The Relationship of Household Consumption and Saving to Income

by

Nicholas S. Souleles

B.S.E., Civil Engineering and Operations Research, Princeton University, 1988
B.A., Mathematics and Philosophy, University of Oxford, 1991

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the
Massachusetts Institute of Technology

September, 1995

Signature of Author . . . .
Department of Economics
June 27, 1995

Certified by . . . .
Olivier J. Blanchard
Professor of Economics
Thesis Supervisor

Certified by . . . .
Jerry A. Hausman
John and Jennie S. MacDonald Professor of Economics
Thesis Supervisor

Accepted by . . . .
Richard S. Eckaus
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

The Relationship of Household Consumption and Saving to Income

by

Nicholas S. Souleles

Submitted to the Department of Economics in June, 1995
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

Chapter 1. (jointly with Tullio Jappelli and Jörn-Steffen Pischke) Previous tests for liquidity constraints using Euler equations have frequently split the sample on the basis of wealth arguing that low-wealth consumers are more likely to be constrained. We propose an alternative test using direct information on borrowing constraints obtained from the 1983 Survey of Consumer Finances. In a first stage we estimate probabilities of being constrained that are then utilized in a second sample, the Panel Study of Income Dynamics, to estimate switching regression models for the Euler equation. Our estimates do not indicate much excess sensitivity associated with the possibility of liquidity constraints. However, we show that the conditional distributions of consumption in the constrained and unconstrained regimes look quite different and argue that these results are weak evidence that liquidity constraints affect consumption. Similar results are obtained using consumption data from the Consumer Expenditure Survey.

Chapter 2. Under the canonical life-cycle/permanent-income theory, consumption is not supposed to respond excessively to the predictable or transitory components of income. Previous tests of this theory on micro-data have usually had difficulty isolating these components. In contrast, this paper provides a clean test by focusing on income tax refunds, which are both predetermined and transitory. Using the Consumer Expenditure Survey, it finds significant and interpretable excess sensitivity. This sensitivity is due in part to sharply increased spending on nondurables at the time of refund-receipt by those likely to be liquidity constrained. However, there is also evidence of substantial increases in spending, mostly on durables, by those unlikely to be constrained; as well as in spending after the receipt of refunds (particularly while on trips in the summer), which is also unlikely to be due to liquidity constraints. The large magnitudes of these responses to refunds evidence a greater efficacy of fiscal policy than found by most previous studies. The responses also serve to evaluate some recent behavioral theories of saving, namely mental accounts and self-control/forced saving.

Chapter 3. Despite the high cost of college, there has been little study of the adequacy of households' savings and other resources available to pay for college. This paper gauges this adequacy, using the Consumer Expenditure Survey, by examining whether households are able to maintain their standard of living as they pay for college. The main finding is that households appear to maintain their consumption up to 6 months into the academic year, despite large expenses. This is consistent with the life-cycle theory of saving and consumption. Furthermore, households appear not to cut their consumption during the 6-9 months before the year starts. For households financing college out of savings, the implication is that their saving seems to have been fully underway at least this many months in advance of their expenditure in the fall. There is, however, evidence of a drop in consumption in proportion to college expenditures for households with children first beginning college. This is consistent with the view that such households learn to save with experience. But even for such households the effect on consumption is rather small in magnitude.

Thesis Supervisor: Olivier J. Blanchard
Title: Professor of Economics

Thesis Supervisor: Jerry A. Hausman
Title: John and Jennie S. MacDonald Professor of Economics

Acknowledgements

I would like to thank first my supervisors, Olivier Blanchard and Jerry Hausman, for their insight, guidance, and kindness. Special thanks go to Steve Pischke for nearly daily advice and encouragement over the past year. The generous input of Ricardo Caballero is also greatly appreciated. More generally, I thank my teachers at M.I.T. for exceptional instruction.

My classmates and friends at M.I.T. I thank for much discussion and support. The regulars in the computer room over the last year (particularly Toshi Baily, Li-Lan Cheng, John Dippold, Jonathan Parker, and Stacey Tevlin) were especially helpful. I would like to single out my "study partners," Dave Gross, Christine Jolls, and Jack Porter, for their long-standing friendship and tuition. My housemates Jason, Michael, and David deserve thanks for their support and patience.

I would like to acknowledge financial support from the National Science Foundation. The N.S.F. also provided funds for the computer, trader, on which I often worked.

My greatest thanks go to my family. This thesis is dedicated στους γονεις μου, με αγαπη και ευχαριστιες.

Contents

Abstract 2

Acknowledgments 4

Chapter 1. Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources 6

Chapter 2. The Response of Household Consumption to Income Tax Refunds 66

Chapter 3. College Tuition and Household Savings and Consumption 120

Chapter 4. Conclusion 151

Chapter 1. Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources[1]

I. Introduction.

Many applied economists now agree that the rational-expectations/permanent-income model of consumption in its simplest form is inconsistent with the data. There is much less agreement, however, as to the source of these inconsistencies. Is it that preferences are more complicated than in the simplest model? Or are other features of the model at fault, like the assumption of perfect credit markets? The existence of liquidity constraints, in particular, has important policy implications, for example with regard to taxation, financial market liberalization, growth and welfare (Hubbard and Judd, 1986; Jappelli and Pagano, 1995). It is therefore important to distinguish whether the empirical rejection of the model results from borrowing constraints or from some other source.

The most influential microeconomic tests addressing the issue of liquidity constraints have relied on sample-splits based on households' assets. However, assets alone are a rather imperfect predictor of potential constraints. Unfortunately, of the leading U.S. data sets, those with consumption data do not have direct indicators of credit constraints, and those with such indicators lack information on consumption. In this paper

[1]Written with Tullio Jappelli and Steve Pischke. We thank Jon Gruber for making an extract of PSID variables available to us; Steve Zeldes for extensive comments on his sample construction and useful discussions about the sample splitting approach; Josh Angrist for detailed discussions on the econometrics; Steve Cecchetti, Don Cox, Jon Gruber, and Jim Poterba for helpful comments on a previous draft; and NATO for financial support.

we combine data from the 1983 Survey of Consumer Finances (SCF) with data from the Michigan Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) using two-sample instrumental variables techniques.

We use a self-reported indicator of liquidity constraints in the SCF to relate the probability that a household is constrained to demographic variables, and then impute the probabilities of credit constraints in each wave of the PSID. We then estimate the Euler equation for consumption in the PSID as a switching regression using the stochastic sample separation information provided by these predicted probabilities. Compared to the asset-based sample splitting procedure of Zeldes (1989) and Runkle (1991), our approach relies on different information to assess the likelihood of a constraint and explicitly acknowledges the uncertainty associated with identifying liquidity constraints in the estimation. It therefore complements this earlier work. In addition, we provide similar evidence using data from the CEX.

We find little evidence that liquidity constraints affect the conditional mean of the distribution of the growth rate of consumption in either the PSID or the CEX. However, our results tend to be rather noisy.

Liquidity constraints have additional implications for the entire conditional distribution of consumption, beyond those for the mean which are exploited in Euler equation estimates. This was pointed out most clearly by Deaton (1991), but has not been the focus of empirical investigations so far. Our setup naturally lends itself to an analysis of distributional issues by applying quantile regressions to the switching regression framework. The empirical results resemble those from theoretical simulations with income uncertainty and liquidity constraints, and accord with intuition. We take this as weak evidence that liquidity constraints do have an impact on intertemporal allocations of consumption.

In Section II we review the recent literature on liquidity constraints and Euler equations, motivate our methodology and describe how it differs from previous approaches. In Section III we describe the data used in the analysis and compare the indicator of liquidity constraints available in the SCF to asset-based sample separation rules previously used in the PSID. Section IV presents the results of estimating the consumption Euler equation by switching regressions in the PSID, while section V reports similar results for the CEX. Section VI motivates the quantile regression approach, presents the theoretical simulation, and compares it with the results obtained from the data. Section VI concludes.

II. Literature Review and Motivation.

While the central implication of Hall's (1978) rational-expectations/permanent-income hypothesis is simple and powerful, there has been ample research documenting its empirical failure. The model has failed strongly in aggregate time series studies (e.g., Campbell and Mankiw, 1989) but much, if not all, of this failure can be explained by aggregation issues.[2] Studies relying on microeconomic data have not had as clear results. For example, Hall and Mishkin (1982) found excess sensitivity in the Panel Study of Income Dynamics, while Altonji and Siow (1987) did not.

The tests that reject the permanent income model do not point directly to the reason why the model fails. In the early literature following Hall, excess sensitivity was generally held to be due to the presence of credit market imperfections, in the form of interest rate differentials or credit rationing (Flavin, 1985; Hubbard and Judd, 1986; Hayashi, 1987). In fact, credit constraints break the powerful implication of Hall's model: current consumption is no longer a sufficient statistic for everything the consumer knows about the future. This leads to an intertemporal dependence in the Euler equation for consumption. However, such dependence would not have to stem from the budget constraint, as pointed out by more recent literature. Similar dependence can be generated by non-separable preferences, durability of goods or slow adjustment of consumers. While the empirical implications for the Euler equations of all these extensions are rather similar to liquidity constraints (Browning, 1991, Attanasio, 1994), intertemporal

[2]See Gali (1990), Attanasio and Weber (1993), Goodfriend (1992), and Pischke (1995).

dependence originating from the preference side has vastly different policy implications than credit constraints.[3]

Recent empirical work has therefore tried to incorporate additional information to detect the presence of liquidity constraints. One such approach, used by Zeldes (1989) and Runkle (1991), relies on an assets-based sample separation rule.[4] They argue that assets can be used to separate households that are likely to be liquidity constrained (the low-wealth group) from those that have access to credit markets or have no need to borrow (the high-wealth group). If the only violation of the model is due to the existence of liquidity constraints, excess sensitivity should arise only in the low-wealth group. If instead excess sensitivity is due to preference misspecification, there is no reason to believe that the results for the two groups should differ. Zeldes indeed finds a violation of the theory in the low-wealth group: the coefficient of lagged income in the Euler equation is significant and twice as large (in absolute value) as for the high-asset group. Runkle, on the other hand, does not find significant effects of lagged income in the Euler equations for either of the two groups.

While adding outside information improves the power of the test for liquidity constraints and ties potential rejections more clearly to a specific alternative, splitting the sample on the basis of wealth has a number of drawbacks. First of all, a simple split on the basis of wealth is a good indicator of liquidity constraints only if there is a roughly monotonic relation between the two. But poor households are not necessarily identical to constrained households. For instance, households which are able to borrow without full collateral have negative wealth by definition, but are obviously not credit constrained. According to Wolff (1994), about 15 percent of the SCF sample has negative net worth (including housing, real estate and pension wealth as part of assets). Even considering

[3]Meghir and Weber (1993) exploit the fact that liquidity constraints should affect all commodities similarly while the same is not true if preferences are non-separable. Comparing within-period marginal rates of substitution and intertemporal Euler conditions, they find no evidence for the existence of liquidity constraints.

[4]Zeldes (1989) attributes this method originally to Juster and Shay (1964).

measurement errors, this indicates that a significant fraction of the population might be able to borrow without full collateral. Figure 1 plots a nonparametric regression function of the indicator for a binding borrowing constraint which we utilize below on the net wealth to income ratio using the SCF. It indicates that households with negative net wealth are more likely to be constrained than households with positive assets. But the relationship is indeed not monotonic. Households are most likely to be constrained if their net wealth is close to zero, but the likelihood initially falls when net wealth becomes negative.

Second, empirical measures of wealth are bound to be highly imperfect because assets and asset income are often poorly measured. Zeldes' and Runkle's sample split is partly based on a direct question in the PSID, asking whether households currently have liquid assets in excess of two months' income. For other survey years lacking this question, a corresponding variable is created based on information on asset income. Liquidity constraints are unlikely to be perfectly correlated with this particular indicator. For example, it is not likely that two month's income is the exact cutoff below which assets are low enough to indicate a binding constraint. Survey measurement error may further obscure the relationship. If this is the case, the high-asset sample contains some constrained households, while the low-asset group is contaminated by the presence of unconstrained households. This reduces the power of the statistical test since it moves the coefficients on lagged income closer together for the two groups. As will be seen, we instead classify an individual as liquidity constrained if he or she was refused loans or discouraged from borrowing. Even though this criterion is subject to problems of its own, it uses different and more direct information on liquidity constraints than the sample split based on assets.

A further problem, inherent to any sample separation rule, is that in empirical work the Euler equation is usually linearized, and so omits second and higher order terms of the conditional distribution of consumption growth. As was pointed out by Carroll (1992), this may create a correlation between consumption growth, lagged income and assets leading to spurious evidence in favor of liquidity constraints. Our method is subject to a similar criticism.

To contrast the exogenous sample splitting technique to our own, it is useful to consider the Zeldes and Runkle model as a switching regression. There are two separate Euler equations

$$(1) \qquad \Delta\ln c_{it+1} = \alpha_1 + \beta_1'X_{it+1} + \gamma_1 \ln y_{it} + \varepsilon^1_{it+1} \quad \text{if } \pi'Z_{it} >= u_{it}$$
$$\Delta\ln c_{it+1} = \alpha_2 + \beta_2'X_{it+1} + \gamma_2 \ln y_{it} + \varepsilon^2_{it+1} \quad \text{if } \pi'Z_{it} < u_{it},$$

the first equation referring to constrained households and the second equation to unconstrained households. The vector X_{it} includes preference shifters, and possibly the interest rate; y_{it} is disposable income. The instruments Z_{it} and the random variable u_{it} indicate whether a household is constrained or not in period t. The split of Zeldes and Runkle is a special case of this, treating the asset-income ratio as an indicator $L_{it} = I(\pi'Z_{it}>=0)$, where I(.) is an indicator function, Z_{it} contains only the asset-income ratio and the cutoff point (two months' income), and var(u) = 0.

In the general case, if the permanent-income model holds and the Euler equation is correctly specified, γ_2 should be zero; if liquidity constraints are responsible for excess sensitivity, γ_1 should be negative and significantly different from zero. The test is even somewhat more robust. Even if γ_2 is not exactly zero, due to non-time-separable preferences, for example, then $\gamma_1 < \gamma_2$ would still indicate the presence of liquidity constraints, if the non-separability affects both groups in a similar fashion.

To address the measurement problems associated with the asset-income ratio, one could recognize explicitly that L_{it} is only an imperfect predictor of liquidity constraints. This leads to a switching regression model analogous to that proposed by Lee and Porter (1984) in a different context. Hajivassiliou and Ioannides (1991) extend this idea one step further. They consider not only uncertainty about regime classification, but also relate to demographic characteristics the cutoff of the asset-income ratio below which the liquidity constraint is assumed to be binding. In terms of the formulation above, these variables become part of the Z_{it} vector. They estimate the Euler equations with a two-step procedure and with full information maximum likelihood. Both cases involve a binary regression of the asset-income ratio on demographic variables and lagged values of the dependent variable. Their approach acknowledges the fact that the sample separation rule is a rather imperfect indicator of liquidity constraints, but retains the assumption that liquidity constraints vary monotonically with the asset-income ratio. As we pointed out, this rules out the possibility that individuals can borrow without full collateral and therefore have negative net worth.

Garcia, Lusardi and Ng (1995) proceed in a similar fashion to Hajivassiliou and Ioannides, but avoid the latter problem using a switching regression model without additional sample separation information. The instruments Z_{it} are again demographic variables; however, the instruments are not related to any prior indicator of liquidity constraints. Instead, the instruments are used directly to find differentials in the slope coefficients of the Euler equations for the two regimes.

There are various limitations to this approach. First, it is not clear which regime should be labeled as the constrained regime and which one as the unconstrained. Garcia, Lusardi and Ng propose to identify the two regimes by comparing the signs of the coefficients of the Z_{it} vector with the logit coefficients of Jappelli (1990), which relate the indicator of liquidity constraints in the SCF to demographic variables. We formally incorporate this eye-balling procedure into our estimation method. A second limitation of their procedure is that its demand on the data is extremely high; nevertheless, they find significant excess sensitivity in one regime but not in the other. Maddala (1986) reports that disequilibrium studies without additional sample separation information have also frequently found surprisingly good results, while Monte Carlo experiments reveal that such results are not likely to be expected.

In this paper, as in Zeldes and Runkle, we use additional information to identify liquidity constrained households, but avoid some of the problems of the previous literature. We use a direct question on liquidity constraints available in the SCF. If this information were also available in the PSID, we could easily split the PSID sample, and apply a switching regression model with known and complete sample separation. Since information on liquidity constraints is only available in the SCF, our approach is best thought of as an application of two sample instrumental variables techniques,[5] although instrumental variables estimation is not necessary for consistency but just serves to link the two samples. It is useful to rewrite the switching regression model in (1) in a single equation as

$$(2) \qquad \Delta \ln c_{it+1} = \theta_2' W_{it+1} + (\theta_1' - \theta_2') W_{it+1} L_{it} + \varepsilon_{it+1},$$

where $\theta_j' W_{it+1} = \alpha_j + \beta_j' X_{it+1} + \gamma_j \ln y_{it}$, and L_{it} refers to the indicator of a binding liquidity constraint from the SCF. Next, take expectations conditional on Z_{it}, the instruments used to predict liquidity constraints. Note that Z_{it} includes the variables in W_{it+1}, therefore

$$(3) \qquad E(\Delta \ln c_{it+1} | Z_{it}) = \theta_2' W_{it+1} + (\theta_1' - \theta_2') W_{it+1} E(L_{it} | Z_{it}) + E(\varepsilon_{it+1} | Z_{it}).$$

The orthogonality condition for the model is $E(\varepsilon_{it+1} | Z_{it}) = 0$, i.e. all variables in Z_{it} and not just those in W_{it+1} need to be orthogonal to the Euler equation residual. Since L_{it} is not available in the PSID or CEX, we need to add a first stage to the model of the form $E(L_{it} | Z_{it}) = \pi' Z_{it}$, which will be estimated within the SCF by regressing the indicator of liquidity constraints on the demographic variables Z_{it}. Equation (2) will then be estimated in the PSID by replacing L_{it} by $\pi\char`^' Z_{it}$, using the predicted π^ ("π-hat") from the first stage. We demonstrate consistency of the estimator in an appendix and show how to construct standard errors.

We also use data from the Consumer Expenditure Survey (CEX) in the second stage, which introduces some additional complications. Note that W_{it+1} will include variables like age and changes in family size to proxy for changes in preferences over the

[5]See Angrist and Krueger (1992) and Arellano and Meghir (1992). Other applications of two-sample techniques in the context of consumption are Lusardi (1995), Carroll and Weil (1994), and Garcia, Lusardi, and Ng (1995). Carroll and Weil also use the PSID and SCF in conjunction.

life-cycle. We will use quarterly changes in consumption in the CEX. Since quarterly changes are not available in the SCF, we cannot construct the change in family size. We circumvent this problem in the PSID by using three year differences in consumption and in the corresponding W_{it+1} variables. Such changes can be constructed by exploiting the panel character of the 1983 and 1986 SCF waves. In using the CEX we have to live with the fact that W_{it+1} is not a proper subset of Z_{it}. This complicates identification of the model somewhat. The simplest way of getting identification is by assuming $E(L_{it} | W_{it+1}, Z_{it}) = E(L_{it} | Z_{it})$. Then, taking expectations of (2) conditional on both W_{it+1} and Z_{it} gives

$$(4) \qquad E(\Delta \ln c_{it+1} | W_{it+1} Z_{it}) = \theta_2' W_{it+1} + (\theta_1' - \theta_2') W_{it+1} E(L_{it} | Z_{it}) + E(\varepsilon_{it+1} | W_{it+1} Z_{it}).$$

In this case, we can again simply replace L_{it} by $\hat{\pi}' Z_{it}$ so that the same two step estimator results.[6] Notice, however, that the assumption $E(L_{it} | W_{it+1}, Z_{it}) = E(L_{it} | Z_{it})$ implies that changes in family size will not help predict whether a household is liquidity constrained, given the variables in Z_{it}, which are other demographic characteristics. This may be a strong assumption when liquidity constraints are directly related to the composition of the family, for example when a child leaves for college. We have also not derived consistent standard errors for this model. Instead we note that unadjusted standard errors will be consistent under the null hypothesis that liquidity constraints do not affect the Euler equation ($\theta_1 = \theta_2$). Thus, we can still carry out consistent tests of the null hypothesis.

We define a liquidity-constrained household as one which gave an affirmative answer to the following question in the 1983 SCF:[7] "In the past few years has a particular lender or creditor turned down any request you (or your husband/wife) made for credit or have you been unable to get as much credit as you applied for?" Some

---

[6]We thank Josh Angrist for very helpful discussions on these issues.

[7]See Avery and Kennickell (1988) for a description of the SCF.

consumers may not apply for credit because they think that, if they did, they would be turned down. So we add to the group of liquidity constrained these "discouraged borrowers," i.e. households who reported an affirmative answer to the question: "Was there any time in the past few years that you (or your husband/wife) thought of applying for credit at a particular place but changed your mind because you thought you might be turned down?" Excluding from the group of credit constrained those who reapplied for a loan and received the desired amount results in 718 households out of a total of 3656 (19.6 percent of the sample) who reported themselves as being liquidity constrained.[8]

Several studies have adopted the same definition of liquidity constraints. Jappelli (1990) describes the characteristics of households for which this constraint binds. With some identifying assumptions, Perraudin and Sorensen (1994) estimate the separate determinants of the supply and demand for loans. Cox and Jappelli (1993) and Duca and Rosenthal (1991) estimate that desired debt for those who reported themselves as liquidity constrained exceeds actual debt. Gale and Scholz (1994) find that borrowing constraints substantially reduce contributions to IRAs. Gropp, Scholz, and White (1995) find a small of effect of state personal bankruptcy laws on the probability of being constraint. Eberly (1994) uses the constraint indicator to split the sample and to test for excess sensitivity in Euler equations for the stock of automobiles.

One objection to using our definition of liquidity constraints is that the questions may pertain mostly to consumers who intend to borrow for the purchase of a house, car or other durable which serves as collateral. The Euler equation for the constrained, by contrast, is about consumers who are unable to obtain consumption loans. Even so, we feel that our procedure yields sensible results. Banks tend to look at similar personal characteristics for personal consumer loans as for mortgage applications. Even in the case of a mortgage application, the bank is mainly interested in borrowers' ability to repay because foreclosures are typically associated with substantial costs. Thus, our first stage regression should capture the correct relation between demographic variables and the

[8]For comparability with the PSID and CEX, below we do not use the full SCF sample, so that this fraction will be slightly different there.

probability of being credit constrained. In fact, our first stage results reported below qualitatively resemble those reported by Boyes, Hoffman, and Low (1989) using data on applications made to a particular credit card company.

Another shortcoming of our approach is that the information on liquidity constraints comes from a single cross-section, whereas we impute the constraint probabilities in the PSID for the whole 1971-86 period. The first-stage model for liquidity constraints is a reduced form reflecting not only desired consumption, but also the nature of the constraint imposed by lenders on consumers. It is reasonable to assume that the consumption rule changes slowly at best; but the hypothesis of unchanging behavior of financial intermediaries is less tenable. The supply of credit, in fact, responds to a variety of factors, such as the monetary and regulatory regimes and the institutional developments in the credit market. For instance, for a given set of characteristics, a household's constraint in the 1970s might have been more severe than in the 1980s, when consumer credit became more liberal and credit card use increased. Unfortunately, there is no evidence on borrowing constraints for the early 1970s similar to the question in the 1983 SCF.[9] Given our operational definition of liquidity constraints, there is no satisfactory way of incorporating any "supply side" effect in the analysis. Thus, in using the estimated coefficients from the SCF in years other than 1983, we must assume that lenders' behavior remains constant through time. If the relationship between household characteristics and borrowing constraints changed over time, the precision of the predicted probabilities of being refused credit is reduced, and our test biased against detecting liquidity constraints.

---

[9]In future research we plan to use the 1989 SCF which repeated the 1983 questions about credit availability (the questions were not asked in the 1986 SCF). This will allow us to take into account some of the changes in the screening procedures used by banks in the allocation of credit, and provide a more reliable indicator of credit constraints. Since the 1989 SCF has a panel section, we will also be able to characterize liquidity constraints over time by exploiting longitudinal information. Unfortunately, the SCF codes to match 1983 and 1989 households are not yet available. Preliminary work using the 1989 SCF indicates that the fraction of constrained households is roughly the same (19 percent) as in 1983. We regard this as indirect evidence in support of the validity of the survey questions.

III. Data and Measures of Liquidity Constraints.

In the construction of the PSID sample and of the relevant variables we follow Zeldes (1989) closely. The main difference is that he uses data up to 1982, while our sample extends to 1987. Here we describe broad characteristics of the data; details are given in the Appendix. The main drawback of the PSID is that the only consistently available consumption measure is food expenditures. In order to estimate the Euler equation we must therefore invoke separability between food and other expenditures. The validity of this assumption is questionable. Meghir and Weber (1993), using data from the Consumer Expenditure Survey, find that food at home, transport and services are not weakly separable either from each other or from food out of home, clothing, and fuel. Lusardi (1995), using the same data set, finds instead that Euler equations are very similar for food and for broader aggregates of nondurable expenditures; in each case she rejects the Euler equation. We also use data from the CEX and find very similar results for food and broader aggregates of nondurable spending.

The final PSID sample contains consumption changes for 1971 and for 1974 to 1984. Other years are lost in forming changes of the variables, taking into account the timing of questions in the PSID, and because the food expenditure question was not asked in 1973. Income and wealth measures in the PSID are deflated by the price index of personal consumption expenditures (base year 1982-84); food expenditures are deflated by the price indices for food at home and away from home. Unlike Zeldes we include the low income subsample in most of our analysis and provide weighted estimates, but selected results using just the representative Census sample are reported as well.

We try to match the SCF sample and variables as closely as possible to their PSID counterparts above. The high-income subsample in the SCF is excluded. Also excluded are the 159 ''uncleaned'' observations with mostly missing values for which no imputations were made. We only use households which were reinterviewed in 1986 and exclude households with changes in marital status in the intervening period. The final SCF sample includes 2,139 observations; details are again given in the Appendix.

Each of the three sample splits proposed by Zeldes is based on the ratio of assets to current disposable income. In most years the PSID does not ask directly about the level of assets. In some years a question is available asking ''Do you have any savings [that] amount to as much as two months' income or more?'' This cutoff is the basis for the first sample split that Zeldes uses: everyone with assets worth less than two months' income is included in the constrained group. The question is not available in all years. In these cases the level of assets is estimated by dividing asset income by the current interest rate, i.e. by ''blowing up'' asset income. The second split is more stringent. The constrained group consists only of households with no asset income, and the unconstrained group of those with assets worth more than six months' income. The intermediate part of the sample is discarded. The third split adds a measure of net housing wealth to liquid assets; otherwise it is similar to the first split, also using the two months' income cutoff. We created similar sample splits in our larger PSID sample.

We similarly matched the SCF to the CEX. In the CEX ''consumer units'' are interviewed four times, three months apart. We use income from the first interview, which refers to the 12 months prior to the interview date. In order to make sure that this income does not contain information on period t+1, we use the consumption change between the third and the second interview. All other variables included in our regression, except for changes in family size, also come from the first interview. Our data span the period 1980 to 1991.

In addition to food expenditures, we constructed two broader measures of spending, nondurable consumption and strictly nondurable consumption. Nondurable consumption basically mimics the CPI's category for expenditures on nondurables and services. It includes such items as clothing with some durability. Strictly nondurable consumption excludes these categories as well as items on which expenditures occur only very infrequently. The major components of strictly nondurable consumption are food; household operations, including monthly utilities and small scale rentals; apparel services and rentals; transportation fuel and services, including public transport; personal services; entertainment services and high frequency fees.

The income and asset variables in the CEX are comparatively poor (see e.g. Lusardi, 1995). In particular, total family income is computed by the Bureau of Labor Statistics as the sum of various income sources. In summing income, individual income items were sometimes set to zero if the household response to this item was missing. This results in a cross-sectional distribution of family incomes with many implausibly small amounts. In order to avoid the most blatant misrepresentations, we excluded households with zero or missing earnings. The reason zeros are excluded as well is that these are often not well distinguishable from missing values.

The composition of the CEX sample differs slightly from that of the PSID. Because the SCF sample should correspond to the same underlying population as that used in the second stage estimation, we adjust the SCF sample to be used in conjunction with the CEX accordingly. Because we cannot construct changes in family size at the quarterly level we just rely on the 1983 cross section of the SCF. We therefore do not exclude SCF households with marital changes and households, who were not interviewed in 1986. On the other hand, we exclude households with zero earnings as in the CEX. The CEX also does not interview households in rural areas so that we make a similar sample selection in the SCF. The final SCF sample has 2,328 observations.

In the remainder of this section we compare sample means of demographic variables in the PSID and in the SCF and CEX, present various measures of assets available in the first two surveys and compare the resulting sample splits with our direct indicator of liquidity constraints. For comparison with the SCF, in some cases we present data from the 1984 PSID wave, which contained direct questions on the level of assets.

Table 1 reports sample means for income and various demographic variables for the PSID, the CEX, and the SCF. The SCF sample is the one comparable to the PSID. Column (1) contains means for the pooled PSID sample over the years 1970 to 1984. Column (2) singles out the 1984 PSID wave with asset stocks. This year should also be more comparable to the SCF, which was conducted in 1983. The SCF sample means are displayed in column (4). The means in the various data sets are reasonably similar and the existing differences do not seem to follow any particular pattern. This is comforting, as the two samples used in the estimation need to stem from the same population.

The final two columns in Table 1 compare sample means for constrained and unconstrained households in the SCF using the indicator of liquidity constraints. There are pronounced differences between these two groups. Constrained households have less income, are younger and more likely to be single and non-white than unconstrained households. 16.3 percent of the households in the displayed SCF sample are constrained, which is fewer than in the full 1983 cross section of the SCF. This is not surprising as the households reinterviewed in 1986 and without marital changes will be more stable households, who are also more likely to obtain credit. This is a much lower fraction than Zeldes classifies as constrained (about two thirds).

We next provide some evidence on the quality of the asset information used by Zeldes and its implications for classifying households as liquidity constrained. The upper panel of Table 2 reports percentiles of the distribution of liquid assets in the 1984 PSID and in the SCF. While the latter does not provide direct questions about the asset-income ratio, it contains detailed information on the level of assets and on asset income. We use these variables to mirror the available data in the PSID as closely as possible. We report only data for the 1984 PSID because we want to compare the ''blown-up'' measure of wealth, derived from asset income, with the direct measures of actual wealth stocks. Asset income in the PSID refers primarily to interest, dividends, and rents; following Zeldes households with substantial business income are excluded. We use two alternative measures of asset levels for comparison. The ''narrow'' definition excludes bonds and stocks because these include holdings in IRAs. The reason for this is that it is unclear whether respondents include interest and dividend income from IRAs when reporting income from assets. Furthermore, IRAs are not very liquid and it is hard to borrow against them. However, the narrow definition excludes also any other holdings of bonds and stocks; these are part of liquid assets, whose income should be reported by PSID respondents. Therefore we also compute a ''wide'' definition of assets that includes all bonds and stocks (inclusive of IRAs).

Comparison of either the narrow or wide definitions of asset levels with the blown-up measure reveals similar patterns. The left tail of the distribution (up to the 50th percentile) of the blown-up measure lies below the stock distributions, suggesting that asset income may be under-reported at low levels of assets. At high levels of wealth (above the 50th percentile) the pattern reverses: the right tail of the distribution of the blown-up measure is thicker than that of the stocks. This might derive from the fact that asset returns vary across households, counter to the constancy assumed in the asset-income inflating procedure. The pattern is the same if one compares the blown-up measure of assets with asset stocks in the SCF.[10] The lower panel of Table 2 shows that one obtains similar results if one adds to each of the three definitions of assets an estimate of the stock of housing.

The potential biases in constructing sample splits arise from the differences in the left tail of the distribution, because the cutoff-point to separate low- and high-wealth groups occurs at fairly low levels of assets. Therefore the comparison above between the various wealth measures suggests that splits that use asset income to impute wealth will tend to overstate the size of the constrained group.

The upper panel of Table 3 shows how sample splits based on a blown-up measure of assets may affect Zeldes' test for liquidity constraints. The panel refers to the 1984 PSID and contrasts the sizes of the constrained and unconstrained groups resulting from splitting the sample according to the actual stock of liquid assets (narrow definition), with the sizes resulting from splitting according to the blown-up asset income (for brevity, "Zeldes' splits"). The samples for this exercise are slightly different from the ones used elsewhere in the paper and refer more closely to the sample used by Zeldes.[11] Each row in Table 3 reports, respectively, the sample size, the number and fraction of households that are constrained according to both sample splits, constrained according only to Zeldes' split, constrained according only to the asset stock split, and unconstrained in both cases. (The percentages in each row therefore sum to 100.) Following Zeldes, constrained and

[10]Because of the way assets are categorized, the definitions in the SCF not exactly identical to the PSID; see the Appendix for details.

[11]Essentially, for Table 3 we drop the requirement that three year changes in the data are available and we do not use the poverty subsample in the PSID. Also, the entire 1983 cross section of the SCF is used.

unconstrained groups obtained with three different splits are reported: split 1 defines as constrained those with liquid assets below two months' income; split 2 those with liquid assets equal to zero; and split 3 those with net worth below two months' income.

The interesting columns are (3) and (4). Column (3) indicates that Zeldes' constrained group is contaminated by people who instead possess considerable amounts of assets. The extent of this contamination is substantial, especially for liquid assets: 14% of the sample is classified as constrained according to split 1 or 2, but unconstrained if the reported stock of assets is used to split the sample. On the other hand, column (4) shows that the contamination of the unconstrained group with households reporting low asset stocks is less severe, in particular for splits 2 and 3 (2 percent). In sum, even if the probability of being liquidity constrained were directly related to the level of assets, using the available asset income in the PSID leads to substantial misclassification.

The lower panel of Table 3 refers to the SCF and presents similar tabulations of the direct indicator of liquidity constraints against Zeldes' splits. The extent of misclassification of each of Zeldes' splits is much larger. Splits 1 and 2 are relatively good proxies for identifying the unconstrained, but poor proxies for the constrained: 40% of those having access to credit or not interested in borrowing are in fact included in the constrained sample (column 3). The contamination of the unconstrained group in splits 1 and 2 is slight; only 5% of the sample reports being denied credit or discouraged from borrowing and yet has assets in excess of two months' income.[12] Split 3, using net wealth, works much better in sorting out constrained households: only 17 percent of the sample constrained according to split 3 reports not being denied credit. However, split 3 does not do quite as well in identifying the unconstrained (9 percent are misclassified).

These patterns are obviously conditional on the fact that credit status is a superior indicator of liquidity constraints than assets-based cutoff points. Given this condition, Zeldes' testing strategy is inefficient. Using splits 1 and 2, only about a quarter of the

[12]In part, this may reflect that some rejected applicants or discouraged borrowers would have liked to borrow to finance the purchase of durables or housing while holding some liquid assets accumulated for a downpayment.

observations with low assets is really unable to borrow (0.15/(0.15+0.40)); about one third using split 3. This implies that the coefficient of lagged income in Zeldes' regression should be three or four times as large, in absolute value, than what he finds in the data.

Table 4 reports linear probability models of being liquidity constrained using the SCF credit indicator (columns 3 to 5), which will later be used to impute the probability of being constrained in the PSID and CEX. For comparison, we report analogous estimates using as the dependent variable Zeldes' split 1 in the 1971-84 PSID (column 1) and in the SCF (column 2). We report two sets of results. The basic specifications in columns (1), (2), and (4) include standard demographics (age, sex, marital status, race, dummies for family size, and changes in family size), five dummies for education, three regional dummies, and the log of household disposable income. Except for income the definitions of these variables in the SCF and in the PSID are basically the same. In column (3) we report a regression based on a smaller set of regressors, only the variables present in the second stage Euler equation (age, changes in family size, and income), and race and region dummies as excluded instruments. In column (5), we add variables on the labor market status of the head to the basic set of regressors: employment, years of full time work experience, union membership, industry and occupation dummies. These variables represent potentially useful information to identify liquidity constraints; for instance, loan applications always contain a section on employment status. There is a certain compatibility problem with these variables, however: the variables used to predict liquidity constraints should be predetermined for the second stage, but current employment status maybe correlated with income and therefore with consumption growth. Since the SCF was a pure cross-section in 1983, we cannot construct lagged variables for the first stage equation. Our strategy is to introduce employment variables that do not change quickly over time (unlike hours worked). However, notice that the same applies to family income which is also included in the base specifications. Income is needed as a first-stage regressor because it is also present in the Euler equation to allow us to test for excess sensitivity.

The pattern of results in Table 4 is similar to those found by Jappelli (1990). Constrained households tend to be younger, non-white, include larger families, and are more likely to live in the West. The education variables have little impact on the probability of being a rejected loan applicant or discouraged borrower once income is controlled for. Income is significant but not of overwhelming importance. This is likely to result from the fact that richer households tend to ask for larger loans, and lenders are primarily interested in obligation ratios (the ratio of loan to income) or ability to pay (see e.g. Munnell et al., 1992). The employment variables are not individually significantly different from zero, although their signs accord with intuition.

In both the PSID and SCF, the coefficients obtained using as the dependent variable the indicator function based on the asset-income ratio (split 1) are somewhat dissimilar from those obtained with the direct constrained indicator. The schooling dummies matter for the split variable while the regional effects are unimportant. The coefficient on income is much stronger, but income is one of the variables used in the construction of the split. Division bias generated by measurement error may therefore bias the results.

IV. Euler Equation Estimates in the PSID.

In this section we present estimates of the Euler equation in the PSID. The specification and estimation strategy differ somewhat from Zeldes'. Unlike Zeldes, we do not perform fixed effects estimation. The reason is that there is not enough variability in the imputed probabilities of liquidity constraints within individual households over time. The main variation, in fact, stems from households' aging; but age is part of the Euler equation, because it captures changes in preferences directly related to the intertemporal allocation of consumption. Controlling for age, the interactions of the predicted probabilities of a constraint with the Euler equation variables (the term $W_{it+1}\hat{L}_{it}$ in our notation above) are basically not identified anymore within households. We must therefore rely on the variability of the constraint probabilities across households, and assume that individual effects, like differential discount rates or differences in the expected variability of consumption, are uncorrelated with the probability of being liquidity constrained. While this is a strong assumption, it is unlikely that these demographic differences follow the same pattern as liquidity constraints.[13]

We also do not include time dummies in the estimated Euler equation. The reason for this is mainly due to computational difficulties in constructing the variance-covariance matrix of the estimates. However, given that we have 10 to 12 years of data including at least one full recession in each data set, macro effects should largely average out in the estimation. We also estimated similar models with time dummies and the coefficients look hardly different. A further difference with respect to Zeldes concerns the variables included in the Euler equation. Instead of using his measure of changes in food needs (capturing changes in family composition) we directly control for the change in the number of adults and the change in the number of children. We found these variables to be superior indicators for the food consumption profile. We omit the interest rate from the Euler equation; this avoids the need to use instrumental variables in our basic formulations for the Euler equation since all regressors are part of the household information set. We control for age and age squared in the Euler equation. These variables are introduced in the previous stage, and we want to rule out the possibility that the imputed probabilities capture age-related changes in preferences. Finally, we include lagged disposable income to test for excess sensitivity.

An additional difference is that we use three year changes in food consumption in the PSID rather than one year changes as Zeldes. The reason for this is compatibility with the 1983-86 SCF. Using three year changes we are able to construct analogues to all the regressors in the SCF. Given that the Euler equation error follows a martingale, using overlapping three year changes of consumption will introduce moving average errors of order 2 (abstracting from additional within year time-aggregation). Because there are missing data so that not all households are present for the same years, the autocorrelation in the errors will be household specific. We adjusted the covariance matrix for this household specific autocorrelation pattern. Since the pattern is known given the years for

---

[13]For example, Hajivassiliou and Ioannides (1991) and Garcia, Lusardi, and Ng (1995) also do not control for fixed effects in the Euler equation for consumption.

which we have data for a specific household, we also performed GLS estimates but these turned out to be almost identical to the OLS estimates. We present OLS estimates in the tables below.

In order to compare our results to the previous literature, we present in Table 5 the coefficients of an Euler equation estimated on the full sample and estimated separately for the low and high wealth samples, using Zeldes' split 1. Column (1) shows the results for the full sample. Columns (2) and (3) refer to the high wealth sample and columns (4) and (5) to the low wealth sample. We included the poverty subsample of the PSID in the estimation since it will contain many households which are likely to be liquidity constrained. To adjust for the oversampling and to achieve comparability with the representative SCF we weighted the second stage estimates by the family weights. These estimates are shown in columns (2) and (4) labeled "poverty." For more comparability with Zeldes' estimates we also present unweighted estimates excluding the poverty subsample in columns (3) and (5), labeled "Census" (for Census sample).

For the high wealth group we find no evidence of excess sensitivity in either sample. For the low wealth sample the coefficient on lagged income is significantly different from zero in the larger sample in column (4) but not in column (5). This absence of significant evidence of excess sensitivity does not stem from the difference between Zeldes' and our specification, but rather from differences in the sample periods. As Mariger and Shaw (1988) pointed out, the income-consumption correlation in the PSID varies substantially from year to year. In particular, consumption growth is characterized by excess sensitivity in the 1970s; inclusion of the 1980s tends to reverse the results. If we restrict the sample to Zeldes' sample period (up to 1982), we find results qualitatively similar to his: a coefficient of -0.019 with a standard error of 0.005 for the low wealth group and an insignificant coefficient equal to 0.003 for the high asset sample.

Table 6 reports our main results, the coefficients of the switching regression model described in Section II. We use three different sets of instruments. The smallest set includes the second stage regressors, a dummy for non-white and three dummies for Census regions only. Estimates are shown in column (1). The basic set in column (2) adds a dummy for male head of household, two dummies for marital status, five dummies for family size in year t, and five education dummies for the head. The large set in column (3) also includes job related variables for the head: employment status, union status, years worked since age 18, 11 industry dummies, and 9 occupation dummies. The coefficients on lagged income differ somewhat depending on what set of instruments is used. Presumably this is due to the imprecision of the estimates. Curiously, the effect of lagged income is more negative for the unconstrained group and even significant in some specifications. However, note that the use of instruments here is not to solve an endogeneity problem but purely to facilitate the link between the two samples. Therefore, adding additional instruments will not lead to biased estimates as in standard two-stage least squares (Bound et al. 1995, Angrist and Krueger 1995). Instead, adding instruments will help the precision as long as these instruments are validly excluded from the Euler equation. This is evident in the standard errors, especially for the constrained regime. Therefore, the estimates in column (3) should be most reliable. They show little evidence that liquidity constraints affect consumption allocations. As in Table 5, results excluding the poverty subsample in column (4) are not substantially different.

Notice that all the coefficients in the lower panel of Table 6 for the constrained group are estimated much less precisely than those for the unconstrained regime. The standard errors are typically three to four times as large. The reason for the relatively poor results in the constrained regime is due to the fact that the imputed probability of a constraint tends to be small. Figure 2 displays a histogram of the imputed probabilities using the basic set of instruments, including the poverty subsample, and weighting the data. The predicted probabilities are clustered between 0 and 0.4, with very few observations having a probability above 0.5. This indicates there is relatively precise information on which households are very likely to be unconstrained but there is much uncertainty on which particular households will face binding constraints. This is reflected in the larger standard errors for the constrained regime.

Notice that we do not condition on other variables like labor supply which might be non-separable from food consumption. We do not do this because we would (ideally) need instruments for the conditioning variables which are available in the SCF. Since the natural instruments for such variables are their lags, this is not feasible. Zeldes (1989) argued that such non-separabilities, as long as they affect both the constrained and unconstrained consumers alike, should show up in the coefficients for the unconstrained regime as well. The differences in the coefficients between the two regimes should still give an idea of the impact of constraints. We find little evidence for such differences between the two regimes.

V. Euler Equation Estimates in the CEX.

We ran similar regressions using the Consumer Expenditure Survey in order to be able to use broader measures of nondurable consumption. The instrument set differs slightly from the one used above for the PSID. We do not have quarterly changes in family size in the SCF. We therefore include the following seven additional dummy variables to help better predict these changes. The first three such variables are interactions of married, wife's age 25-34, and dummies for no kid, one kid, and two kids. These latter variables are highly correlated with births and therefore increases in the number of kids. An interaction of a dummy for one kid and that kid's age being 1 to 3 serves a similar purpose, so is the fourth variable added. The final three such variables are interactions of married, wife's age 43-52, and dummies for one kid, two kids, and three or more kids. These latter variables capture when kids are likely to leave the household or become adult household members (which happens mechanically when the kid turns 18).

In addition to lagged income, all regressions include again a constant, age, age squared, the change in the number of adults and the change in the number of kids. We also include eleven month of the year dummies to control for seasonality. While the expenditure amounts refer to quarters, households are interviewed every month, so that an interview quarter can end in any month. Table 7 reports the coefficients on lagged income, the other coefficients are not displayed. In addition to estimates for the full sample and the two sample switching regression results, we report two sample splits. One is based on the wealth to income ratio where wealth only includes checking and savings accounts. Because this results in much lower assets we define low wealth households as those with wealth less than 1/12 of income, rather than 2/12 as in the PSID. The second split is based on income; following Garcia, Lusardi, and Ng (1995), we split the sample at a real income of $16,000.

The results are easily summarized. Only the wealth to income split reveals evidence for excess sensitivity for the constrained group, and even then only for food consumption. Generally, the coefficients on lagged income tend to be closer to zero for strictly nondurables than for food and closer for durables than for strictly nondurables. In the switching regression results we tend to find positive but insignificant results for the constrained group and slightly negative but significant results for the unconstrained. This seems to mimic some of the results in Table 6 above.

Table 8 reports additional switching regression results restricting the sample further. In the top panel we exclude farmers and self-employed, and we restrict changes in the log of consumption to be less than 2 in absolute value for all three consumption concepts. The next two panels use the same sample but make additional changes. The second panel also excludes the years 1980 and 1981. These were the first years of data collection for the CEX and the data quality tends to be poorer in these years. The last panel does not use income as an instrument in the first stage. The rationale for this is the paucity of the income variable in the CEX. Since this variable seems badly mismeasured this will affect the predicted $L\hat{}_{it}$. Since the latter is so important for our exercise we want to avoid this possible contamination. On the other hand, this makes the identification of the second stage model somewhat questionable. These estimates should therefore be regarded as only indicative. The results in Table 8 show that none of these changes affects the conclusions in any major way. None of the specifications seems to indicate any role for liquidity constraints.

VI. The Conditional Distribution of Consumption Growth.

The Euler equation relates the conditional mean of consumption growth to preferences and, possibly, to households' resources if credit markets are not perfect. But liquidity constraints have implications about the entire conditional distribution of consumption growth (Deaton, 1991). This is most easily seen if income follows a symmetric i.i.d. process. Constrained consumers are perfectly able to smooth above average income shocks; however, they cannot smooth income shocks that are below average. Thus liquidity constraints reduce consumers' ability to smooth income fluctuations in a predictable way: the distribution of consumption levels would be negatively skewed, and the distribution of consumption changes would have fatter tails than under perfect capital markets. The same result holds if a stationary income process is autocorrelated, albeit less strongly. Things are more complicated if income is non-stationary. For example, liquidity constraints have no impact on consumption if income is a random walk without drift because consumers never want to borrow anyway. In practice, households' income in the United States is less persistent than a random walk (Pischke, 1995); and for many occupations income grows over a good part of the life-cycle, thus leaving ample room for liquidity constraints to affect the conditional distribution of consumption.

Quantile regressions represent a simple and convenient way to characterize the conditional distribution of a variable. In this section we estimate the conditional distribution of consumption growth, controlling for preference shifts and lagged income. Using the framework of the switching regression model outlined in Section II, we argue that liquidity constraints lead to two testable implications in quantile regressions. First, constrained consumers, who move from periods in which the constraint binds to periods in which it does not, have more variable consumption growth than unconstrained consumers. Thus, the distribution of consumption growth should spread out with $\hat{L}_{it}$, the imputed probability that household i is subject to a liquidity constraint in period t. This implies that the coefficient of $\hat{L}_{it}$ in the switching regression should be higher at higher quantiles. The second testable implication derives from the fact that in an economy with liquidity constraints a consumer is more likely to face a binding constraint in period t if income in period t is relatively low. Since there is a good chance that the constraint will be relaxed in period t+1 if income is mean reverting, consumption growth between t and t+1 will be more variable the lower period t income. This implies that the coefficient of lagged income should be more negative at higher regression quantiles, i.e. that one should find more excess sensitivity at higher quantiles.

Since a closed form solution for consumption in the presence of income uncertainty and liquidity constraints does not exist, it is difficult to give a formal proof of the above conjectures. In order to investigate the effect of liquidity constraints on the conditional distribution of consumption growth we did some simple simulations. We estimated consumption Euler equations for the 25th, 50th, and 75th quantiles in the artificial samples. The experiment is designed to reveal systematic differences across various consumption growth quantiles, conditional on lagged income. We also provide OLS estimates as a comparison and as a general check for the methodology.

For simplicity, we posit a discrete process where income can take only three values: low, middle, and high. Base income in each of these states is set to 2, 4, and 5, respectively. The following transition matrix describes the probabilities of income changes:

		period	t+1	
		low	medium	high
period t	low	0.5	0.4	0.1
	medium	0.2	0.7	0.1
	high	0.1	0.1	0.8

The three states are supposed to resemble unemployment, a relatively insecure medium income state (such as blue collar work), and a relatively secure high income state. The initial distribution of income in our simulation samples is given by the stationary probabilities of the transition matrix above. The resulting income process is stationary as long as we do not impose growth. To allow for growth, we add a deterministic growth component of one to the stochastic part obtained using the transition matrix. A three period consumption model is then simulated, where instantaneous utility is quadratic and the interest and discount rates are set to zero; thus the only non-linearity in the policy function stems from liquidity constraints. We chose the number of potentially constrained households in the simulations to resemble the empirical probabilities of being liquidity constrained in the PSID.[14] Denoting the empirical probabilities by $L^\wedge_{it}$ again, each household is assumed to be potentially constrained in the simulations if $L^\wedge_{it}$ exceeds a random number generated from a uniform (0,1) distribution. The constraint is binding only if predicted income is higher than current income; thus households in the high income state never face binding constraints because their expected income is lower than current income.

After obtaining income and consumption for every period, we computed two consumption changes and pooled the data. Since discrete data with few points of support are problematic to use in quantile regressions, and since PSID data are surely noisy, we add a normal (0,1) random variable to each consumption and income change.

We ran three simulations. The first does not impose liquidity constraints on any household; however, even in this case we introduce the probabilities $L^\wedge_{it}$ in the regressions as a basic check of the methodology. The second simulation introduces liquidity constraints. The interaction between the inability to borrow and income growth makes the constraint binding in most cases for the households selected as potentially constrained (i.e., those for which $L^\wedge_{it}$ exceeds the random number). In order to allow for cases in which households move between constrained and unconstrained states more often, in a third simulation we eliminate income growth. Every simulation is run 100 times and includes two income changes for 4,175 households, the number of households in our PSID sample, resulting in a total of 8,350 observations.

Table 9 shows the results on applying quantile regressions to the simulated data. The Table reports four coefficients for each of the three simulations. The first coefficient

---

[14]These are obtained from a probit model in this case to avoid negative predicted probabilities.

is that of lagged income in a simple regression of consumption changes on a constant and lagged income. The other three coefficients come from a separate regression where $L\hat{}_{it}$ and an interaction term between $L\hat{}_{it}$ and lagged income are introduced as additional regressors, corresponding to the switching regression model (each regression also includes a constant term). Notice that the coefficient on lagged income corresponds to θ_1 in (2) while the coefficient on the interaction with $L\hat{}_{it}$ corresponds to the difference (θ_2 - θ_1). This differs from the way we reported these coefficients in Tables 6 and 7 above. Panel A reports the results where no households in the simulations are actually constrained. The results are not surprising, but represent a useful check on our assumptions and methodology. In the absence of liquidity constraints all coefficients of the OLS and quantile regressions are close to zero and not statistically significant. Panel B reports the results of the simulation obtained imposing liquidity constraints. The OLS coefficients indicate that the switching regression model works well. In the standard model with no interactions the coefficient of lagged income is negative but small (-0.045 with a t-statistic of -4.5) as is expected when some consumers are constrained. The switching regression model indicates that excess sensitivity is solely due to the presence of liquidity constraints. On average, consumption changes for the constrained group are higher, as witnessed by the coefficient of $L\hat{}_{it}$. Furthermore, the correlation between consumption changes and lagged income is attributed entirely to the constrained group: the coefficient on lagged income is zero while the coefficient on the interaction with $L\hat{}_{it}$ is -0.227 with a t-statistic of -3.5. Separating the sample therefore increases the (absolute value) of the coefficient of lagged income in the constrained group. On the other hand, the precision of the estimate is reduced because the sample separation rule is stochastic.

The quantile regression results in Panel B also conform to the intuition given above. In the standard specification with only lagged income, the distribution of consumption changes becomes tighter for higher values of lagged income, as expected. Most of the tightening occurs in the upper part of the distribution of consumption changes: even the median regression does not show evidence in favor of excess sensitivity. The results for the switching regression model are even more striking. There is a large, positive effect of liquidity constraints on consumption changes at all quantiles, as is evident from the coefficient on $L^{\wedge}_{it}$. More importantly, the distribution of consumption changes spreads out with the probability of being constrained. Furthermore, the coefficient of the interaction between lagged income and $L^{\wedge}_{it}$ is negative over the entire distribution, and decreases across consumption quantiles.

Panel C eliminates income growth from the simulation. In this case households are more likely to switch in and out of the constrained status as their income varies. The results do not fully conform with the intuition developed above. In the absence of income growth, average consumption changes are not clearly larger the higher the probability of a binding constraint. But the distribution of consumption changes still spreads out with $L^{\wedge}_{it}$. The coefficient of lagged income interacted with $L^{\wedge}_{it}$ (the constrained regime) is negative as in Panel B, but its size in absolute value is much smaller. The results for the unconstrained group are puzzling, however. The coefficient of income is negative, and in the median regression excess sensitivity is of the same magnitude as for constrained consumers. The results for the standard specification without interactions (the first row) is also puzzling: the coefficient of lagged income does not increase, in absolute value, across consumption quantiles.[15]

The simulations indicate that even a simple certainty-equivalence life cycle model with liquidity constraints delivers interesting predictions about the entire conditional distribution of consumption changes. Overall, these predictions conform well with the intuition developed at the beginning of the section. On the other hand, the simulations indicate that the exact distributional patterns depend on the income process chosen. We also conjecture that the predictions will be even more fragile when there are non-linearities in the policy function beyond those introduced by liquidity constraints, for example, in the case of CRRA preferences. We have not explored this issue yet, and regard it as a topic for future research.

---

[15]In these regressions there is also a stronger effect of lagged income at the mean and at all quantiles than in the simulations with income growth reported in Panel B.

Table 10 displays the results obtained from quantile regressions in the PSID and CEX. These regressions differ from the results presented in Table 6 for the PSID in that we use one year changes in log consumption here, not three year changes. The CEX results correspond directly to Table 7. Recall that, unlike in the earlier tables, the reported coefficients on the interactions with the probability of being constrained denote differences between the constrained and the unconstrained regime. For comparison with Table 9, we organize Table 10 in the same fashion.[16] As in the previous sections, the OLS coefficients do not uncover excess sensitivity. The quantile regressions exhibit a pattern similar to the simulations of Table 9 with liquidity constraints. In particular, the coefficient of lagged income falls at higher consumption change quantiles and the coefficient on $L\hat{}_{it}$ increases at higher quantiles. Nevertheless, there are also differences in the patterns between the different samples. The results for food consumption in the PSID resemble more closely those for strictly nondurable consumption rather than those for food in the CEX. The results for food in the CEX in panel B exhibit only small differences between the coefficients on lagged income at various quantiles. In all estimation results the coefficient of $L\hat{}_{it}$ is large and negative at the 25th percentile, small at the median and large and positive at the 75th percentile; this conforms most to the no income growth case of panel C in Table 9.

Overall, we interpret these results as weak evidence that consumption allocations are affected by liquidity constraints. As in the estimation of Euler equations, the predicted probability of a constraint may also proxy for households with more variable income and with less ability to buffer income fluctuations. Therefore, precautionary savings may be an alternative interpretation for these results.

---

[16]In the simulations reported in Table 7 the dependent variable is consumption changes, while it is consumption growth in Table 8.

VII. Conclusion.

Relying on sample separation rules based on household wealth to classify households as liquidity constrained and unconstrained is a technique often adopted in the applied consumption literature. Sample separation rules based on wealth pose several problems, however. If the resulting low-wealth subsample includes many unconstrained households, the excess sensitivity coefficient is biased towards zero. We show that for the sample split used by Zeldes (1989) as many as 75 percent of the households in the low-wealth group may in fact have access to consumer credit or mortgage loans. This finding alone is not necessarily damaging to Zeldes' conclusions, because he still finds significant evidence of liquidity constraints in a sample contaminated by unconstrained households. In this paper we check the robustness of the asset-based sample splitting approach using an alternative method.

The alternative that we propose is to identify liquidity constrained households using information on individuals who have been denied loans or discouraged from borrowing. Such information is available in the Survey of Consumer Finances, and allows us to relate the probability of being liquidity constrained to a set of demographic variables. Using a first stage model estimated on the SCF, we can impute the constraint probabilities in second samples, the Panel Study of Income Dynamics and the Consumer Expenditure Survey, which contain information on consumption. We can then estimate Euler equations for constrained and unconstrained households using a switching regression model with imperfect sample separation. Ultimately, the estimation procedure relies on detecting a correlation between consumption growth and the demographic variables used in the first-stage estimation. The method of combining information from two samples is appealing because the instruments (which have to be available in both data sets) must not affect the Euler equation directly. Time-invariant demographic variables are very well suited as instruments because they are certainly part of households' information set. Problems only arise if the instruments are correlated with omitted variables in the Euler equation error, for instance with proxies for the variance of future consumption growth or with individual rate of time preferences.

The point estimate of lagged disposable income in the Euler equation for the constrained regime is close to zero, suggesting that liquidity constraints do not affect the mean of the distribution of consumption growth. On the other hand, our estimates are generally less precise than those based on exogenous sample separation rules. This is not surprising. While Zeldes' asset-based split classifies 62 percent of the sample as constrained, the SCF suggests that only 16 percent of the households do not have access to credit markets and are therefore liquidity constrained. Beyond the smaller effective sample of constrained households, the switching regression explicitly recognizes that one cannot perfectly classify households as constrained or unconstrained. This uncertainty about regime classification is reflected in the standard errors. While the same source of uncertainty is also present in Zeldes' procedure, his standard errors are computed as if one had perfect knowledge on who is constrained and who is unconstrained.

We also present similar results with data from the Consumer Expenditure Survey. While the CEX has better and broader measures of consumer spending than the PSID, the income variables are of poorer quality. This may make it hard to find significant evidence of excess sensitivity using lagged income as a regressor. On the other hand, our results for the CEX are not very different from what we found for the PSID.

Testing for excess sensitivity alone gives a rather incomplete picture of the implications of liquidity constraints. We therefore explore the effect of liquidity constraints on the entire conditional distribution of consumption growth; in particular, we study the relation between the quantiles of the distribution of consumption growth, the probability that an individual is liquidity constrained and the lagged income coefficient in the Euler equation. Simulations indicate that the conditional distribution is wider the higher the probability of a binding liquidity constraint and that the coefficient of lagged income increases (in absolute value) at higher quantiles. This is exactly what one observes in the data if one compares the coefficient estimates of the Euler equation for various quantile regressions. The pattern of results for the conditional distribution of consumption growth is hard to reconcile with alternative explanations; we thus interpret these findings as weak evidence that liquidity constraints affect the intertemporal allocation of consumption.

The reason we do not find excess sensitivity at the mean of the consumption growth distribution remains unclear. One possibility is that the effects of liquidity constraints on consumption are very small, and our estimates are too imprecise to detect them. Future research will have to clarify whether this is a sensible implication for a model with liquidity constraints. In any case, we conclude that imputing the liquidity constraint probabilities from the Survey of Consumer Finances is a practical and useful alternative to sample separation rules based on assets.

VIII. Appendix: the Data.

1. PSID.

As in Zeldes (1989) the value of a PSID variable in "PSID year" t means the value reported in the survey of year t (even if the variable refers to the previous calendar year). Topcoded variables are generally left at their topcodes.

1.1 Constructed Variables.

Variables constructed following Zeldes (1989). We generally followed Zeldes' detailed appendix, sometimes elaborated upon in personal communication, in constructing the variables that appear in his analysis (i.e., food consumption, disposable income, non-housing wealth, and housing equity). Here we discuss only certain key constructions that either are not fully clear from Zeldes' appendix alone, or that go beyond Zeldes' sample period since we extend the analysis to 1987.

Food consumption. In 1968-1972 and in 1974 (when food stamps are not already included in expenditures on food), total real food expenditure in PSID year t includes the real value of food stamps received in the previous calendar year though reported in t. This latter real value is deflated by the average of the previous year's monthly CPI price indices for food at home.

Disposable income. The property tax for 1978 is estimated as the product of the house value in 1978 times the implicit 1977 tax rate (i.e. the ratio of property taxes paid in 1977 to the house value in 1977), if these are available and if the family did not move between April of 1977 and April of 1978. Else the 1979 tax rate is used, if available and if the family did not move between April of 1978 and April of 1979. If neither of these estimates is available, but the family does not own its house in 1978, property taxes in 1978 are set to zero. In estimating social security taxes, the wife's wages were set to zero if missing, to avoid losing observations. The Euler equation for $\ln c_{t+3} - \ln c_t$ contains (the log of) the real disposable income whose components are reported in PSID year t (and so refer to calendar year t-1).

Liquid assets (non-housing wealth). From 1976 on, the asset income that is to be blown-up to estimate liquid assets is the sum of the head's variable for "dividends, interest, rent, trust funds, and royalties", and of the similar variable for the wife (which includes alimony). (Before 1984 it was not possible to distinguish rental income from dividend, interest, etc. income. To preserve comparability with these years, from 1984 on asset income is the sum of the variables for rent and for dividend, interest, etc. but not rent.) Before 1976 these variables are bracketed, so we instead blow-up "total asset income", defined as total taxable income of head and wife, minus the sum of the head's total labor income and the wife's annual wages. (From 1977 when it is first available, the alimony received by the head is also subtracted in forming total asset income, since it is included in total taxable income but does not represent asset income.) If the relevant measure of asset income is negative, liquid assets are set to missing. Otherwise, the first $250 of asset income in PSID year t (so in calendar year t-1) is divided by the annual average passbook rate in calendar year t-1. The rest of such income is divided by the annual average yield of 3-month Treasury bills in t-1.

Since the blowing-up procedure implicitly assumes that most asset income is interest-like income, liquid assets are set to missing if the value of other types of asset income is too large. Before 1976 this is the case if total asset income is less than the sum of the lower brackets of the head's and wife's dividend, interest, rent etc. variables; or if greater than the sum of the upper brackets of these variables. (The upper bracket of the category "over $10,000" is taken as $9,999,999, the largest possible value for total asset income.) Since 1976 liquid assets are missing if the absolute value of "home business income" is greater than $100. Home business income is total asset income minus dividends, interest, rent, etc. If any of these variables for testing whether other asset income is too large is missing, liquid assets are set to missing. Liquid assets are also set to missing if others' asset income is positive or missing. (For years before 1975, they are missing if others' asset income is positive or missing in 1975.)

To maintain compatibility with the SCF, which does not flag imputations, asset income is blown-up even if some of the income variables have undergone major imputations. (Observations with major imputations are still excluded from the Euler equations.)

Net wealth (including housing equity). The outstanding mortgage principal is not reported in years 1973-5 and 1982. The values in 1973-5 are interpolated from the values in 1972 and 1976 when available, and only if the family did not move and did not hold a second mortgage between April 1972 and April 1976 and if the reported principal is not larger in 1976 than in 1972. Else the values are extrapolated from the values in 1976 and 1977 if available, and if the family did not move nor hold a second mortgage between April 1977 and April of the year at hand, and if the reported principal is not larger in 1977 than in 1976. Else the values are similarly extrapolated from the values in 1971 and 1972, if possible. If the reported principal is zero in both 1972 and 1976, it is set to zero for 1973-5; and if the estimated principal is negative, it is set to zero. The principal in 1982 is estimated similarly. Net housing equity in PSID year t reflects the contemporaneous value of equity, i.e. in calendar year t. Therefore it is lagged before being added to liquid assets, which are blown-up from asset income in the previous calendar year, t-1. The real values of liquid assets and net wealth are the nominal values just described, deflated by the annual average of the NIPA personal consumption expenditures deflators for the previous calendar year (t-1).

1984 Stocks of Wealth. To gauge the quality of the blowing up procedure, we compare the value of assets resulting from inflating asset income in 1984 (based on asset income in calendar year 1983) to the actual stock of liquid assets in 1984. Ideally the stock should correspond to the flows that went into the procedure, namely dividends, interest, rent, etc., but such individual stocks are not separately available. The available relevant stocks are: first, ''cash on hand'', the value together of checking and savings accounts, money market funds, CD's, government savings bonds, and Treasury bills, including amounts in IRA's; second, the value of stocks, mutual funds, and investment trusts, including amounts in IRA's; third, the value of bonds, trusts/estates, life insurance, collectibles, etc.. Putting aside rental income, the ideal stock corresponding to the flows should be bounded below by cash on hand, our narrow measure of liquid assets, and above by our wide measure, the sum of all three categories.

The narrow and wide measures of net wealth add housing equity (net of the outstanding mortgage principal, as above) to the narrow and wide measures of liquid assets. The real values of these stocks are the nominal values just described deflated by the average of the monthly PCE deflators for the first quarter of 1984.

Variables used in the First Stage Estimates of Being Constrained. We will describe only the variables that are not self-explanatory. Generally, the variables are constructed to match the variables used in estimating the probability of being constrained in the SCF. All characteristics refer to the household head. In couples, the head is defined as male.

Married/Divorced. The head of the family is ''married'' if formally married or permanently cohabitating with someone, and this partner is present in the family. The head is ''divorced'' if formally divorced or separated, and the partner is not present in the family.

Region. Unlike the SCF, the PSID has some families living outside the continental United States. Families in Alaska and Hawaii are classified as living in the West, those abroad have region set to missing.

Employed. The head is employed if he or she is working part time or full time at the time of the interview.

Number of years head has worked full time since age 18. This question has been asked since 1974. According to the documentation, it is generally asked only when there has been a change in head, though it was asked of all families in 1974, 1976, and 1985. (It also appears to have been updated for most families in 1975.) From 1974 on, values for missing years between successive reports m and n in years s and t are extrapolated as follows: If the head reports in PSID year s+1 more than 1500 hours worked (in calendar year s), the number of years in s+1 is set to m+1; if not, the value stays at m. (If any of this information is missing, the count if not incremented.) Similarly for the remaining years until t, the count is incremented if the head worked more than 1500 hours. (The count is not incremented beyond 98 years, because 99 is the value for ''not applicable'' or ''don't know''.)

From 1973 backwards, for year s the value (if not 99) at s+1 is decremented if the year s+1 PSID reports more than 1500 hours. (Though once the count hits zero, it stays at zero.) If the number of hours worked is missing, the count does not change. Since they seemed unlikely to be accurate, values greater than 80 years were set to missing (as were the 99's).

Occupation. The PSID reports a one-digit occupation code in 1968-1975, two-digit in 1976-1980, and three-digit from 1981. For consistency the latter two (and the three-digit code in the SCF) are transformed into the first. The organization of the two-digit codes generally points to the appropriate one-digit code. The only exceptions are that military personnel (code 52) and public sector protective workers (code 55) are set to missing, because the corresponding one-digit code does not distinguish them from the unemployed and the ''not-applicables''. The not-applicables and ''don't-knows'' (99) are also set to missing. The table in the documentation setting out the two-digit codes gives the corresponding three-digit codes. Again military personnel (code 600) are set to missing, as are protective workers (codes 960-65) who work for the government. The ''inapplicables'' (i.e., the unemployed/not working, but not those temporarily laid-off) are kept together in a separate category, no matter what their previous occupation.

Industry. The two-digit codes reported to 1981 are naturally grouped into 12 categories. The organization of the three-digit codes reported since 1981 points to the corresponding two-digit codes, and thereby to the 12 categories. The exceptions are again the military and ''don't knows/not applicables'', set to missing; and printing, which following the BLS categorization is grouped with manufacturing. Again, the ''inapplicables'' are grouped together.

1.2 Sample Selection.

We generally followed the procedures of Zeldes, extended to 1987 (wave 20). Therefore we limit discussion to just a few salient notes. Even if a family is non respondent in 1987, previous observations of it are included provided they meet all other requirements. A family-observation in year t is excluded if anyone other than the ''wife'' and their children lives with the head.

1.3 Splitting the Sample.

Splits based on blown-up wealth. The three Zeldes splits based on blown-up wealth are constructed following Zeldes's appendix. Again discussion here is limited to a few salient notes. The various wealth to (average) disposable-income ratios used in splits 1-3 are set to missing if either the numerator or denominator is negative, or if the denominator is zero. For splits 1 and 2 (based on liquid assets), the direct questions about whether the household has any savings and whether they amount to two month's income are available in 1971-2, 1975, and 1979-80. In these years, an observation is not assigned to either the constrained or unconstrained groups if the blown-up wealth-to-income ratio is missing, whatever the answers to the direct questions.

There is a typographical error in Zeldes's description of split 2, the stringent split. In years in which the direct questions are reported, an observation is placed in the unconstrained group only if they have at least two month's savings (not just any savings, as printed in the appendix). In years in which the blown-up ratio is also generally computed, i.e. since 1971, that ratio must exist and be greater than 1/2.

Splits based on Actual Asset Stocks. As above, the wealth to (average) disposable-income ratios are set to missing if either the numerator or denominator is negative, or if the denominator is zero. The reported splits use the narrow measure of liquid assets (cash on hand) in the numerator (plus net housing equity, for split 3), which most closely corresponds to the asset income variables. For comparability with the Zeldes splits using blown-up wealth, the stocks in the numerators here are lagged, since unlike asset income the stock values in PSID year t refer to calendar year t. However, unlike the Zeldes splits the direct questions on savings are not used. Accordingly, under split 1 an observation is unconstrained simply if the ratio of (the lag of) narrow liquid assets to average disposable income is greater than (or equal to) 2/12; and constrained if this ratio is less than 2/12. Similarly, under the extreme split 2 an observation is unconstrained if this ratio is greater than 1/2 and constrained if the ratio is zero. Split 3 is like split 1, but with (the lag of) narrow net wealth in the numerator.

2. CEX.

Details can be found in the Appendix to Chapter 2.

3. SCF.

3.1 Constructed Variables.

The variables in the SCF generally are constructed to match the corresponding variables in the PSID, discussed above.

Disposable Income and Blown-up Wealth. Disposable income is total household income (variable b3202) minus nontaxable-interest income (b3207), minus income from the sale of stock, bonds or real estate (b3210), and minus household federal income and social security taxes. We subtract b3202 and b3207 because these seem unlikely to be well accounted for in the components of income that comprise taxable income in the PSID. Unlike the PSID, the SCF does not report property taxes. We recomputed the corresponding PSID variables without property taxes (and without averaging disposable income and its lag, which is not possible in the cross-sectional SCF), which yielded very similar results. These are not reported.

Taxes. Taxable income is adjusted gross income (b3219) minus 1000 times an estimate of the number of exemptions. Taxes are estimated using the 1982 tax tables, similarly to the estimation of the marginal tax in the PSID. The household's tax status is predicted to determine which schedule to use. The earned income tax credit is subtracted from taxes for eligible households. Social Security taxes are computed as for the PSID, using b4546 and b4646 for the head's and spouse's incomes. Even though here we know whether the spouse is self-employed, to maintain comparability with the PSID only the non-self-employed tax rate is used for the spouse.

Blown-up wealth. To match the variables on dividends, interest, rent, trust funds, royalties, and alimony in the PSID, the asset income used here is the sum of dividend income (b3208), taxable interest income (b3209), and income from rent, trusts, and royalties (b3211). Nontaxable interest income, like income from IRA's, was not included, since it seemed unlikely to have been included by many families in answering the PSID question at issue. Alimony was not included because it is not reported separately in the

SCF from inheritance, gifts, and child and other financial support. If positive, this asset income is blown-up as above, using the 1982 passbook and Treasury bill rates, to yield liquid assets. As before, blown-up liquid assets are set to missing if the absolute value of home business income is greater than 100. Here home business income is the sum of business income (b3206), income from asset sales (b3210), and other income (b3216). Net housing equity is the value of a home minus any outstanding principal on a first mortgage (b3709). If there is not a second mortgage (b4024=0), this value is added to the estimated liquid assets to produce blown-up net wealth. Unlike the SCF, it is not possible to lag housing equity before adding. The real values of disposable income and liquid wealth and net wealth are obtained by dividing the nominal values by the average of the 1982 PCE deflators.

Stocks of Wealth. To try to bound the ideal stock that corresponds to the asset income flow just described, we again compute both narrow and wide stocks. Our narrow measure of liquid assets is what the SCF calls ''liquid assets'' (b3301, the sum of checking, savings, and money market accounts, plus IRA's, Keoghs, CD's and savings bonds) minus the value of IRA and Keogh accounts (b3446). Our wide measure adds to this the value of bonds (b3458) and of stocks and mutual funds (b3462). Net wealth adds net housing equity, described above, to these measures of liquid assets. The real values of these stocks are the nominal values divided by the average of the 1983 PCE deflators.

Variables used in the First Stage Estimates of Being Constrained. These generally are constructed to match those from the PSID, so we will make just a few comments here.

Employment status. The head is taken as employed if working full- or part-time (b4511= 1 or 2).

Number of years working full-time. The number of full-time years of work since age 18 is computed as the total number of full-time years of work (b4516) minus an estimate of the number of such years before 18. If this difference is negative, or if b4516 is zero, the difference is set to zero. The number of years before 18 is the difference of 18 and the age of the first full-time job (b4515); if this difference is negative, the number of years before 18 is set to zero.

Occupation and Industry. The three-digit SCF occupation code (b4531) is transformed to the one-digit PSID code discussed above. Again, military personnel (b4531= 600, 996-8) are set to missing, as are protective workers (b4531= 960-5) who work for the government (b4540= 1-3). The transformation of the three-digit industry code (b4536) to the 11 industry dummies follows that of the three-digit codes in the PSID. These variables refer to the current occupation and industry; as in the PSID the unemployed/not working (but not the temporarily laid-off) are grouped together.

3.2 Sample Selection.

We used the version of the 1983 SCF data released with the 1986 data, which adopted a different definition of the head than in the original release of the 1983 data. (Basically, if a household was reinterviewed in 1986, the 1983 variables about the ''head'' are about the 1986 respondent, rather than the 1983 head of household.) We used the variable c1004 to restore the original concept of head, and adjusted the variables for head and spouse accordingly. We dropped the 41 repeated 1983 observations (b3075 = 4) arising from the 41 1983-couples who split and were both reinterviewed in 1986. The 159 ''uncleaned'' observations (b3001 = 3) whose values are largely missing, as well as the high-income sample (b3001 = 1), were also dropped. Furthermore, we only kept households who were present in the 1986 reinterview sample with no marital status change between 1983 and 1986 (c1202 = 1 or 4). Unfortunately we cannot exactly match the sample exclusions of the PSID. There is no indication of which variables have been imputed. Also, we don't know anything about their food consumption. (We did not exclude households with unrelated people in them, since the motivation for this exclusion was to improve the measure of food consumption.)

3.3 Splitting the Sample.

Splits based on blown-up wealth. Under split 1 a household is unconstrained if the ratio of blown-up liquid assets to disposable income is greater than 2/12, and constrained if the ratio is less than this. Under the extreme split 2, it is constrained if the ratio if greater than 1/2, unconstrained if zero. Split 3 is similar to split 1, but uses blown-up net wealth instead. (For all three splits, if either the numerator of the corresponding ratio is negative, or the denominator non-positive, the household does not go into either category.)

Split based on the Direct Measure of Liquidity Constraints. A household is constrained if either a) it has been turned down for a loan, or failed to get as much credit as it desired (b5522 = 1 or 3); and if it reapplied for a loan, it did not then receive as much credit as desired (b5525 ~= 1); or b) it has been dissuaded from applying for credit, because it expected to be turned down (b5526= 1).

4. Standard Errors.

The method for computing the standard errors is set out in the working version of this paper.

Bibliography

Altonji, J. G. and A. Siow (1987), "Testing the Response of Consumption to Income with (Noisy) Panel Data," Quarterly Journal of Economics 102, 293-328.

Angrist, J. and A. B. Krueger (1992), "The Effect of Age at School Entry on Educational Attainment: an Application of Instrumental Variables with Moments from Two Samples," Journal of the American Statistical Association 87, 328-336.

Angrist, J. and A. B. Krueger (1995), "Split-Sample Instrumental Variables Estimates of the Return to Schooling," Journal of Business and Economic Statistics 13, 225-235.

Arellano, M. and C. Meghir (1992), "Female Labour Supply and on-the-Job Search: An Empirical Model Estimated Using Complementary Data Sources," Review of Economic Studies 59, 537-557.

Attanasio, O. P. (1994), "The Intertemporal Allocation of Consumption: Theory and Evidence," NBER Working Paper No. 4811.

Attanasio, O. P. and G. Weber (1993), "Consumption Growth, the Interest Rate and Aggregation," Review of Economic Studies 60, 631-649.

Avery, R. B. and A. B. Kennickell (1988), 1983 Survey of Consumer Finances: Technical Manual and Code-book. Washington: Federal Reserve Board.

Boyes, W. J., D. L. Hoffman, and S. A. Low (1989), "An Econometric Analysis of the Bank Credit Scoring Problem," Journal of Econometrics 40, 3-14.

Browning, M. (1991), "A Simple Non-additive Preference Structure of Household Behavior over Time" Journal of Political Economy 99, 607-637.

Campbell, J. Y. and G. N. Mankiw (1989), "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," 1989 Macroeconomic Annual, S. Fischer and O. Blanchard (eds.). Cambridge: The MIT Press.

Carroll, C. D. (1992), "Buffer Stock Savings and Life Cycle/Permanent Income Hypothesis," mimeo., Federal Reserve Board of Governors.

Carroll, C. D. and D. N. Weil (1994), "Saving and Growth: a Reinterpretation," Carnegie-Rochester Conference Series on Public Policy 40, 133-192.

Cox, D. and T. Jappelli (1993), "The Effect of Borrowing Constraints on Consumer Liabilities," Journal of Money, Credit and Banking 25, 197-213.

Deaton, Angus (1991), "Savings and Liquidity Constraints," Econometrica 59, 1221-1248.

Duca, J. V. and S. S. Rosenthal (1991), "An Econometric Analysis of Borrowing Constraints and Household Debt," Canadian Real Estate Research Bureau, Working Paper 90-ULE-014.

Eberly, J. C. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," Journal of Political Economy 102, 403-436.

Flavin, M. (1985), "Excess Sensitivity of Consumption to Current Income: Liquidity Constraints or Myopia?" Canadian Journal of Economics 18, 117-136.

Gale, W.G. and J. K. Scholz (1994), "IRAs and Households Saving," American Economic Review 84, 1233-1260.

Gali, J. (1990), "Finite Horizons, Life-Cycle Savings and Time Series Evidence on Consumption," Journal of Monetary Economics 26, 433-452.

Garcia, R., A. Lusardi and S. Ng (1995), "Excess Sensitivity and Asymmetries in Consumption: an Empirical Investigation," mimeo, Universite de Montreal.

Goodfriend, M. (1992), "Information Aggregation Bias," American Economic Review 82, 508-519.

Gropp, R., J. K. Scholz, and M. White (1995), "Personal Bankruptcy and the Supply and Demand for Credit," mimeo, International Monetary Fund.

Hajivassiliou, V. A. and Y. M. Ioannides (1991), "Dynamic Switching Regression Models of the Euler Equation: Consumption and Liquidity Constraints," mimeo, Yale University.

Hall, R. E. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy 86, 971-987.

Hall, R. E. and F. S. Mishkin (1982), "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica 50, 461-481.

Hayashi, F. (1987), "Tests for Liquidity Constraints: a Critical Survey," in T. Bewley (ed.), Advances in Econometrics, Fifth World Congress, vol. 2. Cambridge: Cambridge University Press.

Hubbard, R. G., and K. L. Judd (1986), "Liquidity Constraints, Fiscal Policy, and Consumption," Brookings Papers on Economic Activity 1:1986, 1-59.

Jappelli, T. (1990), "Who is Credit Constrained in the U.S. Economy?" Quarterly Journal of Economics 104, 219-234.

Jappelli, T. and M. Pagano (1995), "The Welfare Effects of Liquidity Constraints," CEPR Discussion Paper no. 1108.

Juster, F. T. and R. P. Shay (1964), " Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation," NBER Occasional Paper no. 88.

Lee, L. F. and R. H. Porter (1984), "Switching Regression Models with Imperfect Sample Separation: With an Application to Cartel Stability," Econometrica 52, 391-418.

Lusardi, A. (1995), "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," Journal of Business and Economic Statistics 13, forthcoming.

Maddala, G. S. (1983), Limited Dependent and Qualitative Variables in Econometrics. Cambridge: Cambridge University Press.

Maddala, G. S. (1986), "Disequilibrium, Self-Selection, and Switching Models," in Z. Griliches and M. D. Intriligator (eds.), Handbook of Econometrics, vol. 3. Amsterdam: North Holland, 1633-1688.

Mariger, R. P. and K. Shaw (1988), "Reexamining the Sensitivity of Consumption to Transitory Income," mimeo, University of Washington.

Meghir C. and G. Weber (1993), "Intertemporal Non-Separability or Borrowing Restrictions? A Disaggregate Analysis Using the U.S. CEX Panel," mimeo., University College London.

Munnell, A. H., L. E. Browne, J. McEneaney, and G.M.B. Tootell (1992), "Mortgage Lending in Boston: Interpreting HMDA Data," Federal Reserve Bank of Boston Working Paper 92-7.

Perraudin, W. R. M. and B. E. Sorensen (1989), "Cheats, Banks and Liquidity Constraints: an Empirical Study of Demand and Supply in the Loan Market," mimeo, Harvard University.

Pischke, J. S. (1995), "Individual Income, Incomplete Information, and Aggregate Consumption," Econometrica 63, forthcoming.

Runkle, D. E. (1991), ''Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data,'' Journal of Monetary Economics 27, 73-98.

Wolff, E. N. (1994), ''Trends in Household Wealth in the United States, 1962-83 and 1983-89,'' Review of Income and Wealth 40, 143-174.

Zeldes, S. P. (1989), ''Consumption and Liquidity Constraints: an Empirical Investigation,'' Journal of Political Economy 97, 305-46.

Figure 1

Regression Function of Liquidity Constraint on Net Wealth

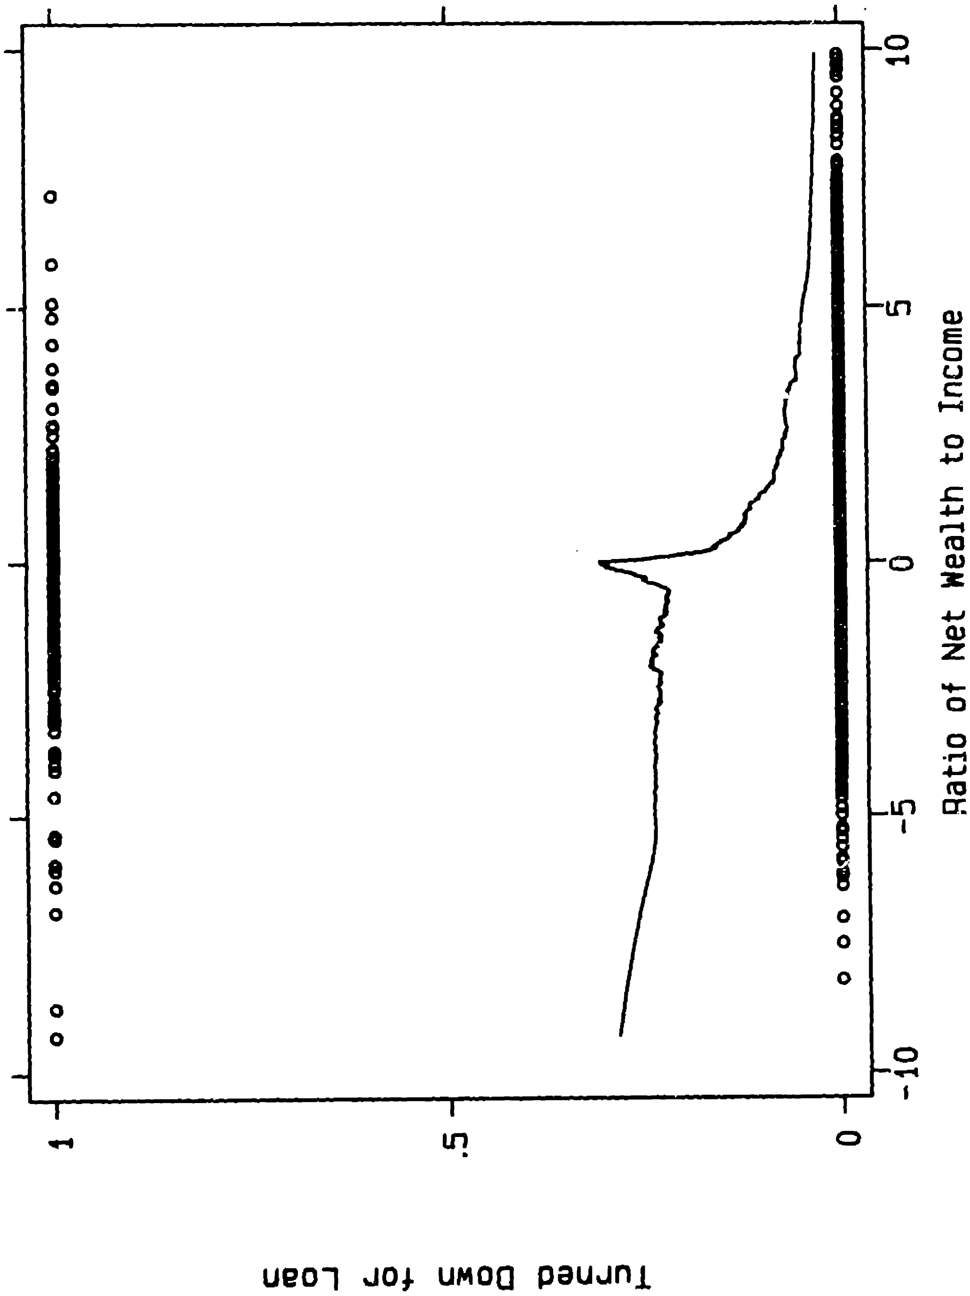

Figure 2

Histogram of Imputed Probability of Constraint

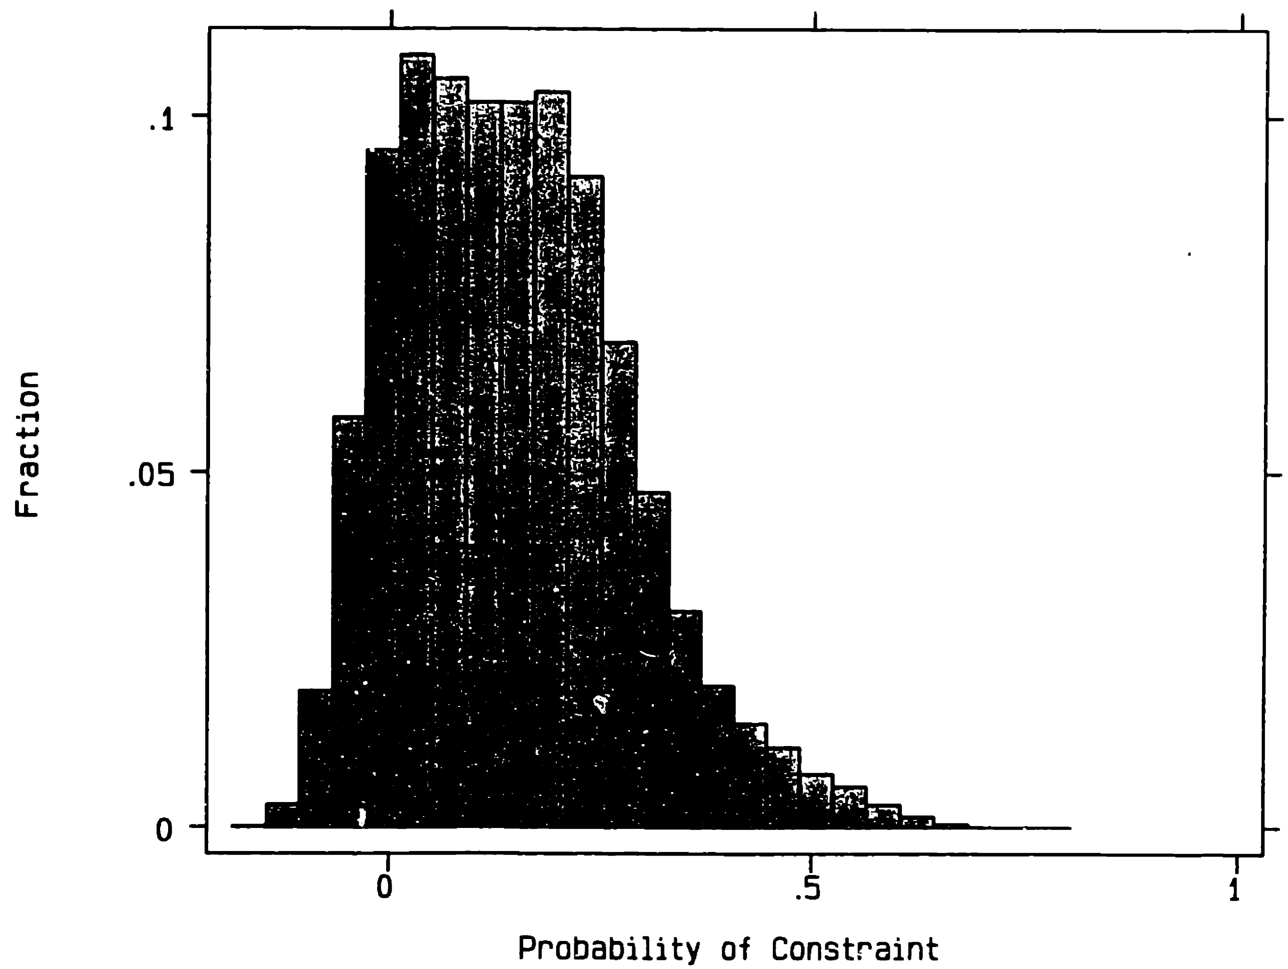

Table 1
Sample means

	PSID (1970 -1984)	PSID (1984)	CEX (1980 -1991)	SCF (full sample)	SCF (constr.)	SCF (unconst.)
Demographics						
Age of head	47.9	48.8	43.8	46.9	37.8	48.7
Married	0.666	0.645	0.704	0.682	0.557	0.707
Divorced	0.107	0.136	0.138	0.121	0.203	0.105
Non-white	0.119	0.140	0.136	0.130	0.282	0.101
Male head	0.747	0.728	0.808	0.770	0.709	0.783
One adult	0.279	0.300	0.213	0.245	0.349	0.225
Two adults	0.583	0.579	0.571	0.569	0.503	0.581
Three or more adults	0.138	0.120	0.216	0.186	0.149	0.193
No kid	0.551	0.584	0.505	0.555	0.429	0.580
One kid	0.168	0.159	0.203	0.173	0.197	0.168
Two kids	0.162	0.161	0.182	0.178	0.234	0.167
Three or more kids	0.119	0.094	0.110	0.094	0.140	0.085
Region						
North east	0.240	0.231	0.214	0.210	0.234	0.206
North central	0.295	0.280	0.272	0.295	0.220	0.310
South	0.295	0.311	0.282	0.337	0.357	0.333
West	0.168	0.177	0.235	0.158	0.189	0.151
Schooling						
No high school	0.148	0.117	0.076	0.140	0.103	0.147
Some high school	0.156	0.154	0.107	0.128	0.149	0.124
High school graduate	0.344	0.354	0.317	0.317	0.291	0.321
Some college	0.145	0.165	0.231	0.193	0.254	0.181
College graduate[a]	0.126	0.135	0.269	0.111	0.091	0.115
Post graduate	0.064	0.060	---	0.112	0.111	0.112
Income and employment						
Disposable income	26,419	26,629	28,055	27,364	21,389	28,532
Head employed	0.715	0.670	---	0.727	0.734	0.726
Union member	0.220	0.152	---	0.206	0.209	0.206
Number of years worked full time	21.8	21.9	---	21.2	14.4	22.6
Number of observations	33,874	3,078	22,253	2,139	350	1,789

a. College graduates include post graduates for the CEX.
Notes: PSID samples include the poverty subsample; means are weighted by family weights. CEX and SCF means are unweighted. Constrained households in the SCF are defined as those who are reported being denied credit or discouraged from borrowing.

Table 2
Percentiles of liquid assets and wealth in the 1984 PSID and 1983 SCF

Liquid assets

Liquid asset measure	Survey	Sample size	10%	25%	50%	75%	90%
Blown up	PSID	3,078	0	0	1,339	26,473	96,563
Narrow definition	PSID	2,836	0	310	2,826	13,564	41,824
Wide definition	PSID	2,466	0	339	3,956	16,956	46,345
Blown up	SCF	1,695	0	0	468	11,502	44,278
Narrow definition	SCF	1,695	23	406	1,973	9,192	29,626
Wide definition	SCF	1,695	23	435	2,263	11,293	36,009

Net wealth

Wealth measure	Survey	Sample size	10%	25%	50%	75%	90%
Blown up	PSID	2,957	0	1,451	32,538	86,132	169,762
Narrow definition	PSID	2,731	0	3,391	30,520	77,430	123,000
Wide definition	PSID	2,382	0	2,825	28,260	71,213	125,471
Blown up	SCF	1,619	0	2,339	31,805	71,426	131,318
Narrow definition	SCF	1,619	174	3,203	29,821	67,061	117,679
Wide definition	SCF	1,619	174	3,511	30,219	69,471	123,673

Notes: The blown up measure of liquid assets is obtained by dividing asset income by the interest rate. The other two measures are obtained directly from asset stocks. The narrow definition excludes bonds, stocks, and IRAs; the wide definition includes these assets. Net wealth is the sum of liquid asset and housing wealth net of outstanding mortgage principal.

Table 3
A comparison of asset based sample splits and the sample split based on the self reported indicator of liquidity constraints

PSID 1984:
Zeldes' split versus split based on asset stocks

Zeldes' splits	Sample size	Constr. accord. to both splits	Constr. in Zeldes' split only	Constr. in stock based split only	Unconstr. in both splits
	(1)	(2)	(3)	(4)	(5)
Split 1: constrained if liquid assets below 2 months income	2,080 (100%)	930 (45%)	296 (14%)	196 (9%)	658 (32%)
Split 2: constrained if liquid assets equal zero[a]	805 (100%)	265 (33%)	114 (14%)	18 (2%)	408 (51%)
Split 3: constrained if net wealth below 2 months income	2,001 (100%)	444 (22%)	104 (5%)	40 (2%)	1413 (71%)

SCF 1983:
Zeldes' split versus split based on self-reported indicator of liquidity constraint

Zeldes' splits	Sample size	Constr. accord. to both splits	Constr. in Zeldes' split only	Constr. in stock based split only	Unconstr. in both splits
	(1)	(2)	(3)	(4)	(5)
Split 1: constrained if liquid assets below 2 months income	3,656 (100%)	532 (15%)	1,470 (40%)	186 (5%)	1,468 (40%)
Split 2: constrained if liquid assets equal zero[a]	2,935 (100%)	454 (15%)	1,158 (39%)	151 (5%)	1,172 (40%)
Split 3: constrained if net wealth below 2 months income	3,656 (100%)	372 (10%)	616 (17%)	346 (9%)	2,322 (64%)

a. Unconstrained group for split 2 has assets greater than six month income, middle group is omitted from this split.
Notes: The samples used here differ from the samples used in the analysis below. The poverty subsample is eliminated from the PSID and the cross tabulations are unweighted. Zeldes' sample selection procedures are followed as closely as possible.

Table 4
Linear probability models for being constrained

	PSID Split 1	SCF Split 1	SCF Self-reported constraint indicator		
Demographics	(1)	(2)	(3)	(4)	(5)
Age	-0.008 (0.002)	-0.002 (0.004)	-0.010 (0.003)	-0.009 (0.003)	-0.009 (0.003)
Age squared/100	-0.004 (0.002)	-0.006 (0.004)	0.004 (0.003)	0.004 (0.003)	0.002 (0.003)
Male head	0.022 (0.024)	-0.014 (0.043)	---	0.041 (0.036)	0.024 (0.040)
Married	-0.087 (0.031)	0.009 (0.052)	---	-0.052 (0.042)	-0.055 (0.041)
Divorced	0.070 (0.021)	0.127 (0.036)	---	0.038 (0.032)	0.038 (0.033)
One adult	-0.123 (0.024)	-0.095 (0.040)	---	0.046 (0.035)	0.049 (0.035)
Three or more adults	0.144 (0.017)	0.049 (0.030)	---	0.018 (0.022)	0.026 (0.023)
One kid	0.090 (0.015)	0.041 (0.031)	---	0.023 (0.024)	0.026 (0.024)
Two kids	0.115 (0.016)	0.062 (0.033)	---	0.055 (0.026)	0.061 (0.026)
Three or more kids	0.161 (0.018)	0.053 (0.040)	---	0.084 (0.037)	0.088 (0.037)
Change in the number of adults	-0.005 (0.006)	0.014 (0.015)	-0.009 (0.011)	0.017 (0.014)	0.017 (0.014)
Change in the number of kids	0.010 (0.007)	0.002 (0.015)	0.001 (0.012)	-0.013 (0.013)	-0.013 (0.013)
Non-white	0.134 (0.015)	0.214 (0.027)	0.193 (0.029)	0.182 (0.029)	0.177 (0.029)
Region					
North east	0.008 (0.019)	-0.030 (0.032)	-0.007 (0.026)	0.002 (0.026)	0.006 (0.026)
North central	-0.005 (0.018)	-0.006 (0.030)	-0.068 (0.024)	-0.056 (0.024)	-0.052 (0.024)
South	0.045 (0.017)	0.000 (0.029)	-0.048 (0.024)	-0.036 (0.024)	-0.033 (0.024)

Table 4 - continued

	PSID Split 1	SCF Split 1	SCF Self-reported constraint indicator		
Schooling	(1)	(2)	(3)	(4)	(5)
Some high school	-0.085 (0.021)	-0.083 (0.037)	---	0.037 (0.028)	0.033 (0.029)
High school graduate	-0.144 (0.020)	-0.163 (0.034)	---	0.002 (0.023)	0.004 (0.024)
Some college	-0.185 (0.022)	-0.206 (0.038)	---	0.042 (0.027)	0.047 (0.029)
College graduate	-0.267 (0.025)	-0.273 (0.044)	---	-0.003 (0.032)	0.011 (0.034)
Post graduate	-0.280 (0.031)	-0.365 (0.045)	---	-0.048 (0.034)	0.059 (0.037)
Income and Employment					
Log income	-0.182 (0.010)	-0.178 (0.019)	-0.067 (0.012)	-0.055 (0.015)	-0.041 (0.015)
Employed	---	---	---	---	-0.014 (0.051)
Union contract	---	---	---	---	-0.022 (0.022)
Full time years since 18	---	---	---	--	0.002 (0.001)
11 industry dummies	no	no	no	no	yes
7 occupation dummies	no	no	no	no	yes
Constant	3.026 (0.108)	1.597 (0.185)	1.192 (0.122)	0.980 (0.150)	0.929 (0.160)
R^2	0.300	0.234	0.126	0.139	0.153
No. of observations	33,874	2,139	2,139	2,139	2,139

Notes: **The dependent variable equals 1 if a household is liquidity constrained, 0 otherwise. In columns 1 and 2 the household is constrained if assets are below two months income. In columns 3 to 5 a household is constrained if he reported being denied credit or discouraged from borrowing. Excluded attributes are households with two adults, no kids, no high school, living in the west. Heteroskedasticity robust standard errors in parentheses. Standard errors for PSID sample are also robust to random individual effects.**

Table 5
Euler equations estimates splitting the sample by wealth
PSID

	Full sample	High wealth group		Low wealth group	
	poverty (1)	poverty (2)	Census (3)	poverty (4)	Census (5)
Constant	0.228	0.281	0.267	0.257	0.249
	(0.031)	(0.067)	(0.071)	(0.038)	(0.054)
Change in the number of adults	0.132	0.135	0.138	0.130	0.133
	(0.004)	(0.008)	(0.008)	(0.004)	(0.006)
Change in the number of children	0.101	0.107	0.087	0.098	0.094
	(0.003)	(0.007)	(0.016)	(0.003)	(0.005)
Age	-0.0056	-0.0090	-0.0087	-0.0046	-0.0044
	(0.0008)	(0.0016)	(0.0016)	(0.0010)	(0.0014)
Age squared/100	0.0033	0.0060	0.0056	0.0023	0.0021
	(0.0008)	(0.0015)	(0.0016)	(0.0010)	(0.0015)
Lagged disposable income	-0.005	-0.001	0.000	-0.011	-0.010
	(0.003)	(0.006)	(0.006)	(0.004)	(0.006)
Number of observations	33,874	8,865	7,877	25,009	12,636

Notes: Dependent variable is the three year change in log food consumption. The two groups are obtained by classifying households as low wealth if they have assets less than two months income (split 1). Columns labeled poverty include the poverty subsample of the PSID and are weighted using the family weight. Columns (3) and (5) exclude the poverty subsample and use unweighted regressions. Standard errors are adjusted for the overlapping data structure. See text for details.

Table 6
Euler equation estimates of the two sample switching regression model
PSID

	OLS small set poverty (1)	OLS basic set poverty (2)	OLS large set poverty (3)	OLS basic set Census (4)
	Unconstrained regime			
Constant	0.480 (0.075)	0.376 (0.061)	0.275 (0.053)	0.384 (0.078)
Change in the number of adults	0.136 (0.006)	0.139 (0.006)	0.138 (0.005)	0.140 (0.008)
Change in the number of children	0.119 (0.007)	0.118 (0.006)	0.115 (0.006)	0.118 (0.008)
Age	-0.0095 (0.0019)	-0.0073 (0.0015)	-0.0065 (0.0013)	-0.0062 (0.0019)
Age squared/100	0.0057 (0.0016)	0.0042 (0.0013)	0.0038 (0.0012)	0.0032 (0.0017)
Lagged disposable income	-0.015 (0.006)	-0.013 (0.005)	-0.007 (0.005)	-0.016 (0.007)
	Constrained regime			
Constant	0.275 (0.231)	0.075 (0.195)	0.132 (0.164)	-0.144 (0.263)
Change in the number of adults	0.108 (0.034)	0.096 (0.031)	0.105 (0.025)	0.114 (0.042)
Change in the number of children	0.019 (0.030)	0.029 (0.026)	0.043 (0.021)	0.003 (0.034)
Age	-0.0033 (0.0071)	-0.0031 (0.0060)	-0.0069 (0.0051)	-0.0027 (0.0079)
Age squared/100	0.0022 (0.0074)	0.0010 (0.0063)	-0.0060 (0.0054)	-0.0007 (0.0083)
Lagged disposable income	-0.032 (0.024)	-0.002 (0.021)	-0.003 (0.017)	0.022 (0.028)
Number of observations	33,874	33,874	33,874	20,513

Notes: Dependent variable is the three year change in log food consumption. Regressions use one of three instrument sets, referred to as small, basic and large set; see text for explanations. Columns labeled poverty include the poverty subsample of the PSID and second stage regressions are weighted using the family weight. Column (4) excludes the poverty subsample and uses unweighted regression. Standard errors are adjusted for two sample estimation and for the overlapping data structure. See text for more details.

Table 7
Euler equation estimates, sample splits and switching regression models
CEX
Only Coefficients on Lagged Income Shown

	Food	Strictly nondurables	Nondurables
	Full sample		
	-0.000	-0.000	0.001
	(0.004)	(0.003)	(0.003)
	Splitting the sample by wealth to income		
High wealth	0.007	0.002	0.002
	(0.005)	(0.004)	(0.004)
Low wealth	-0.014	-0.006	-0.003
	(0.007)	(0.006)	(0.005)
	Splitting the sample by income		
High income	-0.003	0.003	0.003
	(0.008)	(0.006)	(0.006)
Low income	0.002	-0.001	-0.000
	(0.010)	(0.006)	(0.006)
	Two sample switching regressions		
Unconstrained	-0.011	-0.004	-0.004
	(0.009)	(0.006)	(0.006)
Constrained	0.022	0.006	0.004
	(0.018)	(0.013)	(0.013)
Number of observations	22,253	22,253	21,992

Notes: Dependent variable is the quarterly change in log food consumption. Low wealth consumers have checking and savings balances less than one month worth of income in the wealth to income split. Low income consumers have income of less than $16,000. Standard errors shown in parentheses are not adjusted for two sample estimation. See text for more details.

Table 8
Euler equation estimates of the two sample switching regression model
CEX
Only coefficients on lagged income shown

	Food	Strictly Nondurables	Nondurables
	Base sample and specification		
Unconstrained	-0.011	-0.005	-0.002
	(0.009)	(0.007)	(0.007)
Constrained	0.027	0.015	0.008
	(0.18)	(0.013)	(0.013)
Number of observations	20,161	20,161	20,161
	Excluding observations for 1980 and 1981		
Unconstrained	-0.011	-0.009	-0.007
	(0.009)	(0.007)	(0.007)
Constrained	0.021	0.017	0.007
	(0.019)	(0.014)	(0.014)
Number of observations	17,293	17,293	17,293
	Excluding income from the instrument list		
Unconstrained	-0.003	-0.001	0.000
	(0.007)	(0.005)	(0.005)
Constrained	0.008	0.006	0.004
	(0.024)	(0.017)	(0.017)
Number of observations	20,161	20,161	20,161

Notes: Dependent variable is the quarterly change in log food consumption. Base sample in this table excludes farmers and self-employed. Other specifications make additional changes. Standard errors shown in parentheses are not adjusted for two sample estimation. See text for more details.

Table 9
The conditional distribution of consumption growth: simulation results
(100 replications)

A. No liquidity constraints imposed, income growth

	OLS	25th	50th	75th
Standard: y_{it}	0.000	0.040	0.032	-0.003
	(0.009)	(0.013)	(0.011)	(0.012)
Interactions: y_{it}	0.002	0.042	0.033	0.003
	(0.015)	(0.020)	(0.019)	(0.021)
$\hat{L}_{it}$	0.081	0.021	0.018	0.012
	(0.251)	(0.296)	(0.351)	(0.381)
$\hat{L}_{it}$* $\ln y_{it}$	-0.008	-0.008	-0.006	-0.007
	(0.055)	(0.071)	(0.075)	(0.079)

B. Liquidity constraints imposed, income growth

	OLS	25th	50th	75th
Standard: y_{it}	-0.045	0.014	-0.003	-0.059
	(0.010)	(0.013)	(0.011)	(0.011)
Interactions: y_{it}	0.002	0.044	0.037	-0.003
	(0.016)	(0.022)	(0.020)	(0.020)
$\hat{L}_{it}$	1.498	1.049	1.482	1.781
	(0.307)	(0.426)	(0.385)	(0.391)
$\hat{L}_{it}$* $\ln y_{it}$	-0.227	-0.153	-0.221	-0.267
	(0.064)	(0.088)	(0.075)	(0.081)

C. Liquidity constraints imposed, no income growth

	OLS	25th	50th	75th
Standard: y_{it}	-0.064	-0.060	-0.059	-0.060
	(0.018)	(0.023)	(0.022)	(0.023)
Interactions: y_{it}	-0.045	-0.058	-0.050	-0.034
	(0.033)	(0.047)	(0.037)	(0.039)
$\hat{L}_{it}$	0.164	-0.213	-0.010	0.375
	(0.526)	(0.791)	(0.598)	(0.626)
$\hat{L}_{it}$* $\ln y_{it}$	-0.089	-0.020	-0.045	-0.117
	(0.108)	(0.168)	(0.125)	(0.127)

Notes: The dependent variable is the change in consumption. The total number of observations is 8,350 (two for each household). The standard specification is $c_{it+1} - c_{it} = \alpha + \beta y_{it}$. The switching regression specification with interaction terms is $c_{it+1} - c_{it} = \alpha_0 + \alpha_1 p_{it} + \beta_0 y_{it} + \beta_1 p_{it} y_{it}$

Table 10
The conditional distribution of consumption growth:
Estimates from the PSID and CEX

A. PSID, food consumption

	OLS	25th	50th	75th
Standard: $\ln y_{it}$	-0.004	0.021	0.002	-0.033
	(0.003)	(0.004)	(0.003)	(0.004)
Interactions: $\ln y_{it}$	-0.004	0.001	0.001	-0.010
	(0.004)	(0.006)	(0.005)	(0.006)
$\hat{L}_{it}$	-0.033	-1.096	-0.192	1.110
	(0.200)	(0.281)	(0.201)	(0.264)
$\hat{L}_{it} * \ln y_{it}$	-0.018	0.038	-0.005	-0.078
	(0.020)	(0.029)	(0.021)	(0.027)

B. CEX, food consumption

	OLS	25th	50th	75th
Standard: $\ln y_{it}$	-0.000	0.014	-0.000	-0.017
	(0.004)	(0.005)	(0.004)	(0.005)
Interactions: $\ln y_{it}$	-0.011	-0.002	-0.002	-0.006
	(0.009)	(0.009)	(0.007)	(0.009)
$\hat{L}_{it}$	-0.534	-1.101	-0.404	0.209
	(0.291)	(0.335)	(0.259)	(0.312)
$\hat{L}_{it} * \ln y_{it}$	0.033	0.024	0.012	-0.003
	(0.024)	(0.028)	(0.022)	(0.027)

C. CEX, strictly nondurable consumption

	OLS	25th	50th	75th
Standard: $\ln y_{it}$	-0.000	0.054	-0.002	-0.003
	(0.003)	(0.004)	(0.003)	(0.002)
Interactions: $\ln y_{it}$	-0.004	-0.014	0.002	0.020
	(0.006)	(0.006)	(0.005)	(0.007)
$\hat{L}_{it}$	-0.141	-0.829	-0.046	0.704
	(0.208)	(0.218)	(0.196)	(0.235)
$\hat{L}_{it} * \ln y_{it}$	-0.010	0.048	-0.008	-0.054
	(0.017)	(0.019)	(0.017)	(0.021)

Notes: The dependent variable is the one year growth rate in consumption in the PSID and the quarterly growth rate in the CEX. The number of observations is 26,426 in the PSID and 22,253 in the CEX. Each regression includes demographics and the CEX regression include 11 month of the year dummies. Reported standard errors are not corrected for two sample estimation.

Chapter 2. The Response of Household Consumption to Income Tax Refunds[1]

0. Introduction.

High-frequency tests of the canonical life-cycle/permanent-income [LC/PI] theory on micro-data have yielded mixed results. The most prominent of such tests, excess sensitivity tests, have generally been hampered by two difficulties. First, it is hard to isolate the predictable or transitory components of income, to which consumption is not supposed excessively to respond. Most excess sensitivity tests proceed by instrumenting for income, but since the available instruments are usually poor, the results are typically prejudiced against finding significant sensitivity.[2] Second, even when such sensitivity has been found, it has been difficult to pinpoint and interpret its source. It has been especially hard to distinguish violations of the canonical assumptions on preferences, from those on technology, in particular due to liquidity constraints.[3]

---

[1]I thank Jason Abrevaya, Joshua Angrist, Olivier Blanchard, Ricardo Caballero, Jerry Hausman, Jonathan Parker, Jack Porter, and Richard Thaler for many valuable discussions. Special thanks go to Steve Pischke, who has been extraordinarily generous with his help throughout this project; and to Tom Moehrle and Bill Passero of the Bureau of Labor Statistics, for authoritative and patient help with the data. Of course, none is responsible for any errors herein. Financial support from the National Science Foundation is also gratefully acknowledged.

[2]See e.g. Shea [1993].

[3]Deaton [1992, Ch. 6] and Attanasio [1994], among others, discuss this particular difficulty in detail.

This paper avoids both these difficulties by examining the response of households' consumption to their income tax refunds. First, since the receipt and magnitude of a refund depend on events in the previous calendar year, a refund is a predictable kind of income as regards consumption-decisions in the year of its receipt. Consequently consumption should not increase on receipt of a refund, under the LC/PI theory. This consequence assumes that because a refund is in principle predictable once into the year, its recipient is then at least implicitly cognizant of the amount to be received.[4] For consumption-decisions made within the typically 4-12 week lag between the filing of a tax return and refund-receipt, this assumption is obviously sound. But consumption-decisions made earlier might be affected, for any households that find on completing their returns that their refunds are going to be somewhat more or less than expected. Even so, since a refund is also a transitory kind of income, under the theory there should be little response in consumption. That is, any "surprise" on completing a return is unlikely to carry much information about a household's permanent income that was not already known to it by the end of the previous year; and under the theory only innovations to permanent income warrant substantial changes in consumption. Further, this paper also examines consumption well after the receipt of refunds, when uncertainty is no longer an issue.

Second, as regards the source of any excess sensitivity, any response of consumption after refund-receipt is unlikely to be due to liquidity constraints. And even around the time of receipt, as in many other studies the response of households likely to be constrained can be contrasted with the response of those unlikely to be. The interpretation of this contrast here, however, will be more transparent than in most previous studies.[5]

---

[4]Most excess sensitivity tests make an even stronger assumption, namely that the household is aware of anything that the econometrician can predict using lagged information.

[5]As briefly discussed below, I am also exploring the response of consumption to tax payments. Since these decrease income, any response is again unlikely to follow from liquidity constraints.

In addition to testing for excess sensitivity, this paper estimates the magnitude of consumption's response to refunds under a specific alternative hypothesis. These estimates bear directly on the efficacy of fiscal policy, for instance on the recently proposed tax cuts. They bear on Ricardian equivalence in particular, since a refund is not an addition to net wealth, but merely reflects overwithholding in the previous year. The results are therefore informative even about fiscal stimuli of low cost to the government.[6]

The results also bear on some of the leading recent behavioral theories of saving, which have not heretofore received much empirical scrutiny. In particular, the theories of mental accounts and self-control/forced saving have implications for consumption's response to refunds that are evaluated herein.

This paper uses the Consumer Expenditure Survey, which of the leading U.S. micro data sets has the most comprehensive coverage of expenditure. Previewing the results, there is significant and interpretable excess sensitivity of consumption to refunds. This sensitivity is due in part to sharply increased spending on nondurables at the time of refund-receipt by those likely to be liquidity constrained. However, there is also evidence of substantial increases in spending, mostly on durables, by those unlikely to be constrained; as well as in spending after most refund-receipt, particularly while on trips in the summer, which cannot easily be attributed to liquidity constraints or problems of self-control. The effect of refunds on the timing of durables purchases is especially puzzling, not easily explained by the standard models of durables. Counter to some views of mental accounts, the response of consumption is increasing in the size of the refund.

The outline of the paper is as follows. Section I provides background on related studies and about income tax refunds. The specification of the Euler equation and its

---

[6]Since with few exceptions households do not receive interest on the amount overwithheld, refunds actually represent a net loss to households. Slemrod et al. [1994] estimate that the interest that would have been gained in 1989 just from claiming refunds at the very beginning of the year, instead of when they were actually claimed, was on the order of $1 billion. This averages to about $13.50 per tax return.

Of course these results do not themselves justify an activist fiscal policy. Indeed, the opponent of activism might point to the substantial effects of refunds as reason to curtail the degree of overwithholding.

generalization under a specific alternative hypothesis are discussed in Section II. Section III describes the data. Section IV reports the results, for the basic specification as well as for various extensions. The conclusion, Section V, is followed by an Appendix further describing the data.

I. Background: Related Studies and Income Tax Refunds.

1. Related Studies.

Since the path-breaking paper of Hall [1978], excess sensitivity tests have been the most prominent of the high-frequency tests of the LC/PI theory. This tradition of tests looks for any response of the change in consumption to predictable income (or changes in income). An earlier tradition of tests focused on the relative magnitude of consumption's response to permanent versus transitory income. There was particular focus on the response--often in levels, not differences--to windfalls and changes in fiscal policy. In examining the response of the change in consumption to refunds, which are both predictable and transitory, this paper is heir to both traditions. Recent behavioral theories of saving represent a third tradition, concerned less with testing the LC/PI theory than with developing alternatives to it. This Section discusses these three traditions in turn.

The clearest findings of excess sensitivity have been based on aggregate data. However, the assumptions required for the Euler equation to apply in the aggregate are very strong. On the other hand, as Angus Deaton concluded in his recent survey of consumption, "the evidence against the theory in the micro data is weaker and less transparent" [1992, p. 160].

The most influential micro-studies, notably Hall and Mishkin [1982] and Zeldes [1989], and more recently Shea [1993], have made use of the Panel Study of Income Dynamics [PSID]. Unfortunately the information on consumption in the PSID is limited to food, which is a necessity and not obviously separable from the rest of consumption. As a result their findings of excess sensitivity might be due in part to the inapplicability of the Euler equation to food (Attanasio and Weber [1992, 1994], Meghir and Weber

[1993]). Other studies have attributed this excess sensitivity to, for instance, the particular sample used (e.g. Mariger and Shaw [1990], Jappelli, Pischke, and Soulcles [1994]) or measurement error (Altonji and Siow [1987]).[7]

Some of the latest studies of excess sensitivity have employed the Consumer Expenditure Survey [CEX], which has much broader coverage of consumption. One of the first, Lusardi [1992], as well as a number of papers by Attanasio and his colleagues, found the response of nondurable consumption to (instrumented) income insignificant.[8] To find a significant response, Lusardi [1993] had to predict the income-changes of her CEX sample on the basis of income changes of similar households in the PSID.

The test of excess sensitivity closest in spirit to this paper deserves special recognition. In a 1990 working paper, Wilcox investigated the aggregate time-series relationship between the value of refunds disbursed by the Treasury each week, and the monthly NIPA series of personal consumption expenditure. In his baseline specification he found an instantaneous effect on consumption of 7.5% of the refunds.[9] This effect was spread roughly evenly across durables, nondurables, and services. However, his results are rather sensitive to the particular instruments used for other types of income.

---

[7]Further, Runkle [1991]'s attempt to replicate Zeldes's result failed to find excess sensitivity, although Runkle used a different sample within the PSID and different techniques.

Jappelli, Pischke, and Souleles [1994] generally tried to replicate Zeldes's sample and techniques. They found significant excess sensitivity in Zeldes's sample, but this sensitivity disappeared on adding additional years to the sample.

[8]A number of the later papers have argued that nonseparabilities, between different consumption goods or between consumption and labor supply, can account for reported excess sensitivity (Meghir and Weber [1993], Attanasio and Weber [1992,1994]). This paper like Lusardi's uses broad groupings of consumption, like nondurables, which should minimize any problem of the first kind of nonseparability. As for the second, Garcia, Lusardi and Ng [1994] find that the sensitivity persists even after controlling for labor supply. Also, the main story for such a nonseparability does not apply here, because contemporaneous changes in labor supply cannot affect predetermined refunds.

[9]This figure is for the day of receipt. From the tables it is not clear what the accumulated effect is within periods of different length. The specification assumes that any response of consumption is complete within 90 days.

And the response he did find is obviously a rather complex aggregate. For instance, in the same season (the spring) in which most households are receiving refunds, a smaller but still large number is paying its April taxes, which are on average greater in magnitude than refunds. Also, working in the aggregate obscures any differential response to refunds of liquidity constrained households, and any nonlinearity in households' responses.

This paper differs from these tests of excess sensitivity in a number of ways. First, it uses the best available micro-data on consumption, in the CEX. Second, since refunds are predetermined, there is no need to instrument for them or to predict them using other data. Third, the response of households to their tax refunds seems to be a particularly interpretable "experiment", with immediate implications for fiscal policy and models of consumer behavior.[10]

Consider now the earlier tradition, on the response of consumption to changes in fiscal policy and windfalls. A few papers have studied the 1975 tax rebate, all using aggregate data.[11] Modigliani and Steindel [1977] estimated the dynamic effect of the rebate on total consumption. The immediate effect was modest, 5-12% of the rebate being spent within the first quarter; but there was a substantial delayed effect, over 80% within 6 quarters.[12] Blinder [1981] found a somewhat larger immediate effect, 16%, but

[10]In addition to Shea [1993], Wilcox [1990, 1992], and Poterba [1988], a number of other papers look at the response of consumption to "experiments." Eg, Levinson [1993] and Paxson [1992]. A colleague of mine, Jonathan Parker, is looking at the reaction to changes in social security taxes.

[11]To try to stimulate a weak economy, the U.S. government returned to filers 10% of their 1974 taxes, up to a maximum of $200. This amounted to about $8 billion, in the second quarter of 1975.

[12]Modigliani and Steindel compared the estimates of four different models. The models which did not finely disaggregate consumption, their own model and the MPS model, estimated a delayed but ultimately potent response of total consumption. The MPS model found more of the rebate was spent in the quarter of receipt, 16% rather than the 5-12% of Modigliani and Steindel; under either model, only about 24-34% was spent within three quarters, but both models found over 80% spent within six quarters. The two disaggregated models, from DRI and Michigan, gave a much more potent immediate effect, 50-58% within two quarters. In the DRI model, 60% of the total change in consumption was due to nondurables and rent; in the Michigan model, most of it was due

a smaller delayed effect, only 34% within six quarters. Both studies concluded that tax rebates are not an effective way to stimulate the economy quickly.

Deliberately using the framework of excess-sensitivity tests, Poterba [1988] considered the effect of the rebate on monthly aggregate nondurable consumption, conditional on the lag of this consumption. He found an immediate monthly effect of 18-25% of the rebate, much larger than the studies above.[13] Also, considering the leading fiscal changes of the last few decades, he found that consumption generally did not react on the announcement of the changes, as it should under the LC/PI theory, but only after their implementation.[14]

Of the studies in the earlier tradition, all discussed so far have used aggregate data. The post-war literature on windfalls instead examined at the micro-level the consumption of windfall-recipients. In 1950 the U.S. government unexpectedly paid National Life Insurance Dividends, averaging about $175, to certain WWII veterans. Estimating a consumption function in levels on annual data, Bodkin [1959] found statistically significant and huge contemporaneous marginal propensities to consume [mpc's] out of the dividend: .72 for nondurables, and .97 for total consumption.[15] Similarly, Kreinen [1961] examined the response of certain Israelis to restitution-payments from Germany in the 1950s. But he found mpc's of only .16 for nondurables and .17 in total.

One of the leading explanations for the significant difference in these results was that the Israeli mpc was lower because the corresponding windfalls were generally much

---

to furniture in particular. Finding these last figures implausible, Modigliani and Steindel favored the results of the first two models.

[13]With only a few observations for the rebate, the specification is unavoidably constrained. Eg, it does not allow for delayed effects, and it constrains the reaction in a given month to be of equal magnitude but oppositely signed to that in the previous month.

[14]Wilcox [1989] provided a similar result, that social security recipients respond only to the actual increases in their social security benefits, not to the announcement of the increases.

[15]Jones [1960] looked at more disaggregated groupings of consumption. Most notably, for food he found a large mpc of .25, larger than for housing and clothing.

larger relative to the recipients' income. The various analyses of this explanation came to conflicting conclusions, however, and so the difference remains an open issue.[16] This explanation will be tested below.

Two shortcomings of these studies of windfalls are noteworthy. First, the samples were relatively small (only 120 for Kreinen), and the windfall-recipients were rather special people. Second, the estimation of consumption in levels is open to mispecification which working with Euler equations in differences avoids. E.g., Bodkin was well aware that his dividend variable might have been picking up the correlation of the dividend with omitted variables in turn correlated with permanent income.[17] Using Euler equations, by contrast, only innovations to income matter, and these of course are uncorrelated with any predetermined income. In looking at refund-recipients, a much larger and more representative group, and by estimating in differences, this paper avoids both shortcomings.[18]

Consider finally the third tradition of studies. Of the leading behavioral theories of saving, two have direct implications for refunds. First, the "mental accounts" theory, like the argument above about windfalls, suggests that if refunds are large compared to income, they will be classified--i.e., put into one's mental accounts--as a type of income

---

[16]On adding the restitution-payment squared divided by income to his regression, Kreinen got a positive result, not negative as he expected. Using this data and similar data from the subsequent year (for a total of 300 observations), Landsberger [1966] sorted his sample into 5 groups, according to the size of the restitution-payment relative to income. He did find the mpc to decrease. Bodkin [1966] split his American sample into those with dividends greater and less than $300. The group with large dividends turned out to have a larger mpc, the opposite result, though the difference was not significant. He also added the dividend squared to his regression, but it was insignificant.

Note that refunds, although transitory, are not unexpected like windfalls.

[17]On adding likely such variables to Bodkin's original regression, Bird and Bodkin [1965] found the mpc's decreased substantially, to .38 for nondurables and .65 for total consumption.

[18]Also at the micro-level, McNees [1973] analyzed surveys of refund-recipients in particular. 45% of his sample said they spent their 1972 refund, 24% saved it, and 28% used it to pay off debt and bills. The propensity to consume decreased with the size of the refund and was lesser when the refund was larger than expected.

to be saved.[19] As just noted, this nonlinearity in the response to refunds is tested below. Second, the "self-control" theory argues that some people deliberately overwithhold to force themselves to save, perhaps to save up enough for a durable.[20] This paper directly examines the link between refunds and changes in consumption, and so saving; and the composition of these changes between durables and nondurables. It also tests general implications of the theory. E.g., if liquidity constraints are not a problem, people might be able to undo such forced savings, for instance by taking out loans against their refunds, or drawing down their liquid assets.[21] Or, if people are not "splurging" on receipt of their refunds, or if they splurge well after receipt, then there might be more at play than just a straightforward problem of self-control.[22]

---

[19]Eg, Thaler [1990], Thaler and Loewenstein [1992].

A scene from Shaw's Pygmalion colorfully illustrates the idea. Mr. Doolittle, father of Eliza, asks Dr. Higgins to pay 5 pounds to use Eliza in his linguistic experiment. Higgins finds Doolittle to be such a character that he offers to pay twice as much. Doolittle refuses the increase: "Don't you be afraid that I'll save [the 5 pounds] and live idle on it. There wont be a penny of it left by Monday ... Just one good spree for myself and the missus ... [But] she won't have the heart to spend ten; and perhaps I shouldnt neither. Ten pounds is a lot of money: it makes a man feel prudent like; and then goodbye to happiness."

[20]Eg, Neumark [1992]. This view is commonly expressed in anecdotes and in the popular and financial presses as well.

[21]One particular form of undoing has recently become popular: some tax services offer to speed up the receipt of a refund, for a fee.

See Laibson [1993] for a clever model of self-control in the presence of both liquid and illiquid assets.

[22]Some of the recent work in non-expected utility theory, e.g. on loss aversion, also has implications for refunds. In the context of consumption, eg, Shea [1993], Garcia et. al [1994], and Bowman et al. [1993] stress the asymmetry of the response to gains in income, rather than loses.

2. Income Tax Refunds.

Federal income tax refunds are surprisingly widespread and substantial. As Table 1 records, roughly 3/4 of the 100 million or so tax-filers each year receive refunds, amounting in 1992 to over $80 billion. Throughout the 1980's and 1990's, the amount of refunds represented from 15% to over 20% of the total individual tax liability. The average real refund has fluctuated between $700 and $850 (1982-84 $) in these years. The median real refund was about $450 to $500 at the beginning of the 1980's.[23] For many households these amounts obviously represent a sizeable fraction of income, and come in a single check.[24]

In the analysis below the timing of households' receipt of refunds will be important. The Treasury's Daily Statements are informative here, providing the distribution of the total amount of refunds disbursed. Table 2 describes a typical year. About 80% of refunds are mailed out in March, April and May. There is some variation from year to year--e.g. in the middle of the 1980's more refunds went out in May than in March, the opposite of the other years--but the basic pattern remains the same over the period of interest here.[25]

---

[23]Based on Wilcox [1990]'s figure of a median refund for 1979 of $350-400. I presume that this is for refunds received in 1980. Wilcox also reports that the median adjusted gross income in 1979 was about $ 12,000, which comes to about $15,000 - 17,000 in real terms, depending on its timing.

[24]Filers may credit their overwithholding against their taxes for the following fiscal year. The fraction which does so is rather small, eg about 4% of returns in 1991. Further, most of these filers credit their entire overpayment, so receive no refund, and so are dropped in the analysis below.

There remains the possibility that someone who was planning to credit her refund decides not to do so. However, such scenarios are probably relatively uncommon and unsystematic, and so would largely wash out in the cross-section. In particular, they seem unlikely to be primarily responsible for the pattern and magnitude of the results below. (And even if they are, this would be of economic interest in itself.)

[25]While this distribution is not the same as that for the fraction of returns-due-refunds that is processed, it should be close to it.

Corroborating evidence at the level of individual returns comes from Slemrod et al. [1994], who look at the timing of the receipt of income tax returns by the IRS. Over

II. Specification and Estimation.

This Section begins by specifying the Euler equation and a generalization thereof which is used in testing for excess sensitivity. It then describes the timing and other issues involved in the distinct task of estimating under a simple alternative hypothesis the marginal propensity to consume [mpc] out of refunds.

1. The Euler Equation and its Generalization.

As explained in the next Section, the consumption-periods in the CEX are quarters, beginning in different months of the year for different households. Following Zeldes [1989] and Lusardi [1992, 1993], equation (1) approximates the Euler equation that under the LC/PI theory governs household i's change in consumption from one quarter t to the next, t+1:

$$(1) \qquad dC_{i,t+1} = b_0 + d_{iy} + b_1{*}age_{it} + b_2{*}d(adults)_{i,t+1} + b_3{*}d(kids)_{i,t+1} + u_{i,t+1},$$

where d_y represents (the coefficient on) a year-dummy for the year y in which the household's quarter t begins, age_t is the age of the head of the household at t, and $d(adults)_{t+1}$ and $d(kids)_{t+1}$ record the change between t+1 and t in the number of adults and children in the household, respectively. (Adults are those 16 years old and over.) These

---

60% of returns due refunds are filed before April, mostly in February and March, with many of the rest filed in April. Applying the rule-of-thumb that most returns are processed within 6-12 weeks, this distribution also indicates that most refunds are received between March and May.

For more on the timing of disbursements, see Wilcox [1990]. In the few years in which the IRS had a slow start in sending out refunds, it caught up by mid-year. Eg, in 1985 computer problems delayed the processing of refunds in February and March, but the IRS caught up by May. Similarly in 1994, attempts to curtail fraud slowed refunds in the winter and spring, but again the IRS caught up by summer.

demographics allow for the most basic changes in household preferences, namely with family size and over time.[26,27]

In the spirit of Campbell and Mankiw [1990], the natural alternative hypothesis in which to embed the Euler equation is that households consume a fraction b^+_4 of their refunds on receipt. That is, b^+_4 is the instantaneous mpc out of refunds.[28] The term b^+_4 $*d(\text{refund})_{i,t+1}$ would then be added to equation (1).[29]

---

[26]Carroll and Summers [1991] document that consumption profiles are hump-shaped in age for all sorts of households. Attanasio and Weber [1992, 1993] and Attanasio [1993] argue that family composition might account for much of this shape.

Jappelli, Pischke, and Souleles [1994] suggest that many other demographics, even if time-varying, might have less to do with the unconstrained Euler equation, i.e. with preferences, than with liquidity constraints, and so they are not added to equation (1). More generally Deaton [1992] makes the reductio that one can always add enough covariates to kill the significance of the excess sensitivity coefficient.

[27]The absence of a discount factor and (real, after-tax) interest rate is usually justified by one of two arguments. Following Lusardi (and many others) the discount factor can be assumed to be the reciprocal of the interest rate, and so cancel it out. Alternatively, like Attanasio and Weber [1992, 1994] one can use the municipal interest rate, which is tax-free and so the same across households, and is therefore captured by the time dummies. (For quarterly changes in the years of interest, any uncertainty in the inflation rate is hopefully safe to ignore.) If the discount rate varies across households, it must be assumed uncorrelated with the regressors. Both arguments make rather strong assumptions about preferences, but these assumptions seem unlikely to be primarily responsible for the particular pattern and magnitudes of results found below. (And again, if they are they would themselves be of economic interest.)

[28]A more realistic alternative including lagged and seasonal responses is described in the next subsection. Given the aggregate distribution of refund disbursements, however, this instantaneous mpc coupled with indicators for summer-time trips (below) should capture most of such responses anyway.

[29]The (predictable) quarterly change in other income might be added as well. However, this is not available in the CEX. Anyway, it is not clear that refunds are correlated with the high-frequency change in income within the fiscal year after the refund has been determined.

There might though be some correlation for households whose income is regularly seasonal, and so correlated even with predetermined income. This does not however seem to be a particularly salient possibility, and the most likely such households, in farming, are dropped. Alternatively, refunds might contemporaneously affect income; eg, due to

As the next Section details, however, household refunds are not reported quarterly in the CEX--only annually--so instead the level of refunds in year y, $refund_y$, is used in testing for excess sensitivity. The generalization of the Euler equation that is estimated is then

$$\text{(2q)} \quad dC_{iq} = b_{0q} + d_{iq} + b_{1q}{*}age_{iq} + b_{2q}{*}d(adults)_{iq} + b_{3q}{*}d(kids)_{iq} + b_{4q}{*}refund_{iy} + u_{iq},$$

where the equations for the 12 possible calendar quarters t are indexed by **q, the number of the first month in quarter t**, 1-12.[30] So, if q is month number z, the dependent variable is the household's combined consumption in months z+3, z+4, and z+5 minus its combined consumption in months z, z+1, z+2. For q=1 for instance, this is $C_{Apr,May,Jun}$ - $C_{Jan,Feb,Mar}$, January being month 1.[31]

As explained shortly, for q = 1-9 the quarters t and t+1 come entirely after the year in which $refund_y$ was determined. Therefore, under the LC/PI models b_{4q} is zero for these q's.[32] That is, the predetermined refund should not help predict the change in

---

overtime to meet the consumption-response to refunds. The time dummies should pretty much handle this seasonal effect, and anyway it presumes a response to refunds to begin with. Either way, controlling for industry and occupation in the main regressions below did not qualitatively change the responses found.

[30]The relationship between t and q can be written as $t_q + 1 = t_{q+3}$. To keep the notation in the text simple, the subscripts will be omitted.

[31]Shea [1993] and Lusardi [1992, 1993] also face difficulties due to the mismatch in the timing of their consumption and income variables, in the PSID and the CEX respectively. The mismatch forces Shea to use the wage as his regressor, not income. This might in part account for the large magnitudes of his excess sensitivity coefficients.

[32]Strictly speaking, the LC/PI theory allows, as for any windfall, a small reaction to any unexpected portion of a refund, which should be smoothed over the remainder of one's life. As argued above, this portion is likely relatively small.

Adding the interaction of age and refunds to the regressions below led to significant multicollinearity. On bracketing the age groups into about 10-year intervals, it was often the young with the larger response, opposite to the theory. Liquidity constraints or the impatience of youth are possible explanations.

consumption. This is the central null hypothesis this paper tests. There is excess sensitivity if any of these estimated b_{4q}'s differs significantly from zero. In addition to testing for excess sensitivity, however, this paper tries to estimate the structural mpc b^+_4 under the alternative hypothesis. The next subsection illustrates the relationship between b^+_4 and the estimated b_{4q}'s.

2. The Pattern of Estimates under the Alternative Hypothesis.

Equation (2q) is estimated independently for each starting month q, yielding a different b_{4q} for each q. Figures 1a and 1b illustrate the pattern of estimates that might be expected under the alternative hypothesis. As a gross approximation to the actual distribution of refund-receipt, suppose each of three people gets a \$1 refund; one in March, one in April, and the third in May.

Figure 1a displays the coefficients that would be estimated (times the population size 3) if each person immediately consumes the \$1, i.e. if the mpc $b^+_4 = 1$. The complete \$3 effect would be found for q=12, in $dC_{12} = C_{Mar\text{-}May} - C_{Dec\text{-}Feb}$ (December being the 12th month). This is because each person has increased her $C_{Mar,Apr,May}$ by \$1 (the first in March, the second in April, and the third in May), but $C_{Dec\text{-}Feb}$ is unaffected, so $dC_{12} = 3\text{-}0 = 3$. However, for this and "earlier" periods q's (i.e., for q=10-12), the trailing quarter t extends back into the previous year, in which the refund has not been predetermined. E.g., for q=12, innovations to income in December could correlate the refund received in the following year with consumption in December, and so with $C_{Dec\text{-}Feb}$. Consequently, equation (2) will be estimated only for periods q = 1-9.

For period q=1 the observed $dC_1 = C_{Apr\text{-}Jun} - C_{Jan\text{-}Mar}$ would be 2-1 = 1. The \$1 spent by the first person in March comes into the negative, trailing term $- C_{Jan\text{-}Mar}$, and so partly offsets the \$2 spent by the second and third persons in April and May. The net result is an observed effect of \$1. Normalizing by the number of people, the estimated $b_{4,1}$ would be 1/3, a third of the actual mpc b^+_4 of 1. For larger q's, as more of the spending moves into the trailing term $-C_t$, the estimated coefficients become negative. E.g. $dC_2 = C_{May\text{-}Jul} - C_{Feb\text{-}Apr} = 1\text{-}2 = \text{-}1$. The coefficients then die off as the trailing quarter begins after May, after which there is no more spending.

More generally, since for q=1 the leading quarter t+1 (April to June) covers the time of refund-receipt for most households, the estimated $b_{4,1}$ should be positive and would primarily reflect the immediate response to refunds, net of the attenuation (due mostly to refunds in March) just described. The next Section shows how the data can be arranged better to match the employed regressor, $refund_y$, with the desired regressor under the alternative, $d(refund)_{t+1}$. This matching diminishes the attenuation under the alternative in the estimated $b_{4,1}$, giving a better lower bound for the actual mpc b^{+}_{4}.

For larger q's, after most households have already received their refunds, any estimated positive coefficient is probably largely capturing lagged household responses, not just the immediate response of those few households receiving their refunds late. Recall the strong negative effect shown in Figure 1a for q = 2-5, due to the passing of the "spring effect" into the trailing quarter t. For a lagged response to overcome this effect and be detected, it would probably have to be large and synchronized across households. Consider the example in Figure 1b. If two of the three people save a good part of their refunds for late summer, say to spend on vacation, a second hump in the estimated coefficients would appear as the "summer effect" comes into and leaves quarters t+1 and t.[33] Section III describes the timing of the refund variable in greater detail.

---

[33]More generally, the aggregate distribution of the timing of receipt can be used to back out the complete response to refunds. When only the response in the quarter of receipt is important, the attenuation factor for dC_{t+1} is $p_{t+1} - p_t$, where p_t is the probability the refund is received in quarter t. Using Table 2, for q=1 this is about .52 - .38, which is about 1/7. This might overstate the attenuation, however, because a good share of refunds in January to March come in late March, and some recipients will delay before cashing and consuming their checks.

The attenuation allowing for both lags and seasonality can be derived using a variation of the standard errors-in-variables argument. Here, the observed regressor $refund_y$ is some constant r, whereas the desired regressor $d(refund)_{t+1}$ is r with probability p_{t+1}, -r with probability p_t, and zero otherwise. The attenuation factor turns out to be a convolution of these probabilities $\{p_s\}$ and the lagged/seasonal consumption-responses, which can be unravelled using the estimated b_{4q}'s. Work on this is underway. Assuming the immediate response is larger than the lagged responses, since refund-receipt is concentrated in the spring, then using the instantaneous mpc and indicators for summer-time trips should capture most of the general response anyway.

3. Estimation.

For a given starting month q, equation (2q) is estimated by OLS, pooling together observations from different years of the same calendar quarters t and t+1. The standard errors are corrected for heteroscedasticity.[34]

Under the null hypothesis refunds do not belong in the Euler equation at all, so their correlation with other variables is irrelevant. Under the alternative hypothesis, the main orthogonality assumption is that any unmodelled individual effects that are correlated with refund-receipt come additively into the levels of consumption, and so drop out on taking differences.

The case for orthogonality under the alternative is bolstered by exploiting three different sorts of cross-sectional variation in refund-receipt. First, when the sample is limited to those getting a refund, even if there is something special about those receiving refunds that survives differencing, so long as it is independent of the magnitude of refunds it is captured by the constants and so does not take out orthogonality. Second, those not receiving refunds are sometimes added as a **control group**, with a refund-regressor of zero. The resulting response of consumption to refunds is then net of the consumption of non-recipients with the same age and family changes. Like the time-dummies, this also helps control for seasonality. Third, a dummy for those known to have received their refunds later than others, interacted with the size of the refund, will later be added to the regression. This "differences-in-differences"-like coefficient captures variation due solely to the timing of receipt, conditional on a receipt, and thereby controls for anything special about the magnitude of the refund as well.[35]

---

[34]It is assumed that any common shocks within a year that correlate the change in consumption across households are mostly captured by that year's dummies. Allowing different constants and time-dummies across q controls for seasonality.

[35]Slemrod et al. [1994] justify the orthogonality of such variation in timing. Using data on individual tax returns from the IRS, they found the effect of a refund's magnitude on the filing date was rather small (1.6 days earlier per $1000), even smaller when not conditioning on the receipt of a refund (0.22 days). Using those not receiving refunds as a control helps to investigate even these small effects.

III. The Data.

This Section discusses the variables and sample used in this paper, ending with a detailed description of the refund variables. The Appendix provides further details.

The data are drawn from the Consumer Expenditure [CEX] Surveys for 1980 to 1991. The CEX is a rotating panel. "Consumer units" are interviewed four times, three months apart, but starting in different months. Of the roughly 5,000 units in the survey at any given time, some will have their first interview in say February, with the remaining interviews in May, August, and November. For others the quarterly interviews might be in e.g. April, July, October, and January. Etc.

1. Financial and Expenditure Variables.

The income variables, including income-tax refunds, are generally recorded in interviews 1 and 4, for the 12 months up to the start of the month of the interview. However, a good share of the consumer units come into the survey as replacements for those that have dropped out early.[36] The refunds of such replacements are recorded in the interview in which they come into the sample, whatever its number, as well as in interview 4. As regards wealth, one can calculate a consumer unit's liquid wealth at the start of the quarter covered in interview 1.

Expenditures are recorded in each interview (1 to 4) for the preceding three months, the reference period. At the 6-digit universal classification code [ucc] level, these records are rather detailed, including e.g. expenditures on wigs or on the cleaning of septic tanks.[37] Expenditures are assigned to particular months within the reference period. For most ucc's the month is the actual month of purchase. But for ucc's

---

I am also trying to exploit the clearly exogenous timing-variation due to delays in disbursements by the IRS in the mid-1980s. Unfortunately, the sample-sizes seem to be too small.

[36]There is no obvious way to correct for any resulting sample-bias.

[37]For a humorous account of responding at this level of detail, see Goldman [1993].

comprising about a quarter of total expenditure (including much of food), the total expenditure within the reference quarter is allocated evenly across the three months; and for other ucc's expenditures are arbitrarily assigned to one of the months. This paper accordingly treats the data as quarterly.

For each consumer-unit/reference-quarter, expenditures at the ucc-level were deflated by the corresponding CPI sub-indices (1982-4 $)[38], and then aggregated to create the following four consumption-groups used below. **Total consumption [TC]** basically replicates the total expenditure calculated by the CEX, less a few inappropriate ucc's (like "spending" on pensions). Although the Euler equation (1) does not apply to durables, under the alternative hypothesis the response of total consumption to refunds is of great interest. A subset of TC, **nondurable consumption** [ND] generally corresponds to the CPI's nondurables and services. The main exceptions are that certain ucc's that are treated as durable under the NIPA accounting have also been removed, as has spending on health and education for which the link between expenditure and consumption is relatively loose.[39] The mean of consumer units' nondurable to total consumption (before cutting the sample as below) is 0.80. The smallest grouping is "food" (including alcohol, to correspond with the PSID), whose average share of total consumption is 0.26.

The ND grouping includes many relatively durable and lumpy expenditures to which the Euler equation might not strictly apply.[40] In response, following Lusardi [1992, 1993] a subset of nondurables, "**strictly nondurables**" [**SND**s], has been created.[41] E.g. durables and semidurables like wigs and clothing have been dropped.

---

[38]One of the chief purposes of the CEX is the calculation of the weights for the CPI, so the match between ucc's and CPI subindices is close.

[39]Eg, many of the expenditures on health are negative, reflecting reimbursement for past payments.

[40]Such durability could lead to false findings of excess sensitivity. See e.g. Mankiw [1982].

[41]It appears that Lusardi used the already partially-aggregated expenditure groupings available in the CEX, as opposed to the primitive ucc's. Her delineation of SND's is consequently coarser than that here.

So have infrequent expenditures like repairs and septic tank cleaning. Although unavoidably discretionary, this grouping has been guided by the principle that expenditures least likely to represent consumption "smoothable" at high frequency be dropped. The major components of SNDs are food; household operations, including monthly utilities and small-scale rentals (e.g. videos); apparel services and rentals; transportation fuel and services, including public transport; personal services; and entertainment services and high-frequency fees. The mean of SND to total consumption turned out to be 0.53.[42]

2. The Sample.

Sample-selection was generally guided by the goal that a household's reported expenditure tolerably reflect its actual consumption, especially of (strictly) nondurables. A consumer unit is dropped from the sample if there are other consumer units in the same household. The observations are then household-quarters. A household-quarter is dropped if the household lives in student housing, or if the head of household is a farmer; or if the household lacks (non-trip) food-expenditure for any month of the quarter, or if any food was received as pay in the quarter. (The latter food might not have otherwise been voluntarily purchased.) All quarters for a given household are dropped if it paid for board at school or at a boarding house in any quarter, because such expenditures are often lumpy. Further, households are dropped if their head is less than 24 years old, or over 64. (The former might still be dependent on their parents, and might not yet have much experience with withholding. The latter generally do not have much income withheld and might have unusual consumption, e.g. large health expenditures.) Counting only those household-quarters whose refund variable is used below, this leaves 43,492 observations

[42]In addition to being in these groupings, expenditure on food, SNDs and NDs while on trips were also kept track of separately, for the purposes of the analysis below of trips. The CEX defines trips as any overnight trips, and all-day trips travelling over 75 miles (but not commutes).

with positive refunds for estimating the Euler equations, spread across the 12 periods q. (There are 4,160 such observations for period q=1.)

3. Refunds and their Timing.

The refund variables cover the twelve months preceding the interview-month in which they are recorded. As a result, the exact month in which any refund is received is unknown. Nonetheless, the timing can be pinned down sufficiently to get at the structural mpc b^*_4 out of refunds. This analysis is somewhat tedious, but necessary to take full advantage of the information available in the data.

Figure 2 displays the timing of the consumption and income reference-periods for the period of greatest interest, q=1 in the spring. The timing depends on whether the trailing quarter t, covering January to March, corresponds to interview 1, 2 or 3 (n = 1-3, respectively). Depending on q and n, the regressor used is either the refund reported in interview 4, **r4**, or that in interview 1, **r1**. Two guidelines govern this choice. First, the time-span covered by a refund variable may not extend into the calendar year following the consumption quarters as issue, so any refund remains predetermined. Second, the span never begins before July of the previous year. This ensures that the twelve-month variable always covers refunds received in January to June of the year of consumption, which is when about 90% of refunds are received.

Note that even under these guidelines a refund might have been received before the consumption-quarter t, or during that quarter, or after consumption-quarter t+1 (though never in the following year). For the task of estimating the structural mpc, these three possibilities each leads to attenuation: for the utilized regressor refund_y in equation (2) to match the desired regressor $d(\text{refund})_{t+1}$, any refund recorded in the former must have been received during quarter t+1. It is important to remember, however, that these possibilities do not upset the informational assumptions of the excess sensitivity test, since in any case the regressor is still predetermined.

As discussed in Section II, the extent of attenuation under the alternative hypothesis can be gauged by means of the aggregate distribution of refund-receipt. For q=1 two additional means are employed, which take advantage of the second refund report, the one that is not put in the regressor. First, one can control for the first two sources of attenuation by identifying some of the refunds that surely came after consumption-quarter t. These will be called "**late**" refunds. For instance, referring to Figure 2, one can use the refunds r1 (or refunds r2 and r3 for replacement households), to constrain the time-span of the regressor r4 to April and later. I.e., for n=1, if r1 shows no refund, then any refund recorded in r4 must have come after the period of overlap of r1 and r4, and so after March. Similarly if r2=0 for n=2, or r3=0 for n=3. A dummy variable for such late refunds, interacted with the refund-regressor, will be added to equation (2). This interaction will capture the added effect of refunds conditional on knowing they come at least since quarter t+1. Under the alternative, this effect should be positive. This effect resembles a "differences-in-differences" effect, which as already explained controls for most anything special about refund-recipients.[43] Note, however, that nothing can be done about the attenuation due to the refunds that come after quarter t+1, which are about 10% of refunds (Table 2).

Unfortunately the refunds that are identifiably late are relatively few (about 10% of refunds). In response a second, less complete correction for the attenuation is employed for most of the results. It controls for only the first source of attenuation above, by restricting the sample to "**current**" refunds. For instance, refer again to Figure 2. For n=1, the span of r4 is automatically constrained to the current year. For n=2, if r1 shows no refund, then any refund recorded by r4 must also have been received since

---

[43]This assumes of course that those identifiable as the recipients of late refunds are representative. This assumption requires further scrutiny.

I have also experimented some with comparing r4's with r1's (or r2's/r3's) when the latter are non-zero. When the same magnitude, it might be that they refer to the same refund. This would pin the refund down to the time of their overlap. However, since people probably round their reports and memories fade, even reports actually referring to the same refund check might differ slightly. Accordingly I assumed that if the magnitudes r4 and r1 differ by less than $100, they refer to the same refund. The results were hard to interpret, casting doubt on this assumption.

January. Similarly for n=3 if r2=0. To retain observations[44], for n=3 also included were observations with r1=0, for which r4 is constrained to October of the previous year and later. Restricting the sample to current refunds eliminates the refunds most likely to have been received in the previous year, and so under the alternative should increase the estimated coefficients. This correction will be sufficient to investigate the source of the excess sensitivity found below, and to give some rough idea as to its magnitude, despite the remaining attenuation due to refunds received after the leading quarter t+1 and, more importantly, during the trailing quarter t when consumption comes in with negative sign.

Finally, consider the quality of the refund variables used under the guidelines above. As for other income variables, there appear to be too many reports of no refund: of the observations with a valid refund variable (gross of the sample restrictions described above), only about 45% report a positive refund (46% for q=1), a good deal fewer than the approximately 3/4 of returns in the IRS data that receive refunds. One quality-check is to compare these variables to the CEX variables for federal income tax, selected under the same guidelines. 17% of the valid such tax reports are positive, giving a ratio of tax payments to refunds of about 0.38, which is close to that from the IRS data.[45] Another check is that the average of refunds to after-tax income is about 10% (the median, about 3%), roughly comparable to the IRS data.[46] A final check compares the average of the CEX refunds to the national averages in Table 1. The former are somewhat greater. (The

---

[44]The r2's and r3's are many fewer than the r1's, because they are available only for the replacement households.

[45]Eg, in FY 1989 the number of returns with taxes due was about 28m, and the number of returns with refunds in FY 1988 (and so received in 1989) was about 76m, giving a ratio of 0.37.

I have also begun to explore the response of consumption to these tax payments. The informational requirements here are harder to guarantee, however, because the CEX variable on income taxes does not distinguish April taxes based on the previous fiscal year, from estimated taxes for the current year.

[46]For FY 1991 the average refund divided by the average AGI for refund-recipients (using the mid-points of the AGI brackets) is about 4%; for FY 1985, it is about 5%. Dividing instead by after-tax income would increase these figures.

median is also greater than Wilcox's estimate of almost $500.) Reassuringly, however, the CEX averages display the same pattern across years as in the aggregate, peaking in the middle 1980's.

The excess of zero reports, and the large means conditional on a nonzero report, might in large part reflect people's rounding their responses to the Survey, especially rounding to zero (or forgetting) small refunds.[47] These quality-checks suggest that, conditional on a report of a refund, measurement error does not appear to be unacceptably large or systematic.

IV. Results.

This Section begins by estimating equation (2q) for periods q = 1-9. It finds significant excess sensitivity. Attention then turns to investigating the source of this sensitivity, and to the distinct task of estimating the magnitude of consumption's response to refunds under the structural alternative hypothesis. The focus is primarily on the spring, period q=1, when most refunds are received. For this period, the sample is split to examine the role of liquidity constraints. Also, the refund regressor is refined to correspond better to the desired regressor under the alternative, giving a more complete view of the response to refunds. The linearity of this response is also examined. The Section concludes by investigating the delayed effect of refunds, in the summer.

1. Excess Sensitivity Tests, all periods q.

Table 3 records the salient coefficients estimated for equation (2) for periods q = 1-9.[48] Consider first q=1, for $C_{Apr,May,Jun}$ - $C_{Jan,Feb,Mar}$. For expenditure on food, the

---

[47]A histogram of the reported (nominal) refunds reveals that many households round their report to a multiple of $100, but otherwise the distribution seems reasonable.

[48]The changes in family size, as reported for q=1 and unreported for other q's, generally have positive and large effect, often significant. The coefficients on age are of varying sign, but not large and more often insignificant.

coefficient b_4 on refunds is positive, but small at about .01 and insignificant. The coefficient on strictly non-durables (**SNDs**) is unsurprisingly larger at .02, and significant at the 10% level. It is significant at the 5% level on including the **control group**, i.e. on adding to the regression with a refund-regressor of zero those who did not receive any refund. Strictly nondurable consumption is, therefore, excessively sensitive to tax refunds--counter to the LC/PI theory. For non-durables (**NDs**) the coefficient is slightly larger, but insignificant. Since the coefficient on total consumption (**TC**) is much larger, at .16, most of the response to refunds seems to come in durables. This coefficient is significant at the 10% level (actually, its p-value is .056), at 5% with the control group.

Turning to the later periods, there is also significant excess sensitivity for periods q = 2, 3, and 7. Further, the pattern of sensitivity is roughly that in Figure 1b. For q=2, the coefficients that are significant are negative. This is consistent with the passing of consumption's immediate, spring-time response into the negative, trailing term - $C_{Feb\text{-}Apr}$. For q=3 the coefficient is again positive, and for q=7 it is back to negative--consistent with the use of a good part of refunds in the late summer. Even food sometimes responds significantly in these periods.

TC by contrast is not significant after q=1, nor are NDs at q=7, even though the other categories are subsets of them. The inclusion in NDs and TC of durables and semi-durables, which are irregular, might be swamping the effects of the strictly non-durables.[49] The differences in the overall pattern of response of durables and nondurables across q are noteworthy. The response of durables comes mostly immediately, in the spring, when it is sharp. The response of nondurables is prolonged. This delayed reaction is unlikely to be due to liquidity constraints.

Adding the control group does not generally change the coefficients substantially. This bolsters the above results, in particular suggesting that they are not simply artifacts

[49]Lusardi [1993] finds similar behavior for the larger categories of goods.

of seasonality. It also suggests that refund-recipients are not somehow "special", and thereby supports the generalization of the results to other types of income-receipt.[50]

These results constitute a statistically significant rejection of the canonical LC/PI theory. The estimated magnitudes are economically significant, as well. Even without any correction for the attenuation under the alternative hypothesis due to the timing of refunds, the effect on total consumption is over 15% of a refund within the quarter of its receipt. This is already greater than the base estimates of Modigliani and Steindel [1977], and comparable to those of Blinder [1981].[51] (It is important to remember that the timing-attenuation does not affect the test of excess sensitivity, because the refunds used in the regressions remain predetermined.)

To investigate one component of the attenuation, the regressions above were run for q=1 for only the **current** refunds, i.e. dropping the refunds that are most likely to have been received in the previous calendar year. As expected the coefficients on TC do increase, to about .22 (.10), significant at the 5% level.[52] Adding the control again changes the magnitudes rather little.

The point is sometimes made that spending on durables is a form of saving, so that these results are not as counter to the LC/PI theory as they might seem. However, the response of nondurable consumption is itself sufficient for rejection of the theory, and the timing of the expenditure on durables remains unexplained. That is, in the absence of liquidity constraints, why do these households wait for their (predictable) refund check

---

[50]Though recall from Section III that some in the control group with no refund reported must in fact have received refunds. Nonetheless, the control still includes the many who did not in fact receive refunds.

[51]It is not clear how to compare these results to Wilcox [1990]'s, because he appears to report only the response on the day of refund-receipt.

Note that the regressions have not yet been weighted for the sample used, so the comparison with the aggregate studies remains somewhat loose.

[52]The other consumption groupings retain positive coefficients, but they are rather smaller and less significant than before, presumably due to the decrease in observations. E.g., for SNDs the coefficients are .012 (.011), and .016 (.012) with the control.

to buy the durables?[53] In any case, for the purposes of fiscal policy the effect of receiving the check is still expansionary.

The rest of this paper concentrates on periods 1 and 3, with the aim of uncovering the source of the excess sensitivity just reported, and of estimating the full response of consumption under the alternative.

2. Sample Splitting, period q=1.

If liquidity constraints are important, the response of the constrained ought to be especially sharp in the spring, the time of most refund-receipt. Following Zeldes [1989] and others, splits with respect to various financial variables will be used to contrast the behavior of those most likely to be liquidity constrained with that of those least likely to be. Although none of these splits quite matches the theoretical definition of liquidity constraints, both common sense and observation suggest they are closely related to it, and the mismatch does not seem responsible for the intuitive pattern of results found below.[54] The focus will be on the response of strictly nondurables (SNDs) and total consumption (TC). The coefficients on the demographics are not remarkable, so are not reported.

The natural starting point is to look at those with no liquid wealth. However, due to measurement error and the fact that household wealth at the exact time of refund-receipt is unknown, positive cutoffs are chosen. Jappelli [1990] and others have suggested that between 12-20% of U.S. households are constrained. Accordingly, for the wealth-

---

[53]For a standard (S,s) model to explain such a response, it seems that its state variables would have to be a function of liquid assets. But the main motivation for such state variables is liquidity constraints.

[54]Jappelli, Pischke, and Souleles [1994] argue that such splitting can potentially yield spurious findings of sensitivity. It is not clear that it could explain the particular pattern of results here. The next draft of this paper is slated to follow them in using the direct information on borrowing constraints in the Survey of Consumer Finances to identify the constrained. In the meantime it is enough to have a positive, albeit imperfect, correlation between this direct measure of liquidity constraints and financial status.

As explained in the appendix, for each financial variable the earlier of its two reports in the CEX is used in splitting, to minimize any endogeneity. In any case, the refunds are predetermined.

split and other splits, to bracket these numbers it was decided a priori to focus on those below the 10th and 25th percentiles.[55] For wealth the cutoff at 25% was less interesting, so Table 4 begins with the "tight" split, contrasting the results from separate regressions for those below the 10th percentile of real liquid wealth ($100)--the likely constrained--with the results for those above the 75th percentile ($7500)--the likely unconstrained.[56]

For the constrained, SNDs react by about .05, which though insignificant is greater than for the full sample of refunds, above. (As the note in the Table shows, their NDs react by about .11, which is significant.) The unconstrained, by contrast, appear to respond primarily to durables, since only the coefficient on TC is much greater than zero or significant. It is rather large, above .20.

To correct again for part of the timing-attenuation, consider the limitation of the sample to current refunds, i.e. the refunds least likely to be from the previous year. The magnitudes increase quite a bit, as expected under the alternative. The response of the SNDs of the constrained is now significant and large at .08. Even their food responds significantly, at about .05. The TC of the unconstrained now jumps by well over 25% of the refund. These increases in the coefficients illustrate the importance of the timing-attenuation, and suggest that the full response of consumption might be substantially larger

---

[55]The exact cutoffs depend on whether the control group is included or not. For simplicity, a single round number near both cutoffs, with and without the controls, is used. Sometimes tighter and looser cuts were investigated, trading off the likelihood of being constrained with sample size.

Using wealth splits at zero or less led to small sample sizes and somewhat erratic results. The appendix explains that many of these non-positive wealth reports are mismeasured.

[56]Ignoring the middle of the distribution throws away a lot of information. But contrasting just the ends of the distribution makes it more likely that what is being identified is the effect of liquidity constraints. (In the middle there might be many of both constrained and unconstrained households.) The technique of Jappelli, Pischke, and Souleles [1994], by contrast, allows for a continuous probability of being constrained, not just 0 or 1.

than so far estimated. The standard errors are not small enough to make these different responses of the constrained and unconstrained significantly different, however.[57] [58]

To investigate the robustness of these results, consider now splits based on other financial variables.[59] One might also expect large consumption-responses from those with low income. Whatever the future path of their income, casual observation suggests that many low-income households live hand-to-mouth. Table 4 reports interesting results for even the "loose" split, at the 25th and 50th percentiles of earnings (gross of deductions).

Compared to the split on wealth (and without any correction for attenuation), for the constrained the coefficients on food and SNDs are now larger and significant, at about .03 and .08. For the unconstrained, the response of TC is smaller though still large and significant at the 10% level.[60] The most notable difference between the splits comes on restricting the refunds to current refunds. The TC of the constrained is for the first time significant and quite large at about .30, both with and without the control. Also, the

---

[57]The comparisons in the Table of the constrained and unconstrained come from pooled regressions of only the constrained and unconstrained. As argued above, the demographics in the regression are not intended to reflect liquidity constraints, so for the reported comparisons the coefficients (unreported) on the demographics were constrained to be the same for the two groups.

[58]The results using the control group generally do not much change. The main exception is TC for the constrained, which however is insignificant with or without the control. In addition to the irregularity of durables, the small sample size might be responsible.

[59]Another motivation is that the income variables in the CEX are generally considered to be superior in quality to the wealth variables, and so will receive more attention in this Section.

[60]The looser split accounts for much of the smaller magnitude, given the nonlinearity of response (below).

SND-responses of the constrained and unconstrained are at last significantly different from each other.[61]

These two splits identify different households as constrained: some have great wealth but low income, others low wealth but great income.[62] The intuition of Zeldes [1989] is that it is the latter who are more likely to be constrained: they might be only currently illiquid, lacking access to the large NPV of their future income. The sample was accordingly also split on the basis of the wealth to earnings ratio. The results splitting at the 10th percentile are stronger, and so are the ones reported in Table 4.

The point estimates are generally like those using the income-split. The main differences are, first, the nondurable response of the constrained increases somewhat with the control. Second, TC is no longer significant for the constrained even for the current refunds. Again this appears to be at least partly due to the irregularity of durables, because (as recorded in the Table's notes) the response of NDs for the constrained is now significant and even larger than under the income-split. Third, and less easily explained, the coefficient on food for the unconstrained is significantly negative (at the 10% level) for the current refunds. However, its magnitude is quite small at -.008. In sum, insofar as a large response of nondurables is in fact characteristic of the constrained, and a

---

[61]Splitting instead at the 10th percentile, the magnitudes for the constrained are generally larger, at the cost of observations.

The results using after-tax income are generally comparable. The main exception is that the TC of the constrained is rather smaller. E.g. splitting at the 25th percentile, for current refunds the coefficient is still large at about .14, though not significant.

Cross-tabs reveal that for either measure of income, about 5-10% of those labelled as constrained were counted as unconstrained by the other. Eg, some with small earnings have large unearned income. They are unlikely to be constrained. On the other hand, it is not clear that someone with large earnings but small reported after-tax income due to say paper losses is constrained either. Anyway, one reason to prefer the earnings variable is that it might be less noisy, being primitive.

[62]For the different wealth and income measures considered, the two "off-diagonal" cells in the cross-tabs of the constrained/ unconstrained using wealth versus the constrained/ unconstrained using income held between 5 to 20% of the total observations, each, depending on the tightness of the splits.

primary response in durables characteristic of the unconstrained, then these results tend to support Zeldes's intuition.[63]

To summarize this subsection, there is significant and interpretable evidence both of liquidity constraints and of "splurging" on durables by the likely unconstrained. First, for the constrained, nondurable consumption increases significantly in this period of refund-receipt. The immediate response of their strictly non-durables is already on the order of 6-11% of their refunds, despite the attenuation. Even their spending on food jumps by as much as 5%, supporting the supposition that they are in fact more likely constrained. Their spending on durables is also possibly quite large, though not estimated with consistency. Second, as regards the unconstrained, their strictly non-durables do not appear much to respond. But even these households spend at least 17-30% of their refunds in the spring, and so are responding primarily and sharply in durables. This response is puzzling, because these liquid households could have bought their durables before receiving their refunds. Since these households represent a good part of the sample (usually 50%), and even more of total spending, these results suggest that the aggregate effect of refunds is substantial. The magnitudes increased quite a bit on correcting for one source of attenuation, due to refunds possibly received in the previous year. This suggests that the full responses might be appreciably larger.

---

[63]On limiting the unconstrained to the top 10% of the wealth/earnings distribution, the differences between the constrained and unconstrained become even more striking. E.g., even for the full sample of refunds the response of the unconstrained's TC is quite large: .35(.04)** and .34(.07)**++ with the control. (The ++'s mark a significant difference between the constrained and unconstrained, as explained in the Table.) And the response of their NDs is significantly smaller (at the 5% level) than the constrained's, both with and without the control. The next section analyzes the nonlinearity of the response more formally.

The results were qualitatively similar using after-tax income.

3. The Magnitude and Nonlinearity of Response, period q=1.

So far the attenuation from mistiming within the current year, notably due to refunds received in January to March and so accruing to the $-C_t$ term, has not been investigated. To this end a dummy variable for the **late** refunds--those known to have arrived after March--is added to equation (2), interacted with the refund regressor. Any remaining attenuation is then due to the late refunds that come too late: those coming after the leading quarter t+1, so after June. But these are only about 10% of refunds. Capturing variation solely in the timing of refund-receipt, this interaction also controls for most anything special about refund-recipients. However, these late refunds represent only about 10% of the refunds used for q=1, so the sample size might be of some concern.

Table 5 records the main results. For the entire sample of recipients, the interactions for late refunds are insignificant, and the "base" refund coefficients are similar to those before in Table 3. However, one lesson of the preceding subsection is that the behavior of the constrained and unconstrained might markedly differ. Table 5 begins by isolating those with low wealth. The interaction for food is significant (at the 10% level) and large, at 0.10. That is, the fraction of the refund spent on food by the constrained is 10% larger for those known to have received refunds after March, than for refund-recipients in general.

Consider again splitting the sample according to earnings. The interaction terms for SNDs are now much larger and significant, about .37 and .08 for the constrained and unconstrained respectively.[64] Adding the base refund coefficients yields the result that

[64]Note that the coefficients on late refunds are now positive for both the constrained and unconstrained, but small for the entire sample. This is not fully explained by differences in the coefficients for the other, non-refund regressors between the two groups; separate regressions for the middle of the distribution often give negative excess sensitivity coefficients. But neither these coefficients nor those for the base refund variable are significant.

Also, recall that the dummy for late refunds does not include all refunds received after March, but only the relatively small fraction of them that can be identified as coming then. (See Section III for more detail.) This could explain why the base refund coefficient is not always negative: some of the non-"late" refunds were actually received after March, and so push the base coefficient up in value.

those known to have received late refunds spend about 45% and 9% of their refunds on strictly nondurables, if constrained or unconstrained respectively.[65] For the constrained most of this response comes in food, which again suggests that in fact liquidity constraints are at play. Although the interaction terms for TC are negative, they have rather large standard errors and the corresponding coefficients on NDs are often positive.

Consider next the linearity of the response to refunds. The second part of Table 5 shows the results on adding to the regression the refund squared, normalized by earnings. For the whole sample, the quadratic term is insignificant, and the linear response remains close to the original response in Table 3.

So too on splitting the sample, for the constrained in earnings the quadratic terms are insignificant. For the unconstrained, by contrast, the quadratic term on food is negative and significant, as under the mental accounts theory. That is, unsurprisingly the fraction of the refund that goes to food diminishes with the relative size of the refund. However, for total consumption the quadratic coefficient for the unconstrained is positive, significant, and rather large at .42. The response of durables appears to grow with the relative size of the refunds: households seem to be using bigger checks to buy even bigger goods. Clearly the response cannot continue so to grow, so the result applies only for income-receipts of the magnitude in the sample. (But remember that refunds are often a relatively large portion of income.)[66]

---

[65]The reported results limit the sample to current refunds, to avoid artificially inflating the interaction term by adding refunds from the previous year to the base refunds. The results without this limitation are qualitatively similar.

I also ran regressions with only the late refunds. Again the results are qualitatively similar. E.g., for the late refunds that are also "current", the constrained's SNDs respond by .718(.250)**. Given the small sample size, only the pooled estimates are reported in the text.

[66]On again restricting the sample to current refunds, the curvature for TC is now significantly positive for the constrained as well, though smaller.

Also, without normalizing by earnings the quadratic term for TCs is still positive and significant.

To summarize, the estimated response of consumption increased substantially after correcting for an important additional part of the timing-attenuation, due to refunds received early in the year. The remaining attenuation, due to refunds received after June, might still be appreciable. Even so, the estimated response of the strictly non-durables of the constrained is as great as 45% of the refund; even for the unconstrained, it approaches 10%. The response is non-linear, apparently increasing for durables with the size of the refund, counter to the mental accounts theory.

4. Spending on Trips, period q=3.

To investigate further the excess sensitivity found in subsection 1 for the summer (q=3), a dummy variable was created identifying those households (about 60%) who spent any money on trips in the leading consumption-quarter t+1 (June-August).[67] To equation (2) were added separate intercept and slope terms for these households (i.e., the trip-dummy and the trip-dummy interacted with the refunds).

Table 6 reports the results. The interaction term for SNDs is positive and significant at the 5% level. At about .05, its magnitude is rather large. Note that the base refunds coefficients are not significant, suggesting that most of the excess sensitivity reported in Table 3 is in fact due to spending on trips.

On adding the control group though, the interactive coefficient drops to .03 and is no longer significant. Also the trip-intercept jumps up for SNDs and food, which together contain all of trip spending. (I.e., there are no trip-ucc's for durables.) This suggests, unsurprisingly, that spending on trips by non-refund-recipients is also important. To control for the tendency to take trips, in particular to see whether the significant coefficient without the control might in part be picking up an income-elasticity of

[67]The CEX has separate ucc's for spending on trips. The <u>average</u> expenditure for food on trips jumps about $10 when June comes into the quarterly reference period for consumption, another $20 when July comes in, and another $12 when August does. For non-food strictly non-durables, the jumps are about $15, $30, and $20, respectively.

The large magnitude of the trip-intercept below also shows the importance of trips to spending.

spending on trips, household earnings were also added to the regression. As the Table shows, this made little difference: the interaction term remains significant at about .05, and the earnings variable is insignificant.[68]

To summarize, the excess sensitivity found in the summer still appears to be largely related to the spending of refunds on trips. This sensitivity is not likely caused by liquidity constraints: the spending on trips takes place after most filers have already received their refunds.[69]

V. Conclusion.

This paper has found significant evidence of excess sensitivity in the response of households' nondurable consumption to their income tax refunds. Further, some of the likely sources of this sensitivity were identified. In particular, liquidity constraints appear to play an important role, because the nondurable consumption of those likely to be constrained increased substantially at the time of refund-receipt--by perhaps over 45% of the refund within a quarter--and far more than the consumption of the likely unconstrained. However, more than liquidity constraints are at play even at this time, because the consumption of those likely unconstrained also responded substantially, by perhaps over 30%, though primarily in durables. Also, the response of nondurables extends later into the year, which is also unlikely to be due to liquidity constraints. Much of this delayed response appears to stem from the spending of refunds on summer-time trips.

The response of durables, by contrast, is concentrated at the time of receipt. For the unconstrained this response is not easily explained by standard models of durables.

---

[68]The results with the control were as before with the control but without income, ie insignificant but positive.

[69]Unfortunately it is not possible to identify many trip-taking households who are sure to have received the refund "early", namely before the summer.

The basic puzzle is that the unconstrained could have bought their durables before receiving their refunds.

Having rejected the null hypothesis of the life-cycle/permanent-income theory, the paper undertook the distinct task of estimating the marginal propensity to consume out of refunds under a simple alternative hypothesis. The figures just cited for the consumption-responses came from correcting for most of the attenuation due to the uncertainty about the exact month in which a refund was received. These responses are already greater than found by most previous studies; the full responses could be appreciably larger. Since federal income tax refunds now amount to over $80 billion per year, these results imply rather substantial aggregate effects of refunds. More generally, the results evidence the potential efficacy of fiscal policy.

This paper also tested some implications of leading behavioral theories of saving. The increasing mpc out of refunds is counter to the theory of mental accounts. The sharp response of durables in the spring might be explained in part by models of self-control/forced saving. But the largest response comes from those with the greatest wealth and income, who could most easily undo the forced saving, simply by writing a check. Also, self-control might have difficulty explaining the delayed response after receipt.

This analysis might be extended in a number of ways. First, given the importance of the response of durables, it would be useful to disaggregate the response of different durables.[70] Second, the alternative hypothesis was rather simple; a more realistic lagged and seasonal response of consumption can explicitly be modelled. Third, in addition to financial variables, other characteristics of households that are informative about liquidity-constraints might be used in identifying those likely to be constrained. Finally, the behavior of households receiving refunds can be contrasted to that of households paying out taxes. Work on these extensions is underway.

---

[70]See eg Eberly [1994] and Attanasio [1994a].

VI. Appendix: the Data.

1. Consumption.

The ucc's are deflated by the nearest "A level" cpi index (e.g., food at home, gas and electricity, etc.). About 50 such groups were used, following the cpi categorization in Nelson [1992]. The few groups for which there is no cpi available, or whose timing in the CEX is off, are not discounted and hence do not enter any of the consumption groupings. The groupings are described in the text.

In aggregating the ucc-expenditures the following were kept track of: expenditures dated outside the reference period, topcodes, and imputations. Too large a fraction of the consumption groups turned out to have some component imputed to drop the imputations (even after 84:1, when imputations can be distinguished from other data-adjustments). If any component of a group was topcoded or missing its cost, the whole group was set to missing. If any component of a group was missing its date or dated before the reference period, that group was dropped for all interviews for the household at issue. A non-negligible number of ucc's are dated in the month of the interview, i.e. after the reference period. Following the recommendation of the staff at the BLS, for consistency such expenditures were accrued to the following reference period.[71]

2. Refunds.

The survey first asks, "During the past 12 months, did you (or any members of your CU [consumer unit]) receive any refunds from the following- ... a. Federal income

---

[71]The BLS does some checking that such expenditures are not re-reported in the following interview. My own investigation confirmed that many such ucc's were provably not double counted. By contrast, it does not appear possible to prove that two similar expenditures in successive interviews refer to the same purchase.

Note that mistakenly moving a past expenditure forward in time does not endanger the excess sensitivity test, since the expenditure was already known.

tax? [Section 22]."[72] If the answer is no, the refund variable, FEDRFNDX, is flagged as a "valid blank" (as opposed e.g. to "don't know" or a refusal). These form the control group in the paper, with the refund-regressor set to zero. If yes, the follow-up question is: "What was the total amount received by ALL CU members?" The refund variable used in the paper are these amounts in FEDRFNDX when flagged as "valid". They were never zero. The only topcodes came in 1980-81, when the rule for topcoding was that most income variables were topcoded if any was, and all of these were set to zero. Such refunds were changed to missing.

According to the BLS staff, there is no direct indication in the public files of whether a given household is a replacement. But they agreed that it is probably safe to take a refund r2 as newly asked in interview 2 if there is no record of interview 1 anywhere in the data set for this household. Similarly, r3 is used if there is neither an interview 1 nor 2. (r2's and r3's from the beginning of 1986 are not used, because of the change in I.D.'s. at that time.) Since these interviews are used only to help pinpoint the timing of the refund in interview 4, and since the only possible error is that the amount reported refers to an earlier interview, any error here does not endanger the excess sensitivity test.

The refund attributed to a given year is deflated by the average of the monthly cpi for all items averaged over March, April, and May of that year.

3. Other variables.

A household-quarter observation is dropped if any variable in the regression of concern is missing, and it is not grouped with either the constrained nor the unconstrained if the variable used to split is missing. Ages topcoded at 90 were retained.

Observations with income or wealth topcoded at the topcode were generally put in the corresponding unconstrained group. Topcoded variables already at zero in the data

[72]The directions in the survey, as well as some previous questions, explicitly note that the twelve-month time-span ends at the end of the month preceding the interview, not on the day of the interview.

were dropped. Wealth to income ratios were set to missing if the denominator was negative or if either the denominator or numerator was topcoded. Income variables that have earnings as a component, including earnings, were generally not used if earnings were zero. So too wealth from interview 4, and the wealth in interview 1 derived from it, are not generally used if the former is zero. Many of these might be false zero's, and the survey does not ask a separate leading question to distinguish e.g. a savings account with zero balance, and the lack of a savings account.

Real liquid wealth is the discounted sum of the real amounts in checking and savings, when both exist. The liquid wealth used in splitting is that corresponding to interview 1. The income variables used are those recorded in the trailing reference quarter, i.e. t rather than t+1. Some analysis was undertaken using only interviews for which the span of the income variables completely precedes the consumption reference periods, but this eliminated too many observations.

A handful of adjustments to the primitive data were made, according to corrections provided in the CEX documentation for various years. For married households with female respondents, the demographic variables specifically about the respondent were switched to refer to the male, identified as the head of household.

Bibliography

Altonji, J.G., and Siow, A., 1987, "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data," Quarterly Journal of Economics, 102, pp. 293-328.

Attanasio, O.P., 1993, "Personal Saving in the US," Stanford University, mimeo, December.

Attanasio, O.P., 1994, "The Intertemporal Allocation of Consumption: Theory and Evidence," NBER Working Paper, July.

Attanasio, O.P., 1994a, "Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules," Stanford University, mimeo, October.

Attanasio, O.P., and Weber, G., 1992, "Consumption Growth and Excess Sensitivity to Income: Evidence from US Micro Data," Stanford University, mimeo, August.

Attanasio, O.P., and Weber, G., 1994, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," NBER Working Paper, July.

Bird, R.C., and Bodkin, R.G., 1965, "The National Service Life-Insurance Dividend of 1950 and Consumption: A Further Test of the "Strict" Permanent-Income Hypothesis," Journal of Political Economy, 73, pp. 499-515.

Blinder, A.S., 1981, "Temporary Income Taxes and Consumer Spending," Journal of Political Economy, 89, pp. 26-53.

Blinder, A.S., and Deaton, A., 1985, "The Time Series Consumption Function Revisited," in Brookings Papers on Economic Activity: 2:1985, pp. 465-521.

Bodkin, R., 1959, "Windfall Income and Consumption," American Economic Review, 49, pp. 602-614.

Bodkin, R., 1966, "Windfall Income and Consumption: Reply," American Economic Review, 56, pp. 540-546.

Bowman, D., Minehart, D., and Rabin, M., 1993, "Loss Aversion in a Savings Model," U. of California at Berkeley, mimeo, May.

Campbell, J., and Mankiw, N., 1990, "Permanent Income, Current Income, and Consumption," Journal of Business and Economics Statistics, July, pp. 265-279.

Carroll, C.D., 1991, "Buffer Stock Saving and the Permanent Income Hypothesis," Federal Reserve Board of Governors, mimeo.

Carroll, C.D., and Summers, L.H., 1991, "Consumption Growth Parallels Income Growth: Some New Evidence," in Bernheim, B.D., and Shoven, J.B., eds., National Saving and Economic Performance, Chicago: Chicago University Press for NBER, pp. 305-343.

Deaton, A., 1992, Understanding Consumption, Oxford: Oxford University Press.

Eberly, J., 1994, "Adjustment of Consumers Durable Stocks: Evidence from Automobile Purchases," Journal of Political Economy, forthcoming.

Garcia, R., Lusardi, A., and Ng, S., 1994, "Excess Sensitivity and Asymmetries in Consumption: An Empirical Investigation," University of Montreal, mimeo, June.

Goldman, R., 1993, "I, Spy: My Secret Life as a Consumer," New York Times, July 18, op.ed..

Hall, R.E., 1978, "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, 96, pp. 971-987.

Hall, R.E., and Mishkin, F.S., 1982, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, 50, pp. 461-481.

Hayashi, F., 1985, "The Effect of Liquidity Constraints on Consumption: A Cross-Sectional Analysis," Quarterly Journal of Economics, 100, pp. 183-206.

Hubbard, R.G., and Judd, K.L., 1986, "Liquidity Constraints, Fiscal Policy, and Consumption," in Brookings Papers on Economic Activity, 1:1986, pp. 1-59.

Jappelli, 1990, "Who Is Credit Constrained in the U.S. Economy," Quarterly Journal of Economics, pp 219-234.

Jappelli, T., Pischke, J.S., and Souleles, N., 1994, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," M.I.T, mimeo, September.

Jones, R.C., 1960, "Transitory Income and Expenditures on Consumption Categories," AEA Papers and Proceedings, 50, pp. 584-592.

Kreinin, M.E., 1961, "Windfall Income and Consumption: Additional Evidence," American Economic Review, 51, pp. 388-390.

Laibson, D., 1993, "Golden Eggs and Hyperbolic Discounting," MIT, mimeo, November.

Landsberger, M., 1966, "Windfall Income and Consumption: Comment," American Economic Review, 56, pp. 534-540.

Levenson, A.R., 1993, "Do Consumers Respond to Future Income Shocks? Evidence from Social Security Reform in Taiwan," Milken Institute for Job and Capital Formation, September.

Lusardi, A., 1992, "Permanent Income, Current Income and Consumption: Evidence from Panel Data," CentER for Economic Research, Tilburg University, mimeo, December.

Lusardi, A., 1993, "Euler Equations in Micro Data: Merging Data from Two Samples," CentER for Economic Research, Tilburg University, mimeo.

Mankiw, G., 1982, "Hall's Consumption Hypothesis and Durable Goods," Journal of Monetary Economics, 10, pp. 417-425.

Mariger, R.P., and Shaw, K., 1990, "Unanticipated Aggregate Disturbances and Tests of the Life-Cycle Model Using Panel Data," Federal Reserve Board of Governors and Carnegie Mellon University, mimeo.

McNees, S.K., 1973, "Tax Refunds and Consumer Spending," New England Economic Review, January/February, pp. 3-11.

Meghir, C., and Weber, G., 1993, "Intertemporal Non-Separability or Borrowing Restrictions? A Disaggregate Analysis Using the US CEX Panel," Institute for Fiscal Studies, mimeo, June.

Modigliani, F., and Steindel, C., 1977, "Is a Tax Rebate an Effective Tool for Stabilization Policy?," in Brookings Papers on Economic Activity, 1:1977, pp. 175-209.

Nelson, J.A., 1992, "Consumer Expenditure Survey, Interview Survey, for Household-Level Analysis, 1980-1989," produced by Bureau of Labor Statistics, distributed by Inter-University Consortium for Political and Social Research.

Neumark, D., 1992, "Are Rising Wage Profiles a Forced-Saving Mechanism?," NBER Working Paper, November.

Paxson, C.H., 1992, "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," American Economic Review, pp. 15-33.

Poterba, J.M., 1988, "Are Consumers Forward Looking? Evidence from Fiscal Experiments," AEA Papers and Proceedings, 78, pp. 413-418.

Reid, M.G., 1962, "Consumption, Savings and Windfall Gains," American Economic Review, 52, pp. 728-737.

Runkle, D.E., 1991, "Liquidity Constraints and the Permanent-Income Hypothesis: Evidence from Panel Data," Journal of Monetary Economics, 27, pp. 73-98.

Shapiro, M.D., and Slemrod, J., 1993, "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," NBER Working Paper, April.

Shea, J., 1993, "Myopia, Liquidity Constraints, and Aggregate Consumption: A Simple Test," University of Wisconsin, mimeo, June.

Slemrod, J., Christian, C., London, R., and Parker, J., 1994, "April 15 Syndrome," mimeo, June.

Thaler, R.H., 1990, "Anomalies: Saving, Fungibility, and Mental Accounts," Journal of Economic Perspectives, 4, pp. 193-205.

Thaler, R.H., with Loewenstein, G., 1992, "Intertemporal Choice," in Thaler, R.H., The Winner's Curse, Ch. 8.

Wilcox, D.W., 1989, "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis," Journal of Political Economy, 97, pp. 288-304.

Wilcox, D.W., 1990, "Income Tax Refunds and the Timing of Consumption Expenditure," Federal Reserve Board of Governors working paper, April.

Zeldes, S.P., 1989, "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, pp. 305-346.

Bureau of the Census, Department of Commerce, 1994, "Quarterly Interview Survey: 1994 Forms, Consumer Expenditure Survey."

Bureau of Labor Statistics, 1993 and other years, "Interview Survey Public-Use Tape Documentation."

Department of the Treasury, "Daily Treasury Statement," various issues.

Internal Revenue Service (Statistics of Income Division), various years, Individual Income Tax Returns, Washington, D.C.: U.S. Government Printing Office.

FIGURES 1: Examples of the Pattern of Consumption Responses b_{4q}

Fig 1a : b(4q) (times 3): Spring Effect

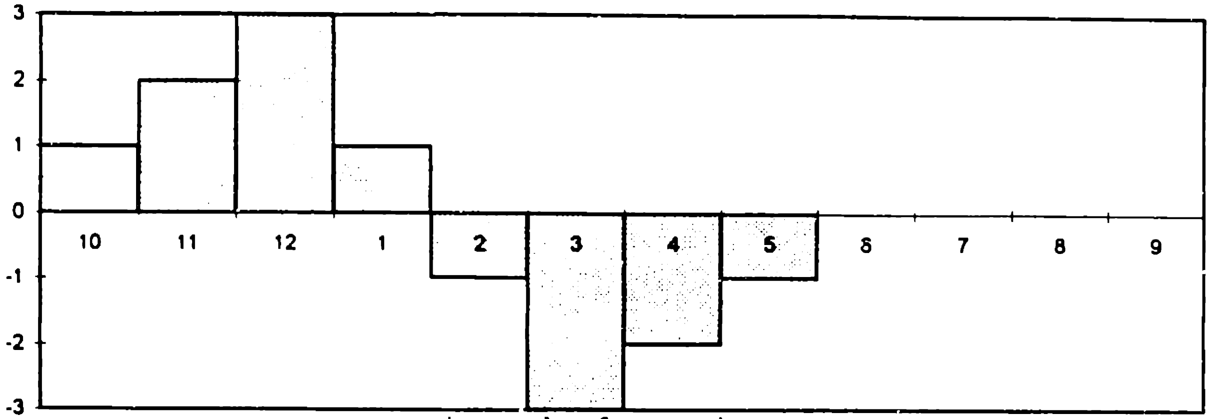

starting month q of consumption quarter t

Fig 1b: b(4q) (times 3): Spring and Summer Effects

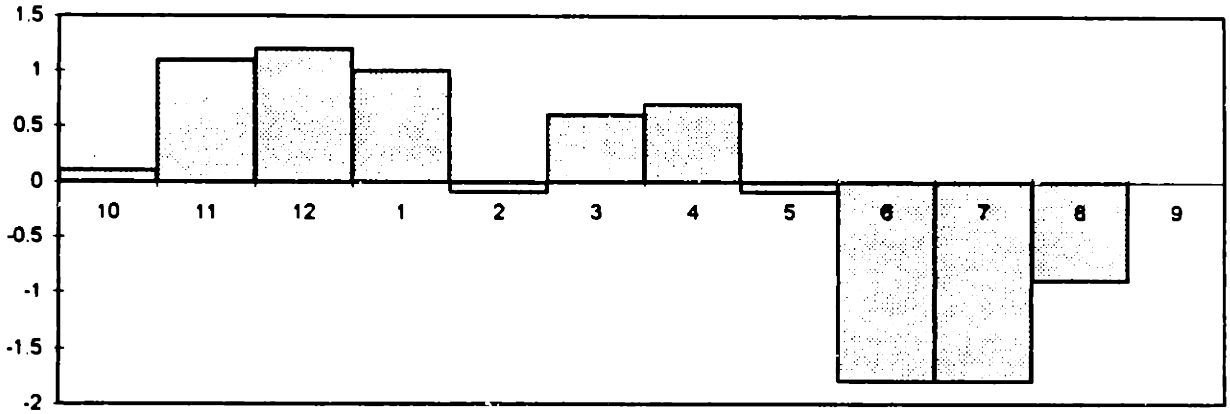

starting month q of consumption quarter t

Figures give observed dC(t+1) under the following patterns of refund-consumption

fig:	1a			1b		
month\person	1	2	3	1	2	3
Jan						
Feb						
Mar	1			0.1		
Apr		1			1	
May			1			0.1
Jun						
Jul				0.9		
Aug						0.9
Sep						
Oct						
Nov						
Dec						

(Eg for 1a and q=1: C(Apr,May,Jun) - C(Jan,Feb,Mar) = 1+1+0 -0-0-1 =1)

FIGURE 2: Timing of Refund Variables for period q =1 [C(Apr-Jun)-C(Jan-Mar)]

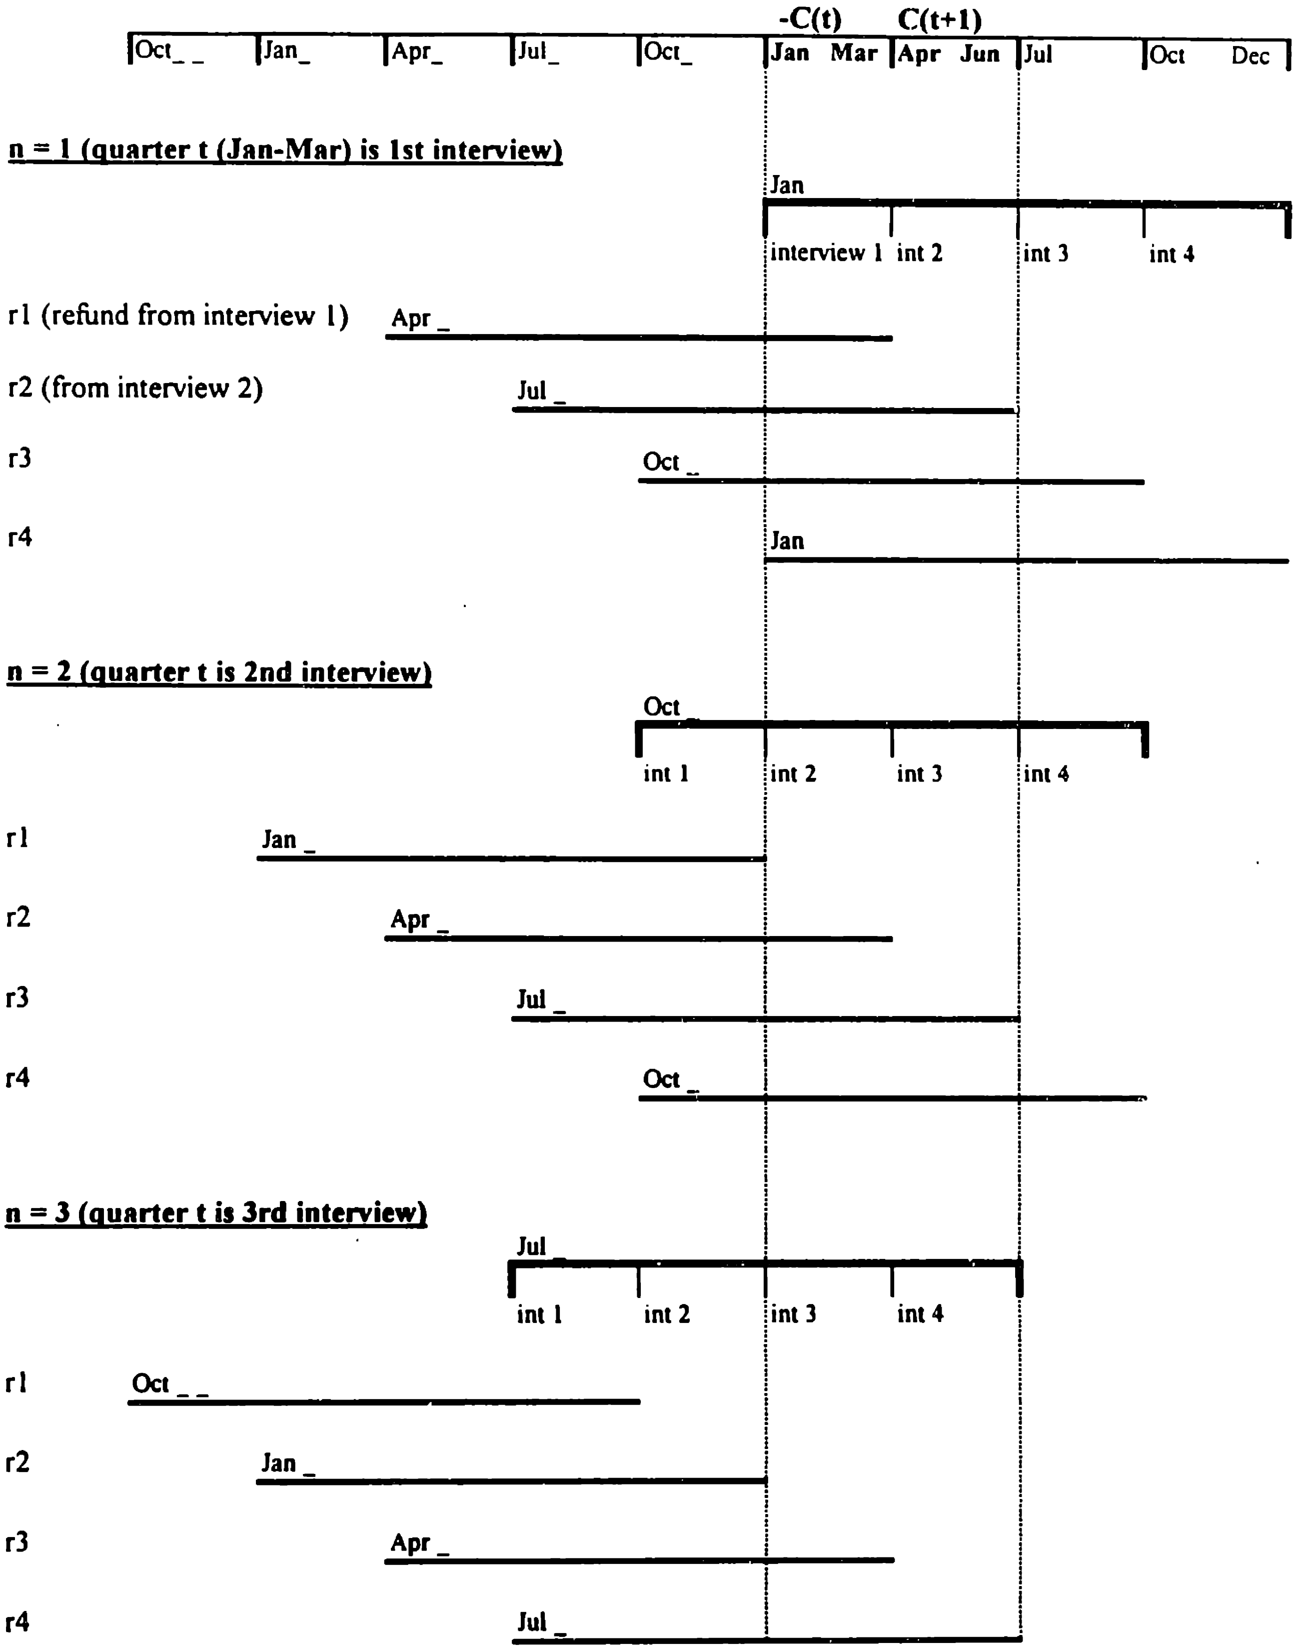

TABLE 1: The Scope and Scale of Income Tax Returns

FY	Returns Receiving Refunds (millions) [% of total returns]	Amount of Refunds (billion $) [% of total tax liability]	Average Real Refund (82-84$)	Average Real CEX Refund (82-84$)
1991	82.4 [71.9%]	84.0 [17.8%]	$730	----
1989	78.3 [69.8%]	71.7 [15.9%]	711	$771
1987	77.6 [72.5%]	69.6 [18.1%]	767	851
1985	74.9 [73.6%]	68.9 [20.3%]	846	912
1983	72.6 [75.4%]	62.0 [22.0%]	829	902
1981	68.3 [71.6%]	53.0 [18.2%]	816	899
1979	69.6 [75.1%]	43.0 [19.5%]	763	845

CEX: all years
mean $874
median $561

Notes:

- Source: IRS, Individual Income Tax Returns, table 3.3 (3.9 for 1979, 3.6 for 1981). The fiscal year is the year before the refunds are received.
- Real values are calculated as described in the data appendix.
- Values for the CEX are based on the refund variables used for particular periods q and interviews n, as described in the text. To maintain comparability with the IRS data, these averages are not limited to the sample used in the estimation (eg, those with unsatisfactory consumption data).

TABLE 2: The Timing of Disbursements of Tax Refunds, 1990

Month	Amount (million $)	Percent of Total
January	195	0.3
February	5,754	7.6
March	22,420	29.7
April	19,346	25.7
May	18,463	24.5
June	1,303	1.7
July	1,390	1.8
August	1,438	1.9
September	1,719	2.3
October	798	1.1
November	2,033	2.7
December	554	0.7
Total	75,412	100.0

% in	Mar - May:	79.8
% in	Jan - Mar:	37.6
	Apr - Jun:	51.9
	Jul - Dec	10.5

Notes:

- Source: Daily Treasury Statements, for Individual federal income tax refunds
- Half of the $5.5 billion of disbursements on the fifth Friday in March, the 30th, were allocated to April to keep the distribution more typical of other years.

TABLE 3: Excess Sensitivity Tests, all periods q = 1 - 9

Coefficient b_{4q} on Refunds for Different Consumption Groups, with and without the Control Group (Heteroscedasticity-corrected standard errors in parentheses)

period\group [# obs]	Food	Strictly Nondurables	Nondurables	Total Consumption
q = 1 [4160]	0.011 (0.009)	0.024 * (0.013)	0.031 (0.021)	0.160 * (0.084)*
w/control	0.010 (0.008)	0.023 ** (0.012)	0.029 (0.018)	0.156 ** (0.073)
q = 2 [2781]	-0.003 (0.012)	-0.038 * (0.020) [b]	-0.060 * (0.033)	-0.059 (0.089)
w/control	0.003 (0.010)	-0.014 (0.017)	-0.033 (0.028)	0.033 (0.076)
q = 3 [2919]	0.015 * (0.009)	0.047 ** (0.019)	0.051 ** (0.026)	-0.033 (0.095)
w/control	0.010 (0.008)	0.035 ** (0.017)	0.043 * (0.023)	-0.001 (0.081)
q = 4 [4350]	-0.011 (0.009)	0.006 (0.015)	-0.007 (0.021)	0.009 (0.056)
w/control	-0.007 (0.008)	0.008 (0.013)	0.002 (0.019)	0.015 (0.050)
q = 5 [2969]	-0.007 (0.011)	0.009 (0.020)	-0.013 (0.031)	0.090 (0.086)
w/control	-0.005 (0.010)	-0.003 (0.018)	-0.035 (0.026)	0.040 (0.074)
q = 6 [2930]	0.008 (0.009)	0.022 (0.015)	0.020 (0.046)	0.013 (0.112)
w/control	0.000 (0.008)	0.006 (0.014)	0.002 (0.041)	-0.037 (0.097)
q = 7 [4216]	-0.009 (0.008)	-0.028 * (0.017)	-0.008 (0.022)	0.009 (0.069)
w/control	-0.014 ** (0.007)	-0.034 ** (0.014)	-0.023 (0.019)	0.003 (0.059)

q = 8 [2757]	-0.000 (0.013)	-0.019 (0.020)	0.005 (0.030)	-0.121 (0.084)
w/control	-0.005 (0.011)	-0.027 (0.017)	-0.000 (0.026)	-0.105 (0.072)
q = 9 [2654]	0.001 (0.010)	-0.005 (0.018)	-0.021 (0.019)	0.053 (0.083)
w/control	0.002 (0.009)	-0.005 (0.016)	-0.022 (0.017)	0.054 (0.073)

* = significant at the 10% level
**= significant at the 5% level
a = p-value of .056
b = p-value of .058

TABLE 3, ctd

Coefficients on Demographics, for period q = 1, with and without the Control
(Heteroscedasticity-corrected standard errors in parentheses)

q = 1	Food	Strictly Nondurables	Nondurables	Total Consumption
age	1.20 ** (0.62)	1.17 (1.03)	2.01 (1.65)	-2.19 (4.84)
w/control	0.80 * (0.47)	0.73 (0.76)	1.16 (1.18)	1.78 (3.39)
d(adult)	47.0 * (28.0)	120.9 ** (42.8)	194.6 ** (71.0)	333.6 ** (157.7)
w/control	90.0 ** (22.1)	138.7 ** (36.5)	160.8 ** (53.6)	120.4 (141.0)
d(kids)	68.9 ** (30.6)	123.8 ** (43.2)	63.6 (121.1)	366.9 (238.0)
w/control	32.2 (25.4)	79.9 ** (37.5)	31.2 (80.5)	369.6 * (211.5)
# obs	4160	4160	4119	4119
w/control	7703	7703	7603	7603

* = significant at the 10% level
**= significant at the 5% level

Notes:

- q is the starting month of the trailing consumption period -C(t). Eg, q=1 refers to C(Apr,May,Jun) - C(Jan,Feb,Mar), January being the first month. The control group are those with no refunds.
- To save space constants and time-dummies, and the demographics and sample sizes for periods besides q=1, are not reported.

TABLE 4: Splitting for Liquidity Constraints, period q = 1 [C(Apr-Jun) - C(Jan-Mar)]

Coefficient b_4 on Refunds, for Samples at the extremes (%) of the Distributions of Financial Variables

Sample Split\ Consumption group [# obs]	Food	Food w/con-trol	Strictly Nondurs	Strictly Nondurs w/con-trol	Total Cons.	Total Cons. w/con-trol
Liquid Wealth						
< $100 (10%)	.014	.018	.057	.063	.082 [a]	-.021 [a]
[315]	(.024)	(.023)	(.039)	(.036)	(.197)	(.181)
> $7500 (75%)	-0.003	-.006	.002	.001	.238**	.205 *
[616]	(.006)	(.006)	(.010)	(.010)	(.114)	(.121)
current refunds						
< $100	.050**	.044**	.081 *	.087**	.131 [b]	-.057 [b]
[134]	(.020)	(.019)	(.043)	(.039)	(.201)	(.186)
>$7500	-.006	-.007	-.004	.003	.296**	.260**
[305]	(.005)	(.005)	(.008)	(.008)	(.090)	(.108)
Earnings						
<$15000 (25%)	.039 *	.032 *	.084**	.077**	.142 [c]	.153 [c]
[897]	(.022)	(.018)	(.029)	(.027)	(.129)	(.121)
>$25000 (50%)	.008	.009	.017	.023 *	.172 *	.173 *
[1934]	(.009)	(.008)	(.012)	(.013)	(.106)	(.092)
current refunds						
< $ 15,000	.044	.041	.085**	.083**	.324**	.296**
[410]	(.030)+	(.025)	(.030++	(.028++	(.135)[d]	(.136)[d]
> $ 25,000	-.000	-.002	.014	.017	.223 *[h]	.203 *
[920]	(.008)+	(.007)	(.013++	(.013++	(.117)[g]	(.111)[g]
Wealth/Earns						
< .01 (10%)	.032	.054**	.076**	.103**	.153 [e]	.050 [e]
[453]	(.021)	(.019)	(.022)	(.029)	(.146)	(.134)
> .08 (50%)	-.001	-.002	.012	.007	.198 *	.199 *
[1247]	(.006)	(.006)	(.010)	(.009)	(.119)	(.112)
current refunds						
< .01	.041*	.057**	.098**	.114**	.068 [f]	-.118 [f]
[199]	(.025)	(.021)	(.045)	(.040)	(.214)	(.193)
> .08	-.008*	-.008 *	-.000	.001	.234**	.233**
[564]	(.005)	(.005)	(.007)	(.008)	(.116)	(.117)+[g]

* = significant at 10%
** = significant at 5%
\+ = coefficients for constrained and unconstrained differ significantly at 10%
++ = coefficients for constrained and unconstrained differ significantly at 5%

a = the nondurables of constrained have coefs .114 (.060)*, .125 (.052)** with control
b = the nondurables of constrained have coefs .164 (.072)**, .152 (.063)** with control
c = the nondurables of constrained have coefs .076 (.037)**, .072 (.038)* with control
d = the nondurables of constrained have coefs .129 (.029)**, .112 (.031)** with control
e = the nondurables of constrained have coefs .137 (.052)**, .170 (.046)** with control
f = the nondurables of constrained have coefs .168 (.063)**, .186 (.055)**+ with control
g = the coefs for constrained's nondurables significantly differ from the unconstrained's
h = p-value is .056
- Heteroscedasticity-corrected standard errors in parentheses

Notes:

- Current refunds are those known to have come within the current year in most cases, or since October of the previous year in a few cases. See text for details.
- \+ and ++ are taken from the unreported pooled regression of (only) the constrained and unconstrained. The demographics were constrained to have the same coefs for both groups. The reported results come from separate regressions for the constrained and unconstrained.
- The cutoffs for the splits are shown in parentheses. Eg, for earnings, at the 25th and 50th percentiles the cutoffs are $15,000 and $25,000 respectively.

- The number of observations is for food and strictly nondurables, without the control group; total consumption has a few fewer observations.
- The control group is those without refunds.
- Financial variables are deflated to 1982-84.

TABLE 5: Late Refunds and Nonlinearity, period q=1 [C(Apr-Jun)-C(Jan-Mar)]

Coefficient on Refunds and Refunds*I(Refund is late)
(Heteroscedasticity-corrected standard errors in parentheses)

Sample \ Consumption group [# obs, # late]	Food		Strictly Nondurs		Total Cons.	
	refs	refs*late	refs	refs*late	refs	refs*late
Wealth split						
< $100 (10%)	.004	.106*	.045	.125	.102 [a]	-.209 [a]
[315, 28]	(.027)	(.063)	(.043)	(.147)	(.207)	(.362)
>$7500 (75%)	-.005	.013	-.003	.042	.271**	-.322
[616, 82]	(.005)	(.032)	(.010)	(.038)	(.103)	(.201)
Earnings split						
current refunds						
< $ 15,000	.040	.368**	.081**	.368**	.329**	-.465 [b]
[410, 68]	(.026)	(.168)	(.027)	(.186)	(.137)	(.797)
> $ 25,000	-.003	.032*	.007	.083**	.236**	-.166 [b]
[920, 133]	(.007)	(.019)	(.011)	(.032)	(.115)	(.191)

* = significant at 10%
** = significant at 5%
a = for nondurables, the coef is .107 (.064)* for refund, .082 (.207) for refund*late
b = for nondurables, the coef on refs*late is .515 (.228)** for the constrained, .024(.121) for the unconstrained

TABLE 5, ctd

Coefficient on Refunds and Refund^2/Earnings
(Heteroscedasticity-corrected standard errors in parentheses)

Sample \ Consumption group [# obs]	Food		Strictly Nondurs		Total Cons.	
	refs	refs^2/ earnings	refs	refs^2/ earnings	refs	refs^2/ earnings
Whole Sample [3780]	.013 (.010)	-.000 (.004)	.026* (.015)	.001 (.004)	.158** (.081)	.025 (.030)
Earnings split						
<$15000 (25%) [897]	.036 (.023)	.001 (.003)	.084** (.033)	-.000 (.004)	.071 (.119)	.027 (.021)
>$25000 (50%) [1934]	.023** (.012)	-.028** (.011)	.033 (.020)	-.029 (.020)	-.053 a (.104)	.417** (.102)

* = significant at 10%
** = significant at 5%
a = for nondurables, the coef for the unconstrained is .038 (.034), .041 (.055) for current refunds

Notes:

- "late" is an indicator for a refund that is known to be "late", ie that is known to have come in April or later. These are about 10% of refunds. See text for more detail.
- Current refunds are those known to have come within the current year in most cases, or since October of the previous year in a few cases.

- The number of observations is for food and strictly nondurables, without the control group; total consumption has a few fewer observations.
- Financial variables are deflated to 1982-84.

TABLE 6: Trips, period q = 3 [C(Jun-Aug) - C(Mar-May)]

Coefficient on Refunds and Refunds*I(trip in Jun-Aug), with and without Earnings
(Heteroscedasticity-corrected standard errors in parentheses)

Sample \ Consumption group	Food	Food w/control	Strictly Nondurs	Strictly Nondurs w/control	Total Cons.	Total Cons. w/control
refunds	.014 (.010)	.009 (.009)	-.014 (.016)	-.010 (.015)	-.037 (.083)	-.030 (.079)
refunds*trip	.003 (.013)	.000 (.012)	.058** (.023)	.033 (.021)	-.003 (.116)	.009 (.106)
trip intercept	85.4** (19.7)	96.5** (13.0)	205.6** (29.1)	268.9** (20.3)	417.4** (136.6)	378.9** (95.8)
4432 obs						
w/ Earnings						
refunds	.017* (.010)		-.019 (.018)		-.025 (.086)	
refunds*trip	-.002 .013		.056** (.025)		.019 (.118)	
trip intercept	100.1** (21.3)		209.2** (31.4)		435.8** (144.5)	
earnings	-.001 (.001)		.001 (.001)		-.006 (.005)	
4085 obs						

* = significant at 10%
** = significant at 5%

Notes:

- "trip" is an indicator for having made expenditures on trips in June-August. The same trip intercept is used for trip-takers among both the refund-recipients and the control group, those without refunds.
- All expenditures on trip are either food, strictly nondurable or nondurable. (But all these are part of total consumption as well.)
- The number of observations is for food and strictly nondurables, without the control group; total consumption has a few fewer observations.
- Earnings are deflated to 1982-84.

Chapter 3. College Tuition and Household Savings and Consumption.[1]

0. Introduction.

Sending the kids to college is obviously a central episode in the life-cycle of many households. Even so there has been surprisingly little study of the adequacy of household savings and other resources available to pay for college. In part this might be due to the difficulty of determining what would count as adequate savings for a given household.[2] It is not enough just to measure the level of assets on starting college, for instance, because these assets might be earmarked for other purposes like retirement or medical expenses. Also such a measure would not distinguish a far-sighted household that began to save early, from a short-sighted one that began to save only at the start of college, by sharply cutting its consumption.

This paper gauges the adequacy of household savings and other resources for college by examining whether households are able to maintain their standard of living--that is, their consumption--as they pay for college. This examination recognizes that,

[1]I thank Olivier Blanchard, Ricardo Caballero, Pierre Gourinchas, Jerry Hausman, Jonathan Parker, Steve Pischke, and participants of the MIT money lunch for helpful comments; and the National Science Foundation for support. None is responsible for any errors herein.

[2]There is a large literature on the effect of household income and wealth on the decision to go to college, but little analysis--discussed below--of the adequacy of savings conditional on going to college.

It is interesting that some of the current proposals for tax cuts link the cuts to the burden of child-rearing.

given the cost of college, what matters for household welfare are distortions in the path of consumption resulting from financing that cost. Such a gauge distinguishes the far- and short-sighted households above, without needing explicitly to disentangle which of their savings are earmarked for college.

The data are drawn from the Consumer Expenditure Survey, which of the leading U.S. data sets is widely regarded to contain the best data on household consumption. The Survey records educational expenditures in detail. Specifically, this paper tests whether households' non-educational consumption decreases in the fall and winter in proportion to their college expenditures in the fall. The change in consumption is measured in relation to consumption in the previous summer, spring, and winter, to determine whether any saving sharply accelerated during these three seasons, or whether it was fully underway by then. Of course households use current income and loans in addition to savings to pay for college. The test here considers the change in consumption given all the resources available to the household. This is the appropriate consideration as regards household welfare.

Although it would be quite interesting to examine the response of consumption over longer horizons as well, the data do not permit this. Nonetheless, the periods examined here are of the greatest interest, since they cover the time when most students choose their colleges and when college payments are begun.

This examination also has general implications for the canonical life-cycle theory of consumption and saving. First, paying for college often substantially reduces the savings a household has going into retirement. More generally, "consumption needs" like child-rearing or buying a house can account for much of the oft-noted paucity of financial assets in mid-life, without requiring a high degree of impatience. And yet such investments in education and housing should properly be thought of in large part as saving.

Second, on paying for college a household's net income, that is its income net of the payment, is typically predictably decreasing. Therefore, assuming separability of "college-services" and of the rest of consumption, the response of non-college consumption to such payments constitutes a test of whether consumption "tracks" income.

Under the life-cycle theory, households should save in advance to meet the costs of college (net of aid and loans), in order to maintain their consumption during college.

This particular test of the theory has two advantages compared to most other such tests. The first is that, since net income is decreasing on payment, the test is free of complications due to liquidity constraints.[3] The second advantage concerns a generic problem of life-cycle tests, namely the need to control for changes in the marginal utility of consumption over the life-cycle. Focusing on a particular, salient life-cycle episode--paying for college--controls for much of the variation in marginal utility.

The outline of this paper is as follows. Section I provides background on college costs and reviews the literature on saving for college. Section II briefly describes the data. Section III sets out the null and alternative hypotheses to be tested, on the extent to which college costs are paid for out of consumption. Section IV reports the results. The conclusion, Section V, is followed by an Appendix further describing the data.

I. Paying for College.

This Section summarizes what is known about the amounts that households paid for college in the 1980's and early 1990's, and the sources of the funds they used to pay.[4]

---

[3]This holds even if a household is constrained from borrowing against its post-college resources: it should still save before college to maintain its consumption during college, in the face of its college expenses net of any loans.

[4]McPherson and Schapiro [1991] provide a good introduction.

Despite many changes in the particulars, the basic picture presented below of the financing of college did not essentially change during the 1980's and early 1990's. See also Mumper [1993].

1. College Costs.

Table 1 reports the costs of different types of colleges. Of the 10 million or so undergraduates each year, about 40% are enrolled in public 4-year colleges. There the average total cost (including room and board) for the 1992-93 academic year was about $6000 (current dollars). About 20% of undergraduates are at private 4-year colleges, where the cost averaged about $15,000 per year. The remaining undergraduates are mostly at public 2-year colleges, where tuition alone was about $1000 per year.

The fraction of these large costs that households actually pay is hard to estimate. Among other things it varies greatly with the financial-aid package the household receives. Miller and Hexter [1985, 1985a] estimate this fraction for full-time students financially dependent on their families, using the 1983-84 Student Aid Recipient Survey. For the most common aid-package, consisting of a Pell grant, and state and campus-based aid, middle-income households paid at least two-thirds of the total cost of college. Under the next most common package, with a Guaranteed Student Loan instead of state aid, the fraction rises to over 4/5, taking into account the repayment of the loan.[5] [6] For low-income households, the fractions are about 1/2 and 2/3 respectively. Miller and Hexter argue that

> these [low-income] families have virtually no discretionary income to tap and can rarely call upon savings for [paying this remaining fraction]. Thus

---

[5]Pell grants and GSL's are the most important forms of financial aid. In 1988-89, over 3 million students received Pells, averaging about $1,400. About 3.6 million (including graduate students) took out subsidized Stafford loans, averaging about $2600. Many of these students took out the maximum Stafford loan. For them and others ineligible for Staffords, there are also guaranteed PLUS (parent loans to undergraduate students) and SLS (supplemental loans to students) loans. In 1988-89, about 800 thousand students (including graduate students) borrowed an average of about $2600 via SLS's, and 200 thousand about $3100 via PLUS's. (McPherson and Schapiro [1991]).

Little is known about other, non-guaranteed loans taken out for college, apart from the self-reported total use of loans described below.

[6]About 1/3 of aid-recipients at public colleges receive only GSL's, so pay 100% of cost (apart from the interest subsidy). For private colleges, the figure is about 1/8.

we can only speculate that these students are living at a lower standard than was allowed for in their student budget or that they are receiving additional support (from grandparents or other relatives, for example) that is not captured in the need analysis process.

The 1986-87 National Postsecondary Student Aid Survey [NPSAS] asked parents how much they contributed to their children's college costs during the academic year. Choy and Henke[1992] summarize the results for dependent undergraduates. 67% of parents gave an average of about $3900 to their children, or directly to the college, for tuition, housing, and other expenses. On adding the 11% of parents who extended loans to their children for these expenses, the average is about $4200.[7] On also adding the 83% giving in-kind gifts (eg housing and clothing), the average is about $6200. The probability and magnitude of gifts are generally increasing in the income and wealth of the parents.[8] For those attending private colleges, the average gift is about $6500, rising to $8800 including loans and in-kind gifts. These amounts, usually multiplied by 2 or 4 (for two- or four-year colleges respectively), are obviously not insubstantial relative to typical household income and savings.[9]

2. Sources of Funds.

According to the NPSAS, of the parents giving gifts (not in-kind gifts, but including loans), about 80% used current income as a source of these funds, 65% used previous savings, 24% used loans, and 30% used additional income from increased work. Unfortunately the amount of funds drawn from each source is not available. Still, there

---

[7]The extent to which the parents will in fact be paid back is unclear.

[8]This does not hold for loans. See also Churaman[1992] for a multivariate analysis of these gifts.

[9]The figures are also important relative to aggregate saving. Gale and Scholz [1994] estimate household contributions for college totalled about $35 billion per year in the mid-1980s.

It is interesting that such contributions represent about 1/3 of the difference between Kotlikoff/Summers and Modigliani in their debate about the relative importance of bequests versus life-cycle savings (Modigliani [1988]).

are two points worth noting. First, 35% of households did not save in advance at all.[10] This fraction rises to about 50% for households of lower income or wealth. Second, the extent to which those using current income are cutting their current consumption, or just diverting other current saving, remains unclear.

Of the parents saving, about 47% report they started saving before their child was in junior high, 44% started during junior and senior high, and 10% still waited until after high school ended (20% for poorer families). So, even of the households that save, over half wait until they are at most 6 years from college. And even those starting earlier might not be saving sufficiently. A 1984 Roper poll found that most households that were saving were not saving enough to meet at current rates their own target for savings by the time of college.[11]

Economists looking at college savings have mostly focused on the substantial tax to saving imposed by means-testing in the financial aid system (Case and McPherson [1986], Feldstein [1992], Edlin [1993]). In the context of the life-cycle model of saving, the focus can be said to have been on the intertemporal elasticity of substitution; not on the adequacy of savings, at issue here.

In sum, it remains unclear whether households have adequate savings and other resources for college, and in particular whether they are substantially cutting their consumption at the time of college. This particular issue is the focus of this paper.

---

[10]These figures still condition on the parents' giving. Unconditionally, over 1/2 of parents with dependent undergraduates did not at all save in advance. And, 14% took out loans, averaging $4000. Most of these were not federal, state, or institutional loans, so little is known about them.

[11]Of those not saving but still expecting to send a child to college, most said that the reason they were not saving was that they could not currently afford to save, but that they would start saving later. Still, about a quarter of the non-savers did not expect to start saving later, either. (Roper[1984])

II. The Data.

This Section discusses the variables and sample used in this paper, focusing on the variables reporting educational expenses. The Appendix provides further details.

The data are drawn from the Consumer Expenditure Survey [CEX] for 1980 to 1993. The CEX is a rotating panel: households are interviewed four times, three months apart, but starting in different months. Eg, some have their first interview in February, with the remaining interviews in May, August, and November; others are interviewed in March, June, September, and December; etc.

1. General Expenditures.

In each interview the CEX records household expenditures over the previous three months, the reference period. For each household reference-quarter, the various non-educational expenditures were deflated by the corresponding CPI subindexes (1982-4$)[12], and then aggregated to create the following consumption groups. The smallest grouping is **food**. The mean of households' food expenditure (before cutting the sample as below) to total expenditures is about .26. The broadest grouping is **nondurables**, which generally corresponds to the CPI's nondurables and services. The main exceptions are spending on education and health, which have been removed. The average share of nondurable to total expenditure is about .80.

Nondurables include many relatively durable and lumpy expenditures which households might not undertake regularly in each quarter, eg spending on wigs and repairs and for septic-tank cleaning. They also include various expenditures, notably clothing, that might not be tolerably separable from educational expenditures.[13] Such expenditures have been removed to create a subset of nondurables, "**strictly nondurables**". Its major components are food; household operations, including monthly utilities and small-scale

[12] One of the chief purposes of the CEX is the calculation of the weights for the CPI, so the match between the component expenditures and the CPI subindexes is close.

[13] Eg, many families send their kids off to college with new sets of clothing.

rentals; transportation fuel and services, including public transportation; personal services, and entertainment services and high-frequency fees. The average share of strictly nondurables to total expenditure is about .53.

2. Expenditure on College.

The CEX records expenditures on college tuition and supplies, and room and board. These have been added together and deflated (1982-84$) for each household-quarter to produce total college expenditure, **COLEXP**. These amounts reflect out-of-pocket outlays, net of expenses covered by non-cash scholarships. Unfortunately, cash contributions directly to students away at college are recorded only annually, so cannot be used below; and there is little direct information about the sources of the funds (eg, loans) used to make the expenditures. Nonetheless, COLEXP captures most of what is at issue here, the out-of-pocket burden of paying for college in each quarter given the households' resources.

For "traditional" collegiate households,[14] the average real COLEXP is about $980, the median $480, for any given quarter in which COLEXP is nonzero. These expenditures are greater and more frequent in quarters covering the start of the two academic semesters.[15] Figure 1 shows the age distribution of the heads of the households making these payments (the lower of the two paths). The frequency of payments (reflected in the size of the bubbles) is greatest for heads in their 40s and early 50s, the ages in which they typically send children to college. The age-distribution is markedly humped, peaking in these years at about $1300 per quarter.[16]

---

[14]For households without adults in education, as described in the next sub-section on the sample.

[15]A survey in El-Khawas [1985] reports that about 45% of colleges offer plans spreading out payment of the annual cost over the entire year.

[16]Figure 1 does not control for family size nor cohort effects. Ages are assigned to "bins" of two years. Only bins with greater than 20 observations have their means graphed. The total number of observations is about 7,500, with about 500 to 900 observations in each bin for those in their 40s and early 50s.

Figure 1 also shows the age-distribution of quarterly nondurable consumption, for the same households and the same quarters in which they have the college expenditures. This distribution follows a rather similar pattern as that for COLEXP below it. This similarity is especially interesting in light of the fact that the path of nondurable consumption itself "tracks" the path of income over time (Carroll and Summers [1991]). But it does not follow that there is a violation of the life-cycle theory here. First, the financial aid system itself leads to larger college payments when income and wealth are larger. Second, when households are paying for college, their net income (net of these payments) is usually decreasing. The analysis below makes the right comparison for testing the theory, comparing the change in consumption to the change in net income.

3. The Sample.

The sample is limited to the households making college payments that are "**traditional**" collegiate households. These are households in which there is a child aged 16-24, and without anyone over 24 in school. The finances and expenditures of other, less traditional households making college payments are harder to analyze.[17]

In particular, this limitation helps control for changes in the marginal utility of consumption over the life-cycle. Attanasio et al. [1994] stress that on controlling with a rich set of demographic variables, the hump in the life-cycle profile of consumption can be levelled. On the other hand, this richness decreases the power of tests of the life-cycle theory, against finding violations. The approach of this paper, by contrast, is to use the usual, sparser set of demographics, but to control for the most important remaining life-cycle characteristics by limiting the sample to households at a particular, salient position in the life-cycle--namely, to households with children in college.[18] The reduction in

---

[17]On the other hand, the "non-traditional" households are quite numerous. Eg, in the 80's about 1/2 of undergraduates were part-time students, and more than half of these were aged over 23. (McPherson and Schapiro [1991]).

[18]Figure 1, for instance, suggests that controlling for the age of the household head is not alone sufficient. Eg, some 40 year olds will have children in college, and others will not.

family size when a child goes away to college is controlled for in the sparse set of demographics, as usual.[19]

It is interesting to see whether the behavior of households sending children to college for the first time differs from the behavior of the those with children in college generally. For instance, if a household has not adequately saved at the start of the first year of college, its consumption might sharply decrease at that time, and remain low into later years. On the other hand, its consumption might drop only later, as savings are used up. Also, households might learn to save for college through experience. Evidence for this possibility comes from a 1990 GAO review of the studies in this area. It concludes that parents know rather little about college costs and financial aid before college, even into the senior year of high school. Parents' most important sources of information about college costs are the colleges themselves, presumably primarily during senior year, and relatives and friends. Thus it might matter whether any of a household's children has already started college.[20]

To get at these issues, a subset of traditional collegiate households most likely to have children starting college for the first time, "**starting**" households, was identified. Starting households are essentially those that have college expenditures in the fall, but not

---

More generally, the logic of the life-cycle theory seems to push for controlling for the life-cycle position of each member of the household, not just of the head. But such a multidimensional control would be hard to implement directly. This paper instead controls for much of these issues by limiting the sample as described.

[19]About 30% of white undergraduates, and 35-50% of minority ones, live at home while in college (Churaman[1992]).

[20]The NPSAS survey cited above has some mixed information on this point. Parents' reliance on current income and previous savings is not monotonic in the grade of the student (freshman to senior and beyond), but does not much vary with it anyway. Use of current income does decrease with the number of other children in college. Use of previous savings increases if there is one other child, though then decreases with more than one child. The largest difference is that parents with other children already in college are much more likely to be saving in general (as opposed to saving for education in particular): from 39% without any other children in college, to 62% with one other and 53% with more than one other child in college.

anytime before; have a child aged 16-19 likely to be starting college; and have no other members apparently in college. This identification is unavoidably imperfect, because the CEX is not intended to trace such family dynamics. Any errors in the classification of starters will attenuate the difference between them and the other households. The Appendix provides further details.

The sample was also limited to households with satisfactory data on consumption, as described in the Appendix. This leaves slightly over 1000 observations of traditional collegiate families, of which just under 200 are identifiable as starters.

II. Specification and Estimation.

This Section begins by specifying the null and alternative hypotheses used in testing the adequacy of savings at the start of the academic year. It then describes the timing and estimation involved in the test.

1. The Null and Alternative Hypotheses.

Following Zeldes [1989] and Lusardi [1992, 1993], the Euler equation that governs household i's change in consumption under the life-cycle theory is specified as

$$(1) \qquad dC_{i,t+1} = b_0 + (\text{month dummies})_{it} + b_1 * age_{it} + b_2 * d(\text{ages } 0\text{-}15)_{i,t+1} + b_3 * d(\text{ages } 16\text{-}24)_{i,t+1} + b_4 * d(\text{ages } 25\text{-}90)_{i,t+1} + u_{i,t+1},$$

where age_t is the age of the head of household at period t, and d(ages 0-15)$_{t+1}$, d(ages 16-24)$_{t+1}$, and d(ages 25-90)$_{t+1}$ record the change between periods t+1 and t in the number of household members less than 16 (eg, births), between 16 and 24 (inclusive), and greater than 24 (eg, deaths) respectively. d(ages 16-24) controls in particular for children's leaving the household for college, or for other reasons. These demographics allow for the most basic changes in household preferences, namely with family size and over time.

The alternative hypothesis is that the household pays for m % of its college expenditures COLEXP out of current consumption. The generalization of the Euler equation that is estimated is then

$$(2) \quad dC_{i,t+1} = b_0 + (\text{month dummies})_{it} + b_1{*}age_{it} + b_2{*}d(\text{ages } 0\text{-}15)_{i,t+1} + b_3{*}d(\text{ages } 16\text{-}24)_{i,t+1} + b_4{*}d(\text{ages } 25\text{-}90)_{i,t+1} - m{*}COLEXP_{i,t+1} + u_{i,t+1},$$

where t+1 is the quarter covering the household's first payments for the fall semester, and no college payments are made in quarter t.

A large m would evidence the inadequacy of the household's savings and other resources to maintain its consumption in the face of its college expenditures. It would also count as a violation of the life-cycle theory that would not be due to liquidity constraints, insofar as COLEXP is a predictable expense.[21] That is, under the life-cycle theory households should save in advance or borrow for college in order to maintain their consumption, so m should be zero.

2. Timing.

Figure 2 illustrates the timing of equation (2). The household reference quarters chosen to be quarters t+1 and t depend on when the household is interviewed and when it makes its college payments. For example, household #1 has a reference quarter of June through August. Suppose it makes its first payment in August. Then quarter June-August is chosen as its period t+1, and March-May as its period t. Household #4 by contrast might first pay in September, so its reference-quarter September-November is chosen as its period t+1. (But, if household #4 first pays in August, it is treated like household #1

[21]The predictable quarterly change in other kinds of income might also be added, but is not available in the CEX. Nonetheless, during the period at issue, COLEXP is likely to be the dominant part of the total change in net-of-college-expenses income, even allowing for some increased income from increased work. In any case, it is the right variable for testing the effect of college expenditures on consumption.

above.) More generally, included is any household with reference quarters t and t+1 as in Figure 2, with college payments in t+1 but not in t. In the basic specification of equation (2), C_{t+1} - C_t can loosely be said to compare consumption in the fall to consumption in the previous spring and summer.

It is interesting to examine other periods as well. Figure 3 shows two households whose change in consumption between t and t+1 is the same, but whose previous path of consumption is quite different. Household #1 starts to save seriously only at t, and so suffers a large drop in consumption between quarters t-1 and t. Household #2 by contrast started saving earlier, lowering its consumption before t to avoid a large drop at t. To distinguish such households, the change in consumption over a longer period, C_{t+1} - C_{t-1} (not $-C_t$) will also be examined. This compares consumption in the fall to consumption roughly in the previous winter and spring. The horizon will also be moved forwards, to quarter t+2--roughly, the winter--to learn about any delayed response of consumption, and to protect further against distortions to C_{t+1} associated with the start of the academic year (like back to school sales).[22] Unfortunately there are not enough observations to look further backwards or forwards. Nonetheless, the examinations just described are sufficient to learn about whether households have saved enough 6-9 months in advance of the academic year, or have sufficient other resources, to sustain their consumption up to 6 months into it.

There are also complications in looking any further back. First, most households learn about admissions-decisions and financial aid packages in the spring, and so face some uncertainty about their college expenditures before then. Also, before a household commits to a college, its expenditures might be endogenous: the choice of a college, and so college expenditures, might depend on the household's savings. Equation (2), by contrast, considers consumption only from the spring/summer on, when these expenditures

[22]Recall that the expenditures most likely to be thus distorted have been removed from strictly nondurables.

Also, these alternative timings for the change in consumption are considered only for the same households as in the base case, namely those having college expenditures in periods t+1 but not in periods t.

have been largely predetermined. Second, the means-testing of the financial aid system substantially distorts households' incentives to save before filling out their financial aid forms, typically done in the first calendar quarter.[23] Again, by limiting its focus to the spring forwards, equation (2) minimizes the distortion to the return to saving.

On the other hand, there remains uncertainty for some households even after the spring, affecting even the modest scope of equation (2). But by the spring any remaining uncertainty should be of second order, for most households.

3. Estimation.

Equation (2) is estimated by OLS, correcting the standard errors for heteroskedasticity. A full set of month dummies (different dummies for the same month in different years) controls for seasonality and common shocks across households. The regressor COLEXP is relatively well measured, compared to the income and wealth variables typically used in life-cycle tests using the CEX.

IV. Results.

This Section begins by estimating the basic specification, equation (2). It then turns to alternative timings, and to the response of households most likely to be starting to pay for college for the first time.

1. Base Case.

Table 2 records the results of estimating equation (2), for consumption in the fall relative to consumption in the previous summer/spring. For nondurables the coefficient on college expenditures, **COLEXP**, is positive, the "wrong" sign, and marginally

---

[23]See Feldstein [1992] and Edlin [1993]. Case and McPherson [1986], by contrast, suggest that the distortion might not in fact be very important. The GAO [1990] study discussed above tends to support them, insofar as households do not typically know much about the financial aid system.

A classmate of mine, T. Kim, is currently further investigating these issues.

significant at the 10% level. As discussed in Section 2, however, some expenditures in nondurables, like clothing, might not be tolerably separable from COLEXP. This gives reason to focus on **strictly nondurables**, which drops such expenditures. The coefficient for strictly nondurables is insignificant and small. The coefficient for food is slightly negative, but again insignificant.

The coefficients are relatively precisely estimated. Eg, the 95% confidence interval for the response of strictly nondurables is about (-.04, .07). That is, with 95% confidence, any decrease in strictly nondurable consumption is less than 5 cents for each dollar of college expenditure. It is worth noting that this test of the life-cycle theory should be relatively powerful compared to many other such tests, because of the large variation in the regressor, college expenditures, and the fact that these expenditures need not be instrumented for.[24] Further, the quality of the variables in the CEX recording these expenditures is likely to be better than that of the regressors used in many other tests of the theory.

As for the demographics, only the coefficients on d(ages 16-24) and d(ages 25-90) are significant. As expected, they are positive, reflecting increased spending with additional household members. The coefficients on d(ages 0-15) have the wrong sign, but are much less precisely estimated. The month dummies, not shown, are together highly significant.

It is of interest to know what assets households use to pay their college expenditures. The CEX provides some information on the change in household liquid assets (checking and savings accounts) between their first and fourth interviews. Using this change as the dependent variable in equation (2), with the other variables the same, yields a coefficient of -.39 (.33) on COLEXP, suggesting that about 2/5 of initial college expenditures come out of liquid assets. This result must be considered only tentative, however, because the wealth variables in the CEX are generally considered to be of lesser quality than the expenditure variables.

---

[24]See Shea [1995].

In sum, one cannot reject the null hypothesis that households maintain their consumption into the fall term, despite large college expenditures. Households appear to have sufficient savings or other resources to finance these expenditures without cutting into their consumption at this time. The next subsection considers their consumption over longer horizons.

2. Alternative Timing.

Table 3 presents the main results on looking further before and after the periods covered by equation (2). Consider first the time after the fall. Column (1) is for a regression like equation (2), but of $C_{t+2} - C_t$ on $COLEXP_{t+1}$.[25] This regression allows consumption an additional quarter to adjust to the college expenditures in period t+1. That is, it compares consumption in the winter to consumption in the summer and spring. (See Figure 2.) The coefficients on COLEXP remain insignificant. Still, it is noteworthy that they are now negative, suggesting that there might be a small delayed response of consumption. The point estimate for strictly nondurables is about -.04, representing a decline in strictly nondurable consumption of about 4 cents for each dollar of college expenditure. The coefficients are again relatively precisely estimated.

In sum, there is some weak indication of a delayed decrease in consumption in the winter, but even so the magnitudes are rather small, less than 5% of college expenditures. Overall, the null hypothesis is still not rejected. It appears that by the start of the academic year households have savings or other resources to finance up to 6 months of college expenditures without much decrease in consumption. It remains possible that households' consumption decreases only after this half-year. But Table 2 suggests that

---

[25]The same covariates are used, though now the changes in family size are for periods t+2 versus t. The demographics recorded in the Table are for the case of strictly nondurables. Those for nondurables and food are comparable and not remarkable.

consumption does not then further decrease with any new round of college expenditures the following fall, because such rounds were already included in the sample.[26]

Consider next the time before the spring/summer. Column (2) presents the main results from regressing $C_{t+1} - C_{t-1}$ on $COLEXP_{t+1}$. That is, it compares consumption in the fall to consumption in the previous spring and winter. This tests whether consumption is cut in the spring, perhaps to "crash save", as by household #1 in Figure 3. The coefficients are broadly similar to those of the base case, in Table 2. These results count against crash saving: insofar as savings are being used to maintain consumption in the fall, this saving appears to be fully underway at least 6 to 9 months in advance.

The results of these first two subsections are consistent with the life-cycle theory. Households appear to be doing a rather good job maintaining their standard of living into the academic year, despite large expenses. Further, households do not seem to cut their consumption in the 6 to 9 months before the year to augment the resources available to them to meet these expenses. For households which rely on savings to pay, the implication is that this saving is fully underway before these months.

The final subsection isolates the **starting** households, which are most likely to be sending children to college for the first time, to see how good a job they do.

3. Starting Households.

Table 4 reports the main coefficients on adding a separate slope and intercept term for the approximately 180 starting households. For the households that are not starters, the coefficient for nondurables is still positive, and larger than in the base case and now significant at the 5% level. However, the coefficient again loses significance and falls near zero on moving to the preferred consumption group, strictly nondurables.

---

[26]With COLEXP correlated across years, even a deterioration of consumption in the spring should have been detected in $C_{t+1} - C_{t-1}$ for the following year, and then should have shown up in column (2). Eg, if COLEXP in the fall of 1990 is depressing consumption in the spring of 1991, consumption in the fall of 1991 minus consumption in the winter of 1990 / spring of 1991 will be negatively correlated with COLEXP in the fall of 1991.

For starting households, the coefficient for nondurables is significantly smaller than for non-starters, by about -.26. This is especially interesting if one expects that any non-separability between COLEXP and nondurables be greater for starting households. (Eg, households might be more likely to buy new sets of clothing for their freshmen than for their sophomores.) Starters' nondurable consumption is then decreasing in COLEXP despite any such nonseparability. The total response is .12 - .26 = -.14, ie starting households' consumption in the fall drops by 14% of their college expenditures. For strictly nondurables, the coefficient for starting households is marginally significantly different, at the 10% level. The total response is .02 - .09 = -.07.

The intercept for starting households is positive for nondurables, so it does not follow that starting households spend less in total than other households.[27] But recall that this paper is not about total spending; the variation exploited is in the magnitude of COLEXP. The intercepts notwithstanding, the nondurable and strictly nondurable consumption of starting households appear to decrease with their college expenditures, counter to the life-cycle theory.

In sum, there is evidence that households have some trouble maintaining their consumption in the face of college expenditures on first sending children to college. This supports the view that they learn to save through experience. But even for these households the net effect is small, less than 15% of college expenditures.

---

[27]At the median college expenditure for starting households, still about $500, the resulting decrease in nondurable consumption is about (.14)(500), or about $75, for the quarter. At the mean expenditure, the decrease is about double this, still less than the about $200 of the intercept.

It is hard to evaluate total spending without taking into account durables and seasonality (via the month dummies). Unfortunately, there are not enough starting households to interact the starting dummy variable with all the month dummies (nor usefully to examine the alternative timings of the previous subsection).

V. Conclusion.

The main finding of this paper is that households appear to maintain their consumption up to 6 months into the academic year, despite large college expenses. Further, households appear not to cut their consumption in the 6-9 months before the year starts. For households financing college out of savings, the saving seems to have been fully underway at least 6-9 months in advance of the expenditure.

There is, however, evidence of a drop in nondurable consumption in proportion to college expenditures for households with children first beginning college, suggesting that such households learn to save with experience. But even for such households the effect on consumption is rather small in magnitude.

The null hypothesis that households are behaving in accord with the life-cycle theory cannot be rejected. They appear to be saving or borrowing enough to sustain their consumption in the face of large, predictable decreases in their net disposable income.[28] As already suggested, the test here should be relatively powerful in detecting violations of the theory, since the regressor has large variation and is of good quality. It is important to note, however, that these households might still have sacrificed quite a bit to pay for college. Even if the slopes of their consumption-paths are not much distorted, the levels might have been rather low in order to accumulate sufficient savings.

On the other hand, one might expect households to do a relatively good job, according to the theory, in the particular case of paying college. They have many years in which to save up in advance of college, and often have some access to guaranteed loans.[29] Moreover, the financial aid system tailors college expenditures to households'

---

[28]It is also worth noting preliminary but similar results on households' ability to smooth their consumption in the face of large tax payments. This is in contrast to significant increases in consumption in response to tax refunds. See Souleles [1995] (Chapter 2), and contrast Shea [1995].

[29]Though recall that many households are at the borrowing-limit of the guaranteed loans. There is also some concern that many households, especially minorities, are "unduly" hesitant to take out loans. (Miller and Hexter [1985, 1985a])

ability to pay. There is also a role for selection: the households which lack the resources are more likely not to undertake large college expenditures.

This paper might usefully be extended in a number of areas. First, it would be interesting to examine the response of consumption at other periods during and around the time of college. Unfortunately the CEX follows households for only a short time, necessarily limiting the horizon that can be examined. But even with a short horizon, periods other than the start of the academic year might be further considered. There was some weak but suggestive evidence presented in Section IV of a delayed response of consumption during the late fall and winter, for instance. Such dynamics of consumption deserve additional attention. Of particular interest might be the end of college, to see whether consumption rebounds after the burden of college expenditures has been lifted; or whether, eg, the need to save persists with retirement around the corner.

Related examinations might be undertaken using data on assets, as well, especially in data sets allowing a longer horizon. Returning to the last example, one might investigate whether households which have used up their early savings for college are able to restore their savings in time for retirement. Or, given the evidence above that many households must be seriously saving at least 6-9 months in advance of college, one might try to determine how early they started seriously saving.

A second extension might consider the effect of college expenditures on household portfolios. Eg, one might identify the types of assets households draw on to pay for college, and in particular the role of loans. Also, there might be interesting effects of these payments, and sending children to college generally, on durables purchases. A final extension might investigate more generally the effects on life-cycle savings and consumption of other consumption "needs", such as other aspects of child-rearing or other large expenditures like cars or homes.[30]

As an aside, note that households also have many years to save in advance of retirement.

[30]Eg, Englehardt [1993] examines the effects of saving up for a down-payment on a house.

VI. Appendix: the Data.

This appendix focuses on the education variables. Additional information on the other variables can be found in the appendix to Chapter 2.

1. College Expenditure.

For college students living at home, COLEXP includes all payments for college expenses made by the household.[31] College students living away, by contrast, are not counted as part of the "consumer unit", and so are dropped by the BLS from the dataset. As seen below, this complicates the identification of such college students. Even so, any payments for such students that a household makes directly to a college appear in COLEXP.[32] But regular cash contributions directly to the student do not. At the quarterly level, these cannot be separated from other contributions, mostly to charities, so are not used in the analysis. They are also recorded separately, but then only annually, so again are not used.

Preliminary analysis uncovered some anomalies in the routines the Census uses to clean the variable for the highest grade attended, for teen-aged household members. Census flags, kindly provided by the BLS, were used to undo most of these anomalies, but some remain. In response, this variable is used only in conjunction with the variables for age and college enrollment.

"Traditional" collegiate households may not have any member over 24 who is enrolled in college or not working for reasons of studies. A child is counted as a candidate to be starting college if it has age in [16,19] and is in the senior year of high

---

[31]Recall that about 30% of white undergraduates, and 35-50% of minority ones, live at home while in college. (Churaman[1992]).

[32]These expenditures are then supposed to count as "gifts". It is unclear whether they are generally correctly so identified, or often left with the "non-gifts". The magnitudes of gifts are on average rather larger than those of non-gifts. Interacting COLEXP with a dummy variable for being a gift did not lead to a significant difference for gifts in the main results, so the distinction was dropped.

school before the fall, and appears to be a cause of all college payments made in the fall. Essentially, this is if it has highest grade 12, but is not enrolled, before the fall; and then has highest grade 13 or is enrolled or leaves the household, in the fall. A household with such a candidate child is considered a "starter" if it does not have college expenditures in any interview before the fall, and if it has no other ("non-starting") members in college. Generally, other members are considered to be in college if a) they are enrolled in college with any age over 19, or are aged 17-19 with grade over 13; or b) they have highest grade over 13, are of typical age for someone in that grade, and that grade increases in the next interview or the member is missing from the household during the spring or fall. Recall that the CEX does not keep track of children after leaving the household, so one cannot tell whether a given household sent older children to college years ago. But as just described one can tell whether children still members of the household are in college, and whether the household is currently paying for college (whether the student is in the household or not). At this point it is worth noting that misclassification of starting households attenuates any difference between starters and other households.

2. Consumption and Other Variables.

To ensure the adequacy of the consumption data, observations are dropped if there are other consumer units in the same household as the family of interest, if the household lives in student housing, or if the head's occupation is farming/fishing.

In aggregating expenditures into the groups food, strictly nondurables and nondurables, if any component of a group was topcoded or missing its cost, the whole group was set to missing. If any component of a group was missing its date or dated before the reference period, that group was dropped for all interviews for the household at issue. By contrast, since COLEXP is never differenced, topcoded components in it were retained. Also, all consumption groups are dropped if the household lacks food expenditures in any month of the quarter. A non-negligible number of expenditures are dated in the month of the interview, i.e. after the reference period. Following the recommendation of the staff at the BLS, for consistency such expenditures were accrued to the following reference period.

For married households with female respondents, the demographic variables specifically about the respondent were switched to refer to the male, identified as the head of household. A household is dropped if the age of any of its members decreases from quarter to quarter, or if it increases by more than one, or if it is otherwise flagged. A handful of other adjustments to the primitive data were made, according to corrections provided in the CEX documentation for various years.

Bibliography

Attanasio, O.P., et al., 1994, "Dynamic Consumption and Saving Behavior in the US and the UK," mimeo.

Carroll, C.D., and Summers, L.H., 1991, "Consumption Growth Parallels Income Growth: Some New Evidence," in Bernheim, B.D., and Shoven, J.B., eds., National Saving and Economic Performance, Chicago: Chicago University Press for NBER, pp. 305-343.

Case, K.E., and McPherson, M.S., 1986, "Does Need-Based Student Aid Discourage Saving for College?," College Entrance Examination Board, January.

Choy, S.P., and Henke, R.R., 1992, "Parental Financial Support for Undergraduate Education," for the National Center for Education Statistics, U.S. Department of Education, May.

Churaman, C., 1992, "Financing of College Education by Minority and White Families," Journal of Consumer Affairs, Winter, pp. 324-350.

Edlin, A.S., 1993, "Is College Financial Aid Equitable and Efficient?," Journal of Economic Perspectives, Spring, pp. 143-158.

El-Khawas, E., 1985, Campus Trends: 1984, Washington: American Council on Education.

Englehardt, G.V., 1993, "Down Payments, Tax Policy, and Household Saving," M.I.T. Thesis, February.

Feldstein, M., 1992, "College Scholarship Rules and Private Saving," NBER working paper #4032, March.

Gale, W.G, and Scholz, J.K., 1994, "Intergenerational Transfers and the Accumulation of Wealth," Journal of Economics Perspectives, Fall, pp. 145-160.

Lusardi, A., 1992, "Permanent Income, Current Income and Consumption: Evidence from Panel Data," CentER for Economic Research, Tilburg University, mimeo, December.

Lusardi, A., 1993, "Euler Equations in Micro Data: Merging Data from Two Samples," CentER for Economic Research, Tilburg University, mimeo.

McPherson, M.S., and Schapiro, M.O., 1991, Keeping College Affordable: Government and Educational Opportunity, Washington: the Brookings Institution.

Miller, S.E, and Hexter, H., 1985, "How Low-Income Families Pay for College," American Council on Education, July.

Miller, S.E, and Hexter, H., 1985a, "How Middle-Income Families Pay for College," American Council on Education, July.

Modigliani, F., 1988, "The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth," Journal of Economic Perspectives, Spring, pp. 15-40.

Mumper, M., 1993, "The Affordability of Public Higher Education: 1970-1990," The Review of Higher Education, pp. 157-180.

Roper Organization, 1984, "A National Study on Parental Savings for Children's Higher Education Expenses," conducted for the National Institute of Independent Colleges and Universities, January.

Shea, J., 1995, "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis," American Economic Review, March, pp. 186-200.

Souleles, N.S., 1995, "The Response of Household Consumption to Income Tax Refunds," M.I.T. mimeo, January.

Zeldes, S.P., 1989, "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, pp. 305-346.

Department of Education, 1993, Digest of Education Statistics, National Center for Education Statistics, NCES93-292.

General Accounting Office, 1990, "Higher Education: Gaps in Parents' and Students' Knowledge of School Costs and Federal Aid," PEND-90-20BR.

FIGURE 1

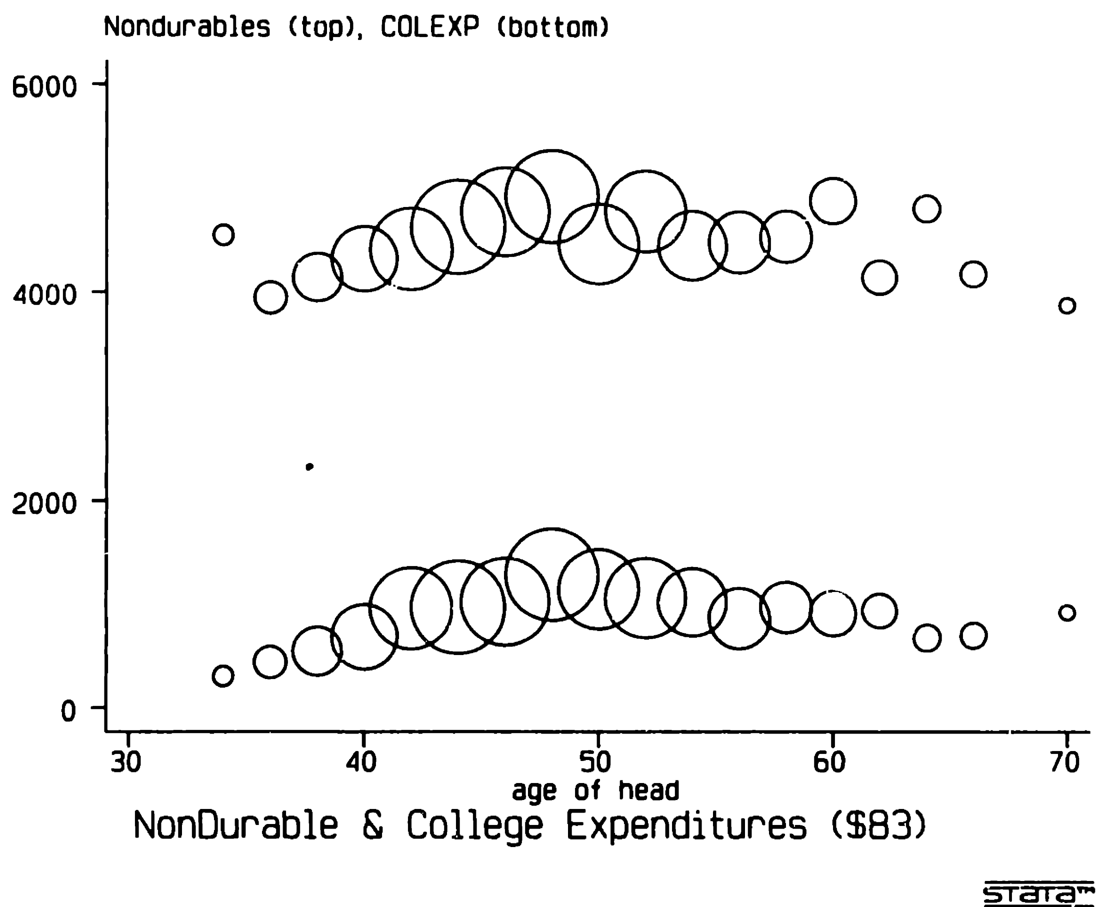
Nondurables (top), COLEXP (bottom)
6000
4000
2000
0
30
40
50
60
70
age of head
NonDurable & College Expenditures ($83)
stata™

FIGURE 2: Timing of Periods

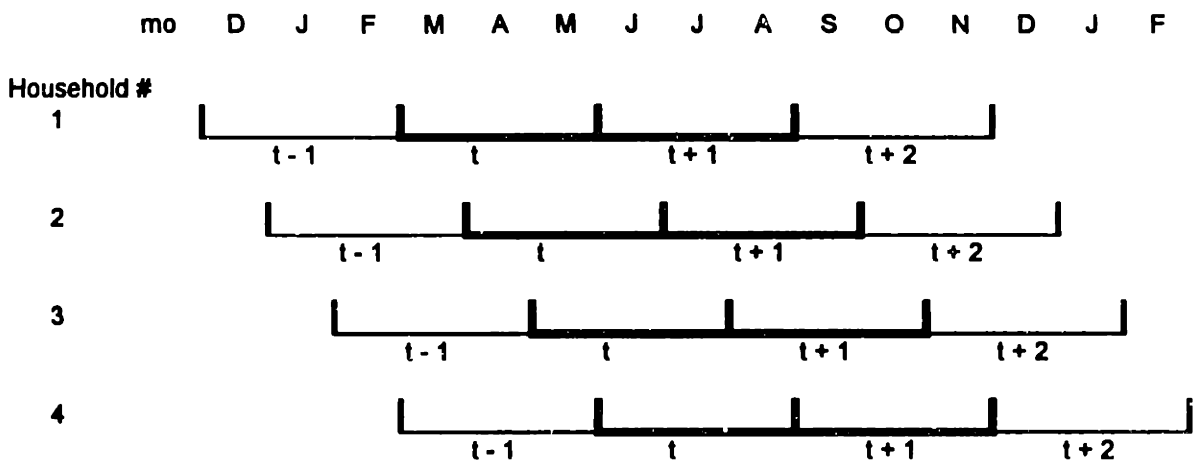

FIGURE 3: Crash Saving

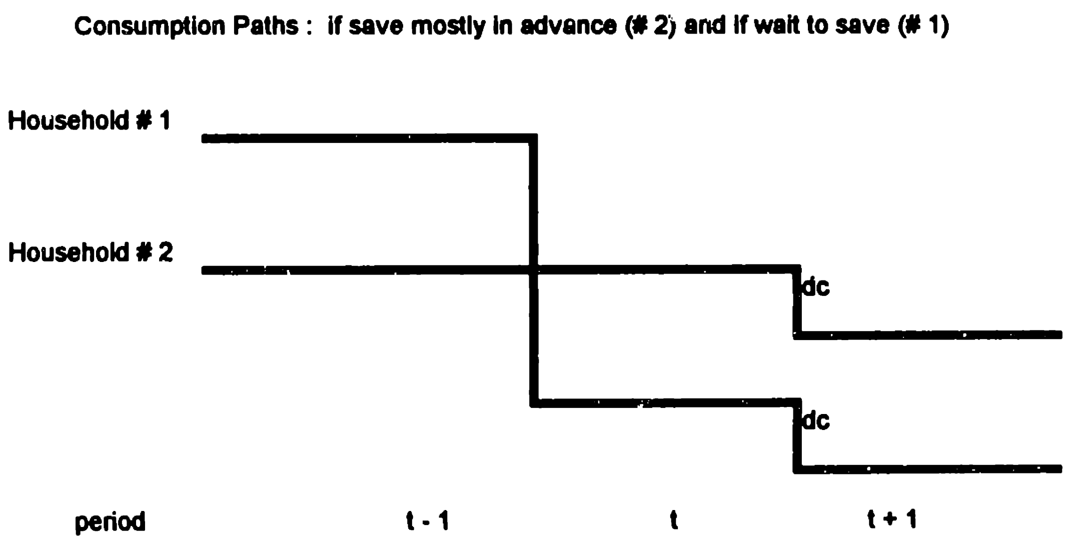

TABLE 1: Undergraduate Enrollment and Costs

Type of Institution	Num. of Undergraduates (% of total) (1991)	avg. annual Tuition (1992-3)	avg. annual Total Costs (1992-93)
Private, 4 year	2.3 million (19%)	$10, 393	$15, 128
Public, 4 year	4.7 million (38%)	2, 352	6, 029
Public, 2 year	5.4 million (44%)	1, 018	-

Notes:

1) Source: DOE[1993].

2) Total costs are in current dollars and include tuition, fees, and room and board assuming full-time enrollment. Tuition is the in-state charge.

TABLE 2: Base Case, Equation (2)

	Nondurables	Strictly Nondurables	Food
Age	10.1 (6.7)	5.0 (4.8)	3.5 (3.4)
d(ages 0-15)	-91.1 (169.5)	-170.5 (114.6)	-102.7 (80.7)
d(ages 16-24)	191.0 * (107.3)	142.9 * (73.5)	98.0 * (57.5)
d(ages 25-90)	769.0 * (429.9)	763.0** (332.5)	600.2 * (348.6)
COLEXP	.085 * (.048)	.012 (.029)	-.009 (.016)
# obs	1047	1065	1065

Notes:

1) month dummies not shown

2) Standard errors corrected for heteroskedasticity

TABLE 3: Alternative Timing

Coefficient on COLEXP by consumption group	(1)	(2)
Nondurables	-.034 (.049)	.060 (.051)
Strictly Nondurables	-.043 (.034)	.038 (.035)
Food	-.023 (.017)	-.001 (.023)
Coefficients on demographics for strictly nondurables		
age	.29 (5.23)	14.79** (6.13)
d(ages 0-15)	3.36 (131.98)	169.7 (133.0)
d(ages 16-24)	230.4** (93.)	114.0 (87.1)
d(ages 25-90)	832.4** (154.3)	336.7 (247.7)
# obs	609	676

Notes:

1) Column (1) is for C(t+2) - C(t) on COLEXP(t+1)
Column (2) is for C(t+1) - C(t-1) on COLEXP(t+1)

2) The coefficients for the demographics in the second panel are for the case of strictly nondurables. Month dummies are not shown.

TABLE 4: Equation (2), with Starting Households

	Nondurables	Strictly Nondurables	Food
COLEXP	.123** (.050)	.024 (.033)	-.008 (.018)
COLEXP * starter	-.258** (.086)	-.087 * (.050)	-.012 (.028)
starter (constant)	206.9 (137.6)	55.3 (88.8)	-64.9 (57.0)
# obs	1047	1065	1065
# starters	180	183	183

Notes:

1) Month dummies and demographics, as in Table 2, not shown.

2) Standard errors corrected for heteroskedasticity

Chapter 4. Conclusion.

Each of the previous three chapters investigates the relationship of household consumption and saving to income. This brief concluding chapter aims to draw some general, somewhat speculative conclusions from their results. The focus is on the role of liquidity constraints and of non-canonical preferences.

Chapters 2 and 3 investigate particular episodes of increases and decreases in disposable income, due to tax refunds and college expenditures respectively. The main test of Chapter 3 finds that households paying for college appear to have sufficient savings or other resources to maintain their consumption despite their large payments. This result is taken to be relatively powerful evidence in favor of the canonical life-cycle theory. Since disposable income is decreasing in this test, the result is free from complications due to liquidity constraints and so is evidence for the "normalcy" of preferences. On the other hand, due to the limitations of the data the time-period covered by the test is short, and so the possibility remains that households are violating the theory outside this period. Or, the result might not generalize; households might be doing a better job--according to the theory--in the particular, low-frequency case of paying for college than in other cases. Shea [1995], for instance, finds the consumption of union members more sensitive to decreases in their wages than to increases.[1] Additional investigation of other episodes in which income decreases would be helpful.[2]

[1]Though the use of wages, as opposed to income, might be affecting his result.

[2]As an extension to Chapter 2, I am also investigating the response of consumption to tax payments. Preliminary analysis has found little excess sensitivity, consistent with

Chapter 2 does find evidence of liquidity constraints in the case of refund-receipt: the nondurable consumption of those most likely to be constrained, i.e. those with relatively low income or wealth, appears to increase with their refunds. However, more than liquidity constraints seem to be at play, because the consumption of those most likely unconstrained, those with relatively large income or wealth, also appears to respond, though primarily in durables. Moreover, the response of nondurable consumption seems to extend late into the year, past the time of most refund-receipt.

Chapter 1, by contrast, takes two different and novel tacks to investigate the role of liquidity constraints. First, instead of focusing on households with low income or wealth, who are not exactly those the theory would classify as liquidity constrained, it uses direct information on borrowing constraints to estimate a switching regression model of the Euler equation. This model does not yield much excess sensitivity associated with the possibility of liquidity constraints. However, the model is somewhat imprecisely estimated, in part because relatively few households are credit-constrained and the information on constraints comes from an outside data source and so has to be matched statistically to the households in the data sets with the consumption data. The second tack is to examine the effect of liquidity constraints on the entire conditional distribution of consumption growth, not just on the conditional mean that appears in the Euler equation. This distribution turns out to look quite different for the constrained and the unconstrained regimes,[3] which is some evidence that liquidity constraints are at work.

In sum, the previous chapters find evidence that liquidity constraints affect consumption, but nonetheless conclude that more than liquidity constraints are at play. Chapter 2 explicitly investigates some of the other possibilities, in particular the leading behavioral theories of saving. The sharp response of durables to refunds might be explained in part by models of self-control/forced saving. But the sharpest response comes from households with the greatest income and wealth, who could most easily undo

---

Chapter 3.

[3]Constrained and unconstrained here are still with respect to the direct indicator of borrowing constraints.

the saving forced by overwithholding, e.g. simply by writing a check. Such models might also have difficulty explaining the delayed response of consumption after refund-receipt. Nonetheless, it remains possible that people think of refunds as a special kind of income, as under theories of mental accounts.[4] One can tell plausible stories in which people do not undo the forced saving of refunds or in which people spend their refunds late in the year. The usual response to such a mental accounts view, however, is a reductio: one could tell such stories to explain most any results. Also, that the results in Chapter 2 do not much change on including the control group (those without refunds) is some evidence that refunds are not an unusual kind of income.

Finally, Chapters 2 and 3 evidence the richness of preferences, which can complicate the relationship of consumption and saving to income.[5] First, Chapter 3 notes that paying for college is just one example of "consumption needs," like child-rearing more generally or buying a house. Such needs are usually incorporated in the canonical model as changes in the marginal utility of consumption. Unfortunately, such changes are hard to control for. The response of Chapter 3 is to focus on one particular life-cycle episode, but even this control is no doubt imperfect. Second, there are hints in Chapter 3 of possible non-separabilities between expenditures on college and other nondurable expenditures. More generally, non-separabilities between consumption goods, or between consumption and labor supply, can make consumption appear sensitive to income (Meghir and Weber [1993], Attanasio and Weber [1992, 1994]). The previous chapters use broad groupings of consumption, like nondurables, to minimize any complications due to the first kind of non-separability. As for the second, Garcia, Lusardi, and Ng [1994] find that excess sensitivity persists even after controlling for labor supply. Also, note that the main "story" for non-separability from labor supply does not well apply to the case of tax

---

[4]Recall that Chapter 2 finds some evidence counter to one prominently expressed implication of mental accounts: the response of total consumption is found to be increasing in the size of the refund. But this evidence need not count against the general thrust of the mental accounts view.

[5]There are of course other possibly important complications that do not receive much attention in this thesis, e.g., precautionary motives and intertemporal dependencies.

refunds, because contemporaneous changes in labor supply cannot affect predetermined refunds, unlike earnings. Finally, in both Chapters 2 and 3 the seasonality of preferences is quite important, since spending on both education and trips varies greatly with the season. The chapters respond with month dummies and control groups, but again these controls are no doubt imperfect. In conclusion, the quantitative importance of these three complications in preferences remains an open issue, requiring additional study.

Bibliography

Attanasio, O.P., and Weber, G., 1992, "Consumption Growth and Excess Sensitivity to Income: Evidence from U.S. Micro Data," Stanford University, mimeo, August.

Attanasio, O.P., and Weber, G., 1994, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," NBER working paper, July.

Garcia, R., Lusardi, A., and Ng, S., 1994, "Excess Sensitivity and Asymmetries in Consumption: An Empirical Investigation," University of Montreal, mimeo, June.

Meghir, C., and Weber, G., 1993, "Intertemporal Non-Separability of Borrowing Restrictions? A Disaggregate Analysis using the US CEX Panel," Institute for Fiscal Studies, mimeo, June.

Shea, J., 1995, "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis," American Economic Review, March, pp. 186-200.