the joint hypothesis that

$$\frac{G_L}{G_I}=\frac{H_L}{H_I} \text{ and } \frac{G_P}{G_I}=\frac{H_P}{H_I} \text{ and } \frac{G_F}{G_I}=\frac{H_F}{H_I} \text{ and } \frac{G_H}{G_I}=\frac{H_H}{H_I} \text{ and } \frac{G_P}{G_L}=\frac{H_P}{H_L} \text{ and } \frac{G_F}{G_L} =\frac{H_F}{H_L} \text{ and } \frac{G_H}{G_L}=\frac{H_H}{H_L} \text{ and } \frac{G_F}{G_P}=\frac{H_F}{H_P} \text{ and } \frac{G_H}{G_P}=\frac{H_H}{H_P} \text{ and } \frac{G_H}{G_F}=\frac{H_H}{H_F}$$

The results for the entire sample were completely negative. It turns out that NONE of the 28 χ^2 statistics was statistically significant at even the 10 percent level (in other words Eq. (10) of Section 2 was never rejected). What may have occurred is that the power of the test is quite weak. First, the number of observations in each pair of equations was relatively low (since only a portion of the data was used). Probably more importantly, other than the coefficients for Income and Liquid wealth the coefficients on the assets are usually not statistically different from zero. Because the values of the different assets do not affect consumption very much and their coefficients may just be reflecting random noise, it is not very surprising that their ratios do not differ in a statistically significant way.

The results change, however, when the two groups are separated. The results for the unconstrained group are similar to the results for the combined group discussed above. For none of the 28 different pairings can one reject the hypothesis that the ratios are the same for both equations. Again, this is not surprising since the assets have relatively little effect on consumption spending.

For the constrained group, however, the results are strikingly different. For eight of the 28 pairings, the hypothesis of equal ratios can be rejected (once at the 10 percent level or better, seven times at the 5 percent level). Vacation spending was most likely to reject the equal ratio test (four out of eight times) perhaps because vacations are more like a durable good than the other seven. The complete list of comparisons can be found in the table below.

	Groceries	Food Out	Charity	Dues	Entertain–ment	Gifts	Transport
Food Out	†						
Charity	†	**					
Dues	†	†	†				
Entertainment	†	†	†	†			
Gifts	†	†	†	*	†		
Transport	**	†	**	†	†	†	
Vacation	**	†	**	**	**	†	†

** Statistically Significant at 5% Level, * Statistically Significant at 10% Level, † Not Statistically Significant

As discussed earlier, constrained consumers' spending is more sensitive to changes in assets than is the spending of unconstrained consumers. Thus this group provides a better test of whether or not the composition of assets that individuals own affects not only the level but also the composition of their spending. For this group, the result is a very strong yes and so provides very strong evidence of framing.

5. Conclusion

This paper has four major results. First, spending seems to be very sensitive to changes in income but much less sensitive to changes in wealth. Second, close examination of the relation of wealth to consumption reveals a pattern in which individuals treat assets as not being fungible (i.e. equivalent to each other). Third, liquidity constraints affect consumption not as the conventional model predicts but in a manner consistent with the existence of either financial or psychological transaction costs. The affect of these constraints may also be evidence of a behavioral bequest motive. Finally, the amount spent on particular goods seems to depend not only on the individual's total resources but also on how those resources are split between different assets. Overall the patterns in the data provide evidence in favor of the behavioral life-cycle model. Any defense of the conventional life-cycle model will have to explain the anomalous results obtained here. Although that may prove possible, the behavioral life-cycle model which incorporates problems of self-control into a standard life-cycle framework, is a relatively simple explanation of the patterns found in this data.

Acknowledgements

I would like to thank Doug Bernheim, Hersh Shefrin, Bill Sundstrom, members of the Applied Micro Workshops at Stanford and U.C. Davis, participants at the Econometric Society meeting on Saving in New Orleans and participants at the Russell Sage Conference on Behavioral Theories of Savings for comments and suggestions on this paper. I would also like to thank Clint Cummins for much computer assistance.

References

Campbell, John Y., Mankiw, Greg N., 1989. Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence. NBER Working Paper # 2924.
Carroll, Chris, Summers, Lawrence, 1989. Consumption Growth Parallels Income Growth: Some New Evidence, Mimeo, Harvard.
Courant, Paul, Gramlich, Edward, Laitner, Edward, 1986. A Dynamic Micro Estimate of the Life-Cycle Model. In: Aaron, Henry, Burtless, Gary (Eds.), Retirement and Economic Behavior. Brookings Institute, Washington D.C.
Deaton, Angus, 1989. Saving and Liquidity Constraints, NBER Working Paper # 3196.
Flavin, Marjorie, 1981. The adjustment of consumption to changing expectations about future income. J. Political Economy 89, 974–1009.
Griliches, Zvi, Hausman, Jerry, 1986. Errors in variables in panel data. J. Econometrics 31, 93–118.
Hall, Robert, Frederick, Mishkin, 1982. The sensitivity of consumption to transitory income: estimates from panel data on households. Econometrica 50, 461–581.
Lars, Hansen., Kenneth, Singleton, 1983. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. J. Political Economy 91, 249–265.
Hayashi, Fumio, 1985. The effect of liquidity constraints on consumption: a cross-sectional analysis. Quarterly Journal of Economics 100, 183–206.
Hsiao, Cheng, Analysis of Panel Data, Cambridge University Press, Cambridge, 1986.
Hurd, Michael, 1987. Savings of the elderly and desired bequests. American Economic Review 77, 298–312.

Mankiw, N. Gregory, 1981. The permanent income hypothesis and the real interest rate. Economic Lett. 7, 307–311.

Mayer, Thomas, 1972. Permanent Income, Wealth and Consumption., University of California Press, Berkeley and Los Angeles.

Shefrin, Hersh, Thaler, Richard, 1988. The behavioral life-cycle hypothesis. Economic Inquiry, XXVI, 609–643.

Shefrin, Hersh, Thaler, Richard, 1981. Gift-Giving and Mental Accounts. Mimeo, Santa Clara University, 1981.

Skinner, Jonathan, 1993. Is Housing Wealth a Sideshow? NBER Working Paper # 4552.

Thaler, Richard, 1990. Anomalies: saving, fungibility and mental accounts. J. Economic Perspectives 4(1), 195–205.

Thaler, Richard, Quasi-Rational Economics. Sage Publications, Los Angeles, 1992.

Zeldes, Stephen, 1989. Consumption and liquidity constraints: an empirical investigation. J. Political Economy 97, 305–346.