

**Unemployment Risk and Precautionary Wealth:
Evidence from Households' Balance Sheets**

Christopher D. Carroll
Department of Economics
The Johns Hopkins University
3400 Charles Street
Baltimore, MD 21218
(410) 516-7602
ccarroll@jhu.edu

Karen E. Dynan
Mail Stop 80
Federal Reserve Board
Washington, DC 20551
(202) 452-2553
kdynan@frb.gov

Spencer D. Krane
Mail Stop 80
Federal Reserve Board
Washington, DC 20551
(202) 452-3702
skrane@frb.gov

April 1999

Abstract

Recent empirical work on the strength of precautionary saving has yielded widely varying conclusions. The mixed findings may reflect a number of difficulties in proxying uncertainty, executing instrumental variables estimation, and incorporating theoretical restrictions into empirical models. For each of these problems, this paper uses existing best-practice techniques and some new strategies to relate unemployment probabilities from the *Current Population Survey* to net worth data from the *Survey of Consumer Finances*. We find that increases in unemployment risk do not boost saving by households with relatively low permanent income, but that a statistically significant precautionary effect emerges for households at a moderate level of income. This finding is robust to certain restrictions on the sample, but not robust across measures of wealth: We generally find a significant precautionary motive in broad measures of wealth that include home equity, but not in narrower subaggregates comprising only financial assets and liabilities.

We are grateful to Dan Bergstresser, Martin Browning, Eric Engen, Steve Lumpkin, Martha Starr-McCluer, Valerie Ramey, and seminar participants at the American Economic Association Annual Meetings, Johns Hopkins University, the NBER Summer Institute, and Georgetown University for helpful comments. We also thank Dan Bergstresser and Byron Lutz for excellent research assistance and Arthur Kennickell, Martha Starr-McCluer, and Gerhard Fries for help with the SCF. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or its staff.

I. Introduction

I.1. Overview

Many recent studies have noted the potential economic importance of precautionary saving. Caballero (1990) and Normandin (1994) have pointed out that precautionary saving may be able to explain certain stylized facts about aggregate consumption such as its excess sensitivity to movements in income. Carroll (1992) and Carroll and Dunn (1997) have argued that precautionary behavior is an important driving force for consumption-led business cycles. And simulations in Hubbard, Skinner, and Zeldes (1994) suggest that precautionary saving could account for almost half of the aggregate capital stock. Yet the empirical evidence regarding precautionary saving is mixed: Kuehlwein (1991), Dynan (1993), Guiso, Jappelli, and Terlizzese (1992) and Starr-McCluer (1996) find little or no precautionary saving, whereas Carroll (1994), Carroll and Samwick (1997, 1998), Engen and Gruber (1997), and Lusardi (1997, 1998) find evidence of a significant precautionary motive.¹

The mixed findings may reflect a number of difficulties in testing for precautionary saving. The problems fall into three general categories: the method of proxying uncertainty, the instrumental variables strategy, and the incorporation of restrictions and insights provided by a theoretical model. In each of these categories, this paper either builds on best-practice techniques and other insights from the existing literature or brings to bear new strategies.

I.2. Proxying Uncertainty

Precautionary wealth is defined as the difference between the wealth that consumers would hold in the absence of uncertainty and the amount they hold when uncertainty is present (Kimball 1990). However, the most appropriate empirical measure of uncertainty is not obvious. Many previous studies have proxied uncertainty with either the variability in a household's income (Carroll, 1994 and Carroll and Samwick, 1997, 1998) or the variability in its expenditures (Dynan, 1993, and Kuehlwein, 1991). But, as Guiso, Jappelli, and Terlizzese and Lusardi (1997, 1998) have pointed out, variability measures may be poor proxies for uncertainty because they can contain large controllable elements in them. For example, a

¹The coefficient of relative risk aversion estimated by Gourinchas and Parker (1997) also implies a substantial precautionary response.

tenured college professor who, by choice, works only every other summer may have much more variable annual income than a factory worker, but does not face the uncertainty of being laid off during a recession. Similarly, differences in the variation in quarterly expenditures between two households may simply reflect differences in the families' preferences towards regular seasonal outlays such as summer vacations or school tuition.

Our measure of uncertainty is the probability of job-loss. Specifically, we estimate the probability that a consumer who currently is employed will be unemployed one year hence.² Future job loss represents a potential major interruption to income over which households generally have little influence, and thus it should provide a much cleaner signal of the uncertainty faced by a household than variation in income or expenditures.

Unfortunately, no single source contains high-quality information on household-level income, wealth, and job-loss risk. Our solution is to use a source of good data on employment and unemployment, the *Current Population Survey (CPS)*, to estimate job-loss risk based on observable household characteristics. We then take the results from the CPS estimation and apply them to predict job-loss risk for households in a data set that contains good information on income and wealth, the *Survey of Consumer Finances (SCF)*. Finally, we relate the resulting predicted unemployment risk to household net worth.

1.3. Instrumental Variables Strategy

Because uncertainty is measured with significant error, most studies instrument for their uncertainty proxy using variables such as the consumer's occupation, education, industry of employment, and demographic characteristics. Econometric identification requires that at least one instrument be related to the dependent variable (wealth, in our case) solely through that instrument's correlation with uncertainty; this instrument can then legitimately be excluded as an independent variable in the second-stage regression of wealth on instrumented

²Lusardi (1998) and Engen and Gruber also use measures of the probability of job loss in their analyses. Lusardi finds significant precautionary wealth accumulation using the household's reported perception of job-loss risk. Engen and Gruber find that the effect of unemployment insurance on precautionary wealth is significantly more pronounced at higher unemployment rates.

uncertainty.

Finding an appropriate instrument to exclude is problematic. For example, suppose that more risk averse consumers both hold more precautionary wealth and choose occupations with lower job-loss risk.³ Then occupation may be a good predictor of job-loss risk, but, if it is excluded from the second-stage regression, the coefficient estimate on the uncertainty variable will be biased because of the correlation between instrumented job-loss risk and the unmeasured risk-aversion portion of the error term. Similar arguments can be made for educational choice and industry choice, and we find some empirical evidence that these concerns may be warranted.⁴

To avoid this identification problem, we include all of the usual variables (occupation, education, etc.) that are used as instruments for uncertainty as independent controls in our second-stage equation. This requires us to find some other instrument that is correlated with job-loss risk and can be excluded from the second-stage regression. We use the region where the household resides. The large variation in regional economic conditions suggests that region will be significantly correlated with an individual's job-loss risk. In addition, based on the assumption that, *ex ante*, most households do not choose where to live on the basis of regional differences in job-loss risk, region should be uncorrelated with unobserved risk-related determinants of wealth.⁵

1.4. Insights From a Structural Model

³Lusardi (1997) emphasizes the link between occupational choice and risk aversion in the 1989 Italian *Survey of Household Income and Wealth*, in which one-half of households mention job security as a reason for choosing their jobs.

⁴A related problem occurs in the branch of the literature that proxies uncertainty with insurance coverage: Risk-averse households may both save more and obtain more insurance, biasing the IV coefficient estimates. In a study focusing on households' health insurance coverage, Starr-McCluer addresses this problem by instrumenting coverage with the percent of the local workforce employed by large firms. Elsewhere, Engen and Gruber consider unemployment insurance coverage, which is determined by state policy makers and thus probably largely exogenous to the individual household's saving behavior.

⁵Engen and Gruber provide a forceful discussion of issues concerning instrument validity; they also argue that regional variables are good candidates for instrumenting uncertainty because they likely satisfy exogeneity requirements.

We solve a theoretical model of precautionary saving, which implies important restrictions on the empirical work. First, because a spell of unemployment causes a household to run down precautionary balances, and because high-risk households are more likely to have recently experienced unemployment spells, at a given point in time we may observe high-risk households holding *less* wealth than their low-risk counterparts, even though high-risk households would hold higher balances in the steady state. Thus, our specification controls for recent shocks that may have depleted precautionary reserves.

Our model also provides guidance on how to transform wealth data. To deal with the extreme skewness of the wealth distribution, many empirical papers have used the logarithm rather than the level of wealth as the dependent variable. Of course, this requires dropping or making ad hoc adjustments to observations with nonpositive wealth. However, a non-trivial proportion of households hold zero or negative net worth--a fact that we show is consistent with optimizing behavior in our model because of the existence of unemployment insurance and differential borrowing and lending rates. To avoid eliminating these households from our empirical work, we instead transform wealth using the inverse hyperbolic sine function suggested by Burbidge, Magee and Robb (1988).⁶ Like the log, the inverse hyperbolic sine downweights large values, but unlike the log it can be applied to positive, zero, and negative numbers; it also allows elasticities to vary with the level of wealth, another property of precautionary behavior implied by standard theoretical models like ours but not allowed by a log transformation.

1.5. Results

Our empirical results provide some support for the proposition that precautionary saving is important. We find that increases in unemployment risk do *not* cause households with relatively low permanent income to significantly boost their net worth, but that a statistically significant and economically sizable precautionary effect emerges for households at moderate and higher levels of income. These results are robust to a number of changes in the specification, but not across subcomponents of wealth: We generally find a significant precautionary motive in broad measures of wealth that include home equity, but not in

⁶We thank Martin Browning for suggesting this transformation.

narrower subaggregates comprised only of financial assets and liabilities. We discuss a number of potential explanations for these findings, both within and outside of the context of precautionary saving.

II. A Model

In this section, we present and simulate a model of household behavior as a way to clarify the relationship between net worth and unemployment risk. To isolate precautionary responses, we consider a model in which there is no saving for retirement or other purposes. We view this stylized model as a qualitative guide to the response of net worth to changes in employment risk and to the dynamics of wealth for households that suffer spells of unemployment.

II.1. The Household's Problem

We assume that the household's problem at time t is to

$$\max_C \left(u(C_t) + \sum_{j=t+1}^{\infty} \beta^{j-t} E_t[u(C_j)] \right) \quad (1)$$

subject to

$$X_{t+1} = R(X_t - C_t) + Y_{t+1}, \quad (2)$$

where C_t represents consumption, X_t is the total resources available to the household at time t , and Y_t is income (which is assumed to be received at the beginning of the period). β^{-1} is one plus the rate of time preference and R is one plus the (constant) rate of return, r .

The utility function is of the constant relative risk aversion (CRRA) form,

$$u(C_t) = \frac{C_t^{1-\gamma}}{(1-\gamma)}, \quad (3)$$

where γ is the coefficient of relative risk aversion. This utility function has a positive third derivative, and thus a mean-preserving spread in consumption uncertainty raises expected

marginal utility. Because households want to smooth expected marginal utility over time, such an increase in risk will cause a household to reallocate resources from consumption today to a precautionary reserve that partially insures consumption tomorrow against potential negative draws of income (see Kimball).

We assume that the household's income, Y_t , equals a permanent component, Y_t^P , multiplied by a transitory shock V_t :

$$Y_t = Y_t^P V_t . \quad (4)$$

The permanent component evolves as:

$$Y_{t+1}^P = G Y_t^P N_{t+1} , \quad (5)$$

where N_{t+1} is a serially uncorrelated shock to permanent income that is lognormally distributed with variance σ_n^2 , so that the $\log(Y_t^P)$ is a random walk with drift $g = \log(G)$.⁷

We will concentrate on the transitory shock to income, V_t , which captures both relatively small year-to-year fluctuations in wages and occasional large drops corresponding to periods of unemployment. Specifically, we assume that with probability ρ the household is unemployed and $V_t = V_{min}$ where V_{min} captures, in a simple way, the “safety net” provided by formal and informal insurance markets.⁸ With probability $(1 - \rho)$ the household is employed, in which case V_t is lognormally distributed with variance σ_v^2 and mean $(1 - \rho V_{min})/(1 - \rho)$. (This last assumption ensures $E_t V_{t+1} = 1$, so that changes in the probability of unemployment

⁷ Note that in this model permanent income is not the discounted value of future cash flows, but rather the path around which earnings exhibit transitory fluctuations, consistent with Friedman's (1957) original definition. Thus, while V_t affects cash flow, it does not influence Y_t^P .

⁸Although very simple, this characterization of social insurance is sufficient for our model to mimic the aspects of the distribution of wealth that we are interested in. See Hubbard, Skinner and Zeldes (1995) for a more sophisticated formulation of social insurance.

affect the variance but not the expected value of income.)

End-of-period wealth is defined as the difference between resources and consumption,

$$W_t \equiv X_t - C_t. \quad (6)$$

A nontrivial fraction of actual households have very low or even negative levels of wealth: For example, in the 1983, 1989, and 1992 waves of the SCF, 8 to 10 percent of households have net worth less than one month's income and 4 to 6 percent of households have negative net worth. If $V_{min} = 0$, no optimizing household would hold zero or negative net worth because of the possibility that $C_t = 0$ and $U'(C_t) > -\infty$. (See Zeldes, 1989, for a more detailed discussion of this mechanism.) Thus, in our model, $V_{min} > 0$ is necessary to characterize the distribution of wealth. Further, to ensure that optimizing behavior produces the clustering of net worth around zero observed in the data, we assume that consumers face different interest rates for borrowing and saving. In particular, consumers ending period t with positive wealth earn a return that applies to lenders, R^{lend} , while those who choose to borrow (end the period with $W_t < 0$) will pay out at a higher rate, R^{borrow} . For lenders, the Euler equation for consumption will be

$$u'(C_t) = \beta R^{lend} E_t \left[u'(C_{t+1}) \right]. \quad (7)$$

while for borrowers, the Euler equation is:

$$u'(C_t) = \beta R^{borrow} E_t \left[u'(C_{t+1}) \right]. \quad (8)$$

Because $R^{borrow} > R^{lend}$, there is a gap between the right-hand-sides of equations (7) and (8), so that there will be a range of X_t for which consumers choose $C_t = X_t$ and $W_t = 0$.

II.2. Parameterization

We select the time period to be a year and set $\sigma_n^2 = \sigma_v^2 = 0.01$ and $g = 0.03$; as discussed in Carroll (1992), these values are roughly consistent with evidence from the *Panel*

Study on Income Dynamics (PSID). Consumers are assumed to discount future utility at a rate of 4 percent per year and have a coefficient of relative risk aversion of 2. We assume the baseline probability of job loss, ρ , is 0.02.⁹ We set V_{min} equal to 0.2, implying that the social safety net ensures resources equal to 20 percent of normal income. Finally, we assume that the lending rate is zero and the borrowing rate is 20 percent--producing a gap between R^{borrow} and R^{lend} that is close to the difference between the after-tax rate of return on Treasury bills and (at least until recently) the interest rates on many credit cards. These parameter values ensure that households will not desire to accumulate assets without bound and thus that the model has a steady state.

II.3. The Consumption Function

The model is solved using standard numerical dynamic stochastic programming techniques. The ratio of resources to permanent income, $x_t = X_t/Y_t^P$, is a sufficient statistic for the ratio of consumption to permanent income, $c_t = C_t/Y_t^P$ (see Carroll, 1996). The solid line in figure 1 depicts the predicted relationship between x_t and c_t under our baseline parameter values.¹⁰ The figure shows that consumers with very low levels of resources--those to the left of the 45 degree line where $c_t = x_t$ --will consume more than their current resources. These will be consumers who have been hit with negative income shocks and are borrowing to smooth consumption through the rough patch; their behavior satisfies equation (8). Those to the right of the 45 degree line satisfy equation (7) and spend less than their current resources, leaving some wealth to both finance planned future consumption and to buffer future shocks. There also is a range--the segment that coincides with the 45 degree line--in which no choice of c_t satisfies either Euler equation because of the discontinuity between R^{lend} and R^{borrow} . In this range, the borrowing rate is high enough to prevent the consumer from going into debt, and the lending rate low enough that it is not worthwhile for the household to save. Hence,

⁹We purposely chose a probability lower than the actual unemployment rate in order to offset the fact that we have normalized spells of unemployment to last a full year; for comparison, Clark and Summers (1979) suggest that the average spell of adult unemployment (excluding people who leave the labor force) lasts between 3-1/2 and 4 months.

¹⁰Note that the minimum value of x equals the present discounted value of V_{min} .

consumption is equal to current resources.

The dashed curve in figure 1 shows the consumption function that would be optimal if the probability of job loss were 4 percent annually, rather than our baseline assumption of 2 percent. With the exception of a portion where $c_t = x_t$ in both regimes, the new consumption function lies below the old function, indicating that at most levels of wealth a consumer facing a 4 percent unemployment risk would consume less than one facing a 2 percent unemployment risk.¹¹ For households consuming less than their current resources, the precautionary response shows up as increased saving, while for households consuming more than their current resources, the precautionary response shows up as a reduction in borrowing.

II.4. The Distribution of Net Worth

How do these precautionary responses affect the observable variable in our data, the cross-sectional distribution of net worth? Given our other baseline parameters, the gap between R^{borrow} and R^{lend} is such that if a household faced no income uncertainty ($\rho = 0$), it would maintain wealth of exactly zero. But, as can be seen in the solid locus in figure 2 and the first column of table 1, which show the steady-state distribution of the ratio of net worth to permanent income under $\rho = 0.02$, income uncertainty causes households to carry some precautionary assets to serve as a buffer against potential bad draws of income in the future.¹² At the mean of the distribution, wealth is about 6-1/2 percent--or roughly one month's worth--of permanent income. Because this figure represents the change in average desired wealth due to the introduction of income uncertainty, it also can be interpreted as the model's average level of precautionary balances associated with a 2 percent chance of job loss. Note also that negative shocks to income leave a noticeable proportion of households with

¹¹ The entire shift in the consumption function reflects precautionary behavior because we adjust the mean of the transitory shock so that the change in ρ does not affect the expected level of income.

¹² We construct this distribution by simulating our model for a large number of consumers. Each starts with an initial endowment of wealth (zero). We then draw random income shocks for them according to the parameters specified above and calculate consumption and assets using the optimal consumption rule. We repeat this process for 10 periods, by which time the distribution has essentially stabilized at what we call the steady-state.

nonpositive net worth: Even with a 20 percent borrowing rate, optimal behavior implies that about 13 percent of households are in debt. And the difference between borrowing and lending rates is large enough that 11 percent of households have wealth of exactly zero.

The dashed line in figure 2 and the second column in table 1 show the steady-state distribution of wealth when the probability of job loss is 4 percent. Reflecting the precautionary response, this distribution lies to the right of the distribution when the unemployment rate is 2 percent. As shown in the third and fourth columns of table 1, the increase in the average wealth ratio amounts to about 5 percentage points--or 2.5 weeks--of permanent income; this additional precautionary saving represents an 80 percent increase in the average ratio of net worth to permanent income. As can be seen by comparing the responses across the percentiles of wealth-to-income, the dollar amount of extra saving (or reduced borrowing) rises somewhat with W_t/Y_t^p (column 3). But, because the bottom part of the wealth distribution is at such a low level of net worth, even a small absolute shift in the distribution represents a much larger *percentage* increase in wealth at the lower percentiles than it does at higher levels of wealth (column 4).

Finally, table 2 illustrates how experiencing a spell of unemployment affects net worth in our model. All households who are unemployed in the year of observation (column 2) have negative net worth; not only have they spent their precautionary reserves, they have borrowed an amount that averages one-third of permanent income.¹³ Furthermore, spells of unemployment leave lasting scars on households' balance sheets: Even two years after a spell of unemployment (column 4), households have substantially lower-than-average levels of wealth and significantly higher-than-average probabilities of being in debt.

In sum, the model has several important implications for our empirical work. Under reasonable parameters, a substantial fraction of optimizing households may have low or

¹³Prior to their spell of unemployment, the assets of these households would average 6-1/2 percent of permanent income (the population mean). Since the social insurance system provides income of 0.2 and these households borrow 0.33 on average, their consumption will average $0.066 + 0.2 + 0.33 = 0.6$ of permanent income during the year of unemployment. Thus, the social insurance system together with the availability of credit (even at a 20 percent interest rate) provide substantial insurance against major interruptions to consumption while the household is unemployed.

negative net worth. For households with similar recent employment experiences, those facing a greater probability of job loss should hold more wealth. However, a household with high unemployment risk may hold less assets than a household with low risk if the high-risk household has recently experienced a spell of unemployment. Finally, precautionary responses may vary depending on the level of the wealth-to-permanent income ratio.

Of course, the model is meant only to highlight the relevant considerations for designing a test of the relationship between wealth and unemployment risk. Besides ignoring several important motives for saving, the model does not allow for differences in preference parameters across households or for the possibility that some households' behavior may be explained better by non-standard models. These issues are discussed further in section IV.

III. Data and Econometric Methodology

Our empirical goal is to construct a model of the relationship between W/Y^p and unemployment risk that takes into account the salient features of the cross-section wealth distributions generated by the simulations in section II. We use a two-step procedure. The first step is to construct estimates of unemployment risk and permanent income for each household, and the second step is to estimate the relationship between W/Y^p and unemployment risk, controlling for other explanatory variables that reflect the influence of tastes and the income process on wealth accumulation.

III.1. Data Sources

Like Starr-McCluer, Engen and Gruber, and Carroll and Samwick, we focus on the cross-sectional relationship between the stock of household wealth and uncertainty. We use wealth data from the Federal Reserve's *Survey of Consumer Finances* because it represents the best available information on households' balance sheets.¹⁴ We use three cross-sectional waves of the SCF: 1983, 1989, and 1992. Each wave contains very detailed and accurate information on the assets and liabilities of a household as well as data on their current income, employment status, and demographic characteristics. Each SCF observation records

¹⁴Curtin, Juster, and Morgan (1989) compare the SCF wealth data to the wealth data from the PSID and the SIPP and conclude that "the unique design characteristics of the SCF give it the highest overall potential for wealth analysis of the three data sets examined." (p.58)

data on assets, liabilities, and income for the household as a whole. The 1983 and 1992 surveys contain roughly 4000 observations, while the 1989 survey covers a bit more than 3000 households.

The model in section II was limited to one type of financial asset and one type of unsecured liability. The real world offers a large range of assets and liabilities. While some assets clearly are more costly than others to use as buffers against adverse shocks to income, home equity lines of credit and the ability to borrow against defined contribution pension plans allow even relatively illiquid assets to provide cash within a matter of weeks. Accordingly, our baseline empirical model will focus on households' total net worth, which we define as the sum of financial assets, real estate, noncorporate business equity, and vehicles, less mortgage and consumer debt. Real estate, businesses, vehicles, mutual funds, and equities are valued at respondents' assessments of their current market value; other assets are valued at face value. This wealth measure includes the value of defined contribution pensions plans, but excludes the value of defined benefit pension plans (which is difficult to calculate accurately) and durable goods other than vehicles (which are not recorded in the SCF).

Unfortunately, because the SCF at most covers only about 4000 households in a given year, it contains too few observations on unemployment to accurately estimate the probability of job loss for the regional breakdown that we work with below. Consequently, we use data from the much larger *Current Population Survey* (CPS) to determine the relationship between household characteristics and unemployment risk. The CPS has a quasi-panel structure: A randomly selected household is interviewed for four consecutive months, then rotates out of the sample for eight months, and then rotates back for four more consecutive months. For our analysis, we use data from each household's fourth and eighth interviews.¹⁵ Thus, for

¹⁵The fourth and eighth interviews are commonly referred to as the "outgoing rotation groups;" data from these have been compiled in a format that is much easier to use than the full CPS. To create our data set, we used programs provided by Welch (1993) to match records from CPS outgoing rotation group extracts for the same months in consecutive years. We were able to match about 75 percent of the households across years (the rest were lost to attrition) and about 95 percent of the individuals within the matched households.

each household, we have two records of employment status and demographic characteristics taken a year apart.

III.2. Sample Selection

Parameter consistency of our two-sample estimator requires that the samples be randomly drawn from the same population. Both the CPS and the SCF area-probability samples are random draws from the noninstitutional U.S. population. However, the SCF also includes a special nonrandom "list" sample designed to oversample wealthy households. Since these list samples are not random and covers a very different population than the CPS and SCF area-probability samples, we excluded them from our analysis. The area probability sample comprised about 90 percent of the total SCF sample in 1983, 72 percent in 1989, and 63 percent in 1992.

The simulations in section II indicate that, on average, households who have recently experienced a spell of unemployment or other major earnings disruption will have low net worth. Since consumers with high current unemployment risk also are more likely to have recently experienced a bad draw to income, even in the presence of a strong precautionary saving motive we could observe a negative correlation between wealth and the *ex ante* probability of unemployment due to the influence of earlier shocks to earnings. Unfortunately, the SCF contains no explanatory variables that can be used to directly control for such major disruptions to income. The SCF, however, does record time at the current job, and we thus decided to limit our analysis to households with similar recent employment experiences by including only households whose head currently is employed and has been at the same job for at least three years.¹⁶

Next, we excluded households whose heads are younger than 20 or older than 65, since employment risk likely is irrelevant for people who have yet to enter the workforce permanently or who are beyond a normal retirement age. Finally, to remove some obviously

¹⁶This restriction introduces complications of its own: Our sample inherently will have fewer high risk households than the population at large. This is particularly true for the 1983 and 1992 surveys, which were done near cyclical downturns. With fewer high risk households, we may have more difficulty identifying a precautionary effect. However, it is unclear what bias may arise from this restriction.

extreme outliers in our data, we also dropped households with the very highest or very lowest (0.1 percentiles) readings for wealth or income.¹⁷ These restrictions leave us with 1,689 households from the 1983 SCF, 1,025 households from the 1989 survey, and 1,032 households from 1992.

For the CPS, we used the same age restrictions as for the SCF. Although the SCF interviewing took place over periods of less than a year, in order to avoid potential problems due to seasonality in unemployment rates, we predict unemployment risk by pooling information from the CPS households whose fourth interview fell sometime during the twelve months preceding the end of the sampling period for each SCF wave.¹⁸ We are left with CPS samples of 59,252 in 1983, 60,026 in 1989, and 63,351 in 1992.¹⁹

III.3. Wealth Transformation

Some summary statistics for net worth, assets, and liabilities are shown in table 3. The distribution of wealth is highly skewed, with the net worth of the median households being only about half the size of the average in each year. As discussed earlier, many households have very low or negative net worth: between 8 to 10 percent have net worth less than one month's income and 4 to 6 percent have negative net worth.

¹⁷Specifically, we excluded households in the lowest and highest 0.1 percentiles (calculated on a weighted basis) of the four measures of wealth we consider in section IV--total net worth, net worth excluding real estate, net worth excluding primary residence, and net financial assets--and in the highest and lowest 0.1 percentiles of income and the ratio of debt service to income.

¹⁸Interviewing for the 1983, 1989, and 1992 SCFs took place between February 1983 and August 1983, August 1989 and March 1990, and June 1992 and November 1992, respectively. Thus, our three different CPS samples are for respondents whose fourth interview fell in the twelve months preceding August 1983, March 1990, and November 1992.

¹⁹Although the SCF observations are at the household level and we generally used the demographic information only for the head of each household, we did not limit the CPS sample to only heads of households because it would have reduced our sample size substantially. We instead included a dummy variable in the CPS regressions for whether the record is for the head of household (SCF definition), as well as terms that interact the head dummy with age and race. We tried interacting other variables with the head dummy as well, but these did not enter the first-stage regression significantly.

Clearly, some transformation of the wealth data is necessary in order to avoid undue influence from extreme observations. The distribution is sufficiently dispersed that the potential solution provided by the model in section II--that is, using the ratio of net worth to some proxy for permanent income as the dependent variable--does not work. For example, the top panel of figure 3 plots a histogram of the residuals from a linear regression of a W/Y^p proxy on the explanatory variables used in the model estimated in table 7 for 1983 (base sample) along with a normal density with the same mean and variance as the residuals. The substantial mass in the tails and the skewness of the distribution indicates that the residuals are very far from normal.²⁰

As noted earlier, most previous studies of household wealth have downweighted large observations by taking logarithms, imposing some additional transformation or restriction on the sample to deal with zero or negative readings of net worth. Some examples include:

Paper	Transformation	Restrictions
Diamond and Hausman (1984)	$\ln(W)$	$W > \$4000$
King and Dicks-Mireaux (1982)	$\ln(W/Y^p)$	$W > \$2500$
Starr-McCluer (1996)	$\ln(W)$ if $W > 0$ 0 if $W \leq 0$	none
Carroll and Samwick (1997, 1998)	$\ln(W - \min(W, 0) + 1)$	none

Transformations that truncate or censor the lower part of the wealth distribution may eliminate many households for whom employment risk is high or for whom a significant share of net worth potentially is accounted for by precautionary behavior. Restricting the dependent variable also adds the complication that the estimation technique should in principle account for sample selection. The log transformation also is problematic because it assumes a constant elasticity of wealth with respect to changes in explanatory variables. Recall that our model simulations imply that at the low end of the distribution, small increases in dollar terms in precautionary wealth can represent extremely large increases in

²⁰For exposition, we eliminated the eight largest positive residuals before plotting this histogram, so the actual distribution is even more skewed than the one presented in figure 3.

percentage terms. Accordingly, average precautionary effects estimated under a log transformation could give undue weight to the responses at the lower end of the wealth distribution.

As an alternative, we use the inverse hyperbolic sine function to transform our net worth variable. This transformation was proposed by Johnson (1949) and suggested for use with wealth data by Burbidge, Magee, and Robb. The inverse hyperbolic sine of z is:

$$g[z, \theta] = \ln[\theta z + (\theta^2 z^2 + 1)^{1/2}] / \theta \quad (9)$$

where θ is an estimated damping parameter. Like the log transformation, $g[z, \theta]$ downweights large values of z . However, it has several advantages over the use of logs: The transformation allows zero and negative values of z , estimates the degree to which large values are downweighted, and does not impose constant elasticities. The effect of a change in any independent variable x on z is given by $\frac{\partial z}{\partial x} = \frac{\partial g[z, \theta]}{\partial x} \cdot \frac{\partial z}{\partial g[z, \theta]}$, where the first term on the right hand side equals the regression coefficient on x and the second term equals $(\theta^2 z^2 + 1)^{1/2}$. Thus, $\partial z / z = (\theta^2 + 1/z^2)^{1/2}$, so that elasticities are decreasing functions of z .²¹

The middle panel of figure 3 shows $g[z, \theta]$ with z equal to our W/Y^p proxy and with $\theta = 3.87$ (the value we estimate for the model with the 1983 total sample in table 7). The bottom panel of figure 3 shows the residuals from a regression of $g[W/Y^p, \theta]$ on the same explanatory variables used in the linear regression cited above. This distribution is much closer to normal than that for the linear model, although the tails still contain a little more mass than a normal distribution with the same variance.

III.4. The Likelihood Function

Indexing households by j , the model we wish to estimate is:

$$g[W_j/Y_j^p, \theta] = \beta_0 + \beta_u \Pr(u_j) + \beta_y \ln Y_j^p + C_j \beta_c + \epsilon_j, \quad (10)$$

²¹Note that for large z , $\partial z / z$ approaches a constant elasticity of θ . As discussed in Burbidge, Magee, and Robb, $\lim_{\theta \rightarrow 0} g[z, \theta] = z$, while $\lim_{|\theta z| \rightarrow \infty} g[z, \theta] = \text{sign}[\theta z] \ln(2|\theta z|) / \theta$, so that the transformation encompasses a variety of functional forms.

where $g[\cdot]$ is the inverse hyperbolic sine function with parameter θ , W_j is net worth, Y_j^p is permanent income, $\Pr(u_j)$ is the probability of the household head becoming unemployed, and C_j is a row vector of control variables. We include $\ln Y_j^p$ as an explanatory variable because a growing body of evidence suggests that saving behavior may vary across levels of permanent income (see Dynan, Skinner, and Zeldes, 1997, Carroll and Samwick (1997, 1998), and Lillard and Karoly, 1997).

Assuming $u_j \sim N(0, \sigma^2)$, the log likelihood function of the W_j/Y_j^p is

$$L[\beta, \theta, \sigma^2] = K - (n_s/2) \ln \sigma^2 - (1/2\sigma^2) \sum_{j=1}^{n_s} \epsilon_j^2 - 1/2 \sum_{j=1}^{n_s} \ln[1 + \theta^2(W_j/Y_j^p)^2] \quad (11)$$

where n_s is the size of the SCF sample and K is the usual constant. The last term derives from the Jacobian of $g[W/Y^p, \theta]$. Nonlinear maximization of (11) produces estimates of the β 's, θ , and σ^2 .

III.5. First-stage Regressions for Unemployment Risk

For currently employed individual j , we assume there exists a latent variable $u_j^* = Z_j^u \alpha_u + v_j$ such that $u_j^* > 0$ if the person will be unemployed one year hence and $u_j^* \leq 0$ if the person will be employed. v_j is a logistically distributed idiosyncratic shock that is uncorrelated with Z_j^u , a row vector of observable characteristics for individual j . Thus, $\Pr(u_j | e_j)$, the probability of a currently employed person becoming unemployed, is

$$\Pr(u_j | e_j) = \frac{\exp(Z_j^u \alpha_u)}{1 + \exp(Z_j^u \alpha_u)} \quad (12)$$

We estimate this probability using data from the CPS. The dependent variable is an indicator that takes on a value of 1 if individual j is employed in month t (observed in the fourth CPS interview) and unemployed in month $t+12$ (observed in the eighth CPS interview), and it takes on a value of 0 if individual j is employed in both periods. Thus, we are exploiting the quasi-panel structure of the CPS to estimate the probability of *becoming* unemployed rather

than simply the probability of *being* unemployed. Of course, our procedure is imperfect in that it does not capture the length of the unemployment spell and it misses spells completed before $t+12$. For notational convenience, let $\Pr(u_j)$ equal $\Pr(u_j|e_j)$.

To proxy for the probability of an employed SCF household head becoming unemployed, we calculate $\Pr(\hat{u}_j) = \exp(Z_j^{uS} \hat{\alpha}_u) / [1 - \exp(Z_j^{uS} \hat{\alpha}_u)]$, where $\hat{\alpha}_u$ is the CPS-based estimate of α_u and Z_j^{uS} are the values of Z_j^u of the SCF household heads. Because the SCF area-probability sample and the CPS sample are both randomly drawn from the same underlying population, this procedure produces consistent estimates of the probability that a household from the SCF area-probability sample will become unemployed (see Angrist and Krueger, 1992).

For Z_j^u , we are restricted to variables that are common to both the CPS and SCF. Our first-stage equation contains regressors for occupation, industry, region, education, age, age-squared, age interacted with occupation, age interacted with education, marital status, race, gender, a dummy for head of household, the head dummy interacted with age, and the head dummy interacted with gender.²²

III.6. First-stage Regressions for Permanent Income

Because every household in the SCF reports income and a wide range of other variables, our first-stage estimates for permanent income can be done entirely within the SCFs. The log of permanent income is assumed to be a function of observable characteristics, Z_j^{yS} :

²²The SCF data sets contain six major occupational headings: managerial and professional; technical, sales and support; services; precision, craft, and repair; operators and laborers; and farming, forestry and fisheries. They contain seven industry classifications: agricultural; mining and construction; manufacturing; wholesale and retail trade; FIRE and business services; other private service producing; and public administration. The SCF does not provide more detailed information about occupation and industry for reasons of confidentiality. We use the nine Census subdivisions as regional identifiers: This information is not available for the 1989 and 1992 area-probability samples in the SCF public-use data set, but the SCF staff helped us by running our programs on their internal data set which contains the regional identifiers.

$$\ln Y_j^p = Z_j^{yS} \alpha_y . \quad (13)$$

We use the fitted value from an OLS regression of observed $\ln Y_j$ on Z_j^{yS} as our estimate of the log of permanent income, $\ln \hat{Y}_j^p$. We include in Z_j^{yS} all of the Z_j^{uS} along with a dummy for home ownership, number of children in the household, number of earners, a dummy for retirement, the log of retirement income, a dummy for whether the head has a defined benefit pension, and dummies for whether the household has ever been turned down for credit or has had problems servicing loans. We designate the variables that are in Z_j^{yS} but not in Z_j^{uS} as $Z_j^{\bar{u}S}$.

Effectively, each household's proxy for permanent income equals average income for all households with similar characteristics. This approach has been widely used, beginning with Friedman (1957). Note, though, that such measures of permanent income do not capture any unobservable individual-specific components of permanent income.

III.7. Control Variables

Finally, we need to specify the control variables, C_j , which are meant to capture factors that may affect wealth through some channel other than unemployment risk or permanent income. As we noted in the introduction, many studies identify the effect of the instrumented uncertainty proxy on wealth by excluding some variable such as occupation or industry as independent regressors in specifications similar to equation (10). However, because the household's choice of such variables can affect the expected life-cycle profile of income or reflect idiosyncracies in discount rates and risk aversion, they may be correlated with wealth through some avenue other than their influence on $\Pr(\hat{u}_j)$ or $\ln \hat{Y}_j^p$. Thus, we include them as control variables in equation (10).

Indeed, because we cannot rule out *a priori* that most variables in Z_j^{yS} might have some independent influence on wealth, we have included all but one of them in C_j . The exception is region, which we believe is less likely to be correlated with the preference parameters that determine the saving behavior of individual households. Furthermore, Blanchard and Katz (1992) show there are no permanent differences in unemployment rates

across regions. This makes it unlikely that a significant number of households have chosen *ex ante* to live in a particular region because of perceived *permanent* disparities in job-loss risk. Thus, regional variation in unemployment is likely quite “exogenous” to an individual household and probably provides a cleaner signal of the effect of $\text{Pr}(u_j)$ on wealth than a measure based solely on the other variables in Z_j^{uC} .

III.8. Other Estimation Issues

As a result of the assumptions on Z_j^{uC} , C_j , and Z_j^{yS} , β_y is identified from region, while β_u is identified from both region and the nonlinear functional form of equation (12). Because, $Z_j^{yS} = Z_j^{uS} \cup Z_j^{\bar{u}S}$, a sufficient condition for parameter consistency is for all of the variables in Z_j^{yS} to be uncorrelated with v_j . This presumes that the $Z_j^{\bar{u}S}$ have no predictive power for unemployment risk independent of the Z_j^{uS} . If, in contrast, these variables were correlated with $\text{Pr}(u_j)$, the estimate of β_u still would be consistent because the $Z_j^{\bar{u}S}$ are included in C_j ; the β_C for these variables, however, would be inconsistent as they would now capture both direct effects on W_j/Y_j^p and the effects of their predictive power for $\text{Pr}(u_j)$.

Instead of maximizing (11), we maximize the constructed likelihood:

$$L[\beta, \theta, \sigma^2] = K - (n_s/2) \ln \sigma^2 - (1/2\sigma^2) \sum_{j=1}^{n_s} e_j^2 - 1/2 \sum_{j=1}^{n_s} \ln[1 + \theta^2 (W_j/\hat{Y}_j^p)^2] \quad (14)$$

$$e_j = g[W_j/\hat{Y}_j^p, \theta] - \beta_0 - \beta_u \text{Pr}(\hat{u}_j) - \beta_y \ln \hat{Y}_j^p - C_j \beta_c$$

Obviously, the use of first-stage equations to estimate variables in equation (14) introduces some complications. The appendix discusses parameter consistency and adjustments we make so that standard errors and other statistics will (asymptotically) account for the first-stage estimation of $\text{Pr}(\hat{u}_j)$ and \hat{Y}_j^p .

Finally, note that the relative unemployment risk associated with the different regions changes over time. For example, when looking at simple top-to-bottom orderings, the average changes between adjacent survey years in the ranking of a region's unemployment rate is

about 1-3/4 positions.²³ Accordingly, the use of three years of data should reduce the odds that our results are just capturing regional fixed effects on wealth.

IV. Empirical Results

IV.1. First-Stage Results

Table 4 presents results from the first-stage CPS logit equations for the probability of becoming unemployed in 1983, 1989, and 1992. The rows present p-values for F-tests of the joint significance of the corresponding independent variables. The industry, region, education, and race dummies are all statistically significant (as groups) at the 1 percent level or better in the first-stage CPS regressions. The occupation dummies and female head variable provide power in estimating $\Pr(\hat{u}_j)$ in 1983 and 1992, but not in 1989. Finally, the age variable does not explain $\Pr(\hat{u}_j)$ in any year, but age interacted with occupation is highly significant in 1983 and marginally significant in the other years, while age interacted with education is significant at 7 percent level or better in all three years. The lower part of the table provides summary statistics for the predicted probability of becoming unemployed. The mean probabilities are 2.7 percent, 1.9 percent, and 2.2 percent for 1983, 1989, and 1992, respectively. More importantly, in order to obtain precise estimates of β_w , we need the variation in predicted job loss to be large. Fortunately, the models produce considerable variation in $\Pr(\hat{u}_j)$, with the standard deviations running between 70 and 80 percent of the means.²⁴

Table 5 presents results from the first-stage income regressions. The adjusted R-squared statistics indicate that these variables explain close to half of the variation in the log of reported income in 1983, and between 30 and 40 percent in 1989 and 1992. Most of the explanatory variables are statistically significant; a notable exception is education in 1983 and 1989. A simple regression of income on education dummies alone produces significant coefficients and explains a nontrivial percentage of the variation in income (for example, 15

²³If regional effects were fixed, we would see no change in relative rankings. If positions were completely random, we would observe an average change of 4-1/2 positions.

²⁴Recall that the mean probabilities are not some unemployment rate *per se*, but rather the probability that a household head currently employed will be unemployed one year hence. For reference, the time series variance in the quarterly aggregate unemployment rate between 1983 and 1992 was 1.3 percentage points.

percent in 1983). Accordingly, the insignificance in the overall equation likely results from collinearity issues and not the irrelevance of education to permanent income.

IV.2. Second-Stage Results: The Basic Model

The results for our second-stage equation are found in table 6. For exposition, we show estimated coefficients and (asymptotic) t-statistics only for selected variables. The left hand columns for each year present results from the total area samples, while the right hand columns show estimates excluding households whose head is self employed from the samples.²⁵ We separate these households because their balance sheets can be heavily influenced by business holdings and because self-employed households may have different attitudes toward risk.

The results provide little evidence for the hypothesis that households accumulate more wealth in response to an increased probability of becoming unemployed. In the total list sample, the point estimate of the coefficient on $\text{Pr}(\hat{u}_j)$ is positive in 1983, but it is close to zero in 1992 and is negative in 1989. In no year is the coefficient statistically different from zero. To translate this result into an economically and statistically meaningful estimate of the effect of a change in $\text{Pr}(u_j)$ on W_j/Y^p , we must take account of the inverse hyperbolic sine transformation. Following the results in section III.3, we have

$$\frac{\partial W_j/\hat{Y}_j^p}{\partial \text{Pr}(\hat{u}_j)} = \left((\hat{\theta} W_j/\hat{Y}_j^p)^2 + 1 \right)^{\frac{1}{2}} b_u \quad (15)$$

where b_u is the estimated coefficients on $\text{Pr}(\hat{u}_j)$. The table presents $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ and asymptotic t-statistics that correspond to a one percentage point increase in $\text{Pr}(\hat{u}_j)$ at the median W_j/Y_j^p . (Recall that $\text{Pr}(\hat{u}_j)$ averages between 2 and 2-3/4 percent and the standard

²⁵The runs that exclude self-employed households also exclude them in the first-stage estimates of unemployment risk and permanent income. Note that the question determining self employment in the CPSs and 1983 SCF asks if the respondent is an *employee* of a private business, government, or is self employed. The 1989 and 1992 SCFs ask if the respondent *works* for someone else or is self employed. Conceivably, some respondents, notably private contractors, could answer these questions differently.

deviations are in the 1-1/2 to 2 percentage point range; thus, a one percentage point increase in $\Pr(\hat{u}_j)$ represents a fairly large change in our metric of employment risk.) In the 1982 sample including the self employed, a one point increase in $\Pr(u_j)$ increases \bar{W}_j/Y^P by 0.04, or about 1/2 month of income, but with a t-statistic of 0.5, this effect is not statistically significant. The effects for 1989 and 1992 also are small and not statistically different from zero. The samples excluding the self employed produce similar small and statistically insignificant estimates of b_u and $\partial(\bar{W}_j/Y^P)/\partial\Pr(u_j)$.

Our simple theoretical model in section II does not imply any correlation between permanent income and the ratio of wealth to permanent income. However, the coefficient on $\ln \hat{Y}_j^P$, β_y , is positive for all years and samples and is of at least marginal statistical significance (p-values of 12-1/2 percent or better) in all but the 1983 sample including the self employed. These estimates suggest either some misspecification of the model in section II or a mismeasurement of permanent income.

The signs of the coefficients on the other independent variables generally make sense and are reasonably stable across SCF years and between the samples that include and exclude the self employed. The ratio of net worth to income is decreasing in the number of earners per household, consistent with the notion that having multiple earners in the family lowers precautionary reserves because both earners are unlikely to become unemployed at the same time. As documented by previous studies, homeowners have higher net worth relative to income. Having a defined benefit pension plan has a negative effect on the net worth ratio, as would be predicted by a standard forward-looking model.

Finally, identification of β_y and, as a practical matter, of β_u requires that the error term be orthogonal from the excluded instruments--the regional dummies. In econometric terms, the critical assumption is that conditioned on the C_j , the regional dummies are correlated with the dependent variable only via their correlation with $\Pr(\hat{u}_j)$ and $\ln \hat{Y}_j^P$. This assumption is formally tested via the overidentifying restrictions (OID) tests presented in table 6. These suggest that region is indeed a valid excluded instrument: none of the p-values is close to statistical significance. Of course, these findings are subject to the usual caveats concerning the low power of such tests.

IV.3. Second-Stage Results: An Extended Model

A possible explanation for the weak results in table 6 is that the model is too stylized to fully characterize actual household behavior. A logical extension of the model--and one suggested by the positive and marginally significant coefficients on $\ln \hat{Y}_j^p$ --would be to allow the strength of the precautionary response to vary with income. Notably, our model's simple specification of V_{min} may not capture all of the salient features of means-tested social insurance--programs which can, as shown by Hubbard, Skinner, and Zeldes (1995), lead households at the lower end of the income distribution to eschew saving because accumulated wealth might cause them to forfeit social insurance benefits. Alternatively, some households may use a "rule of thumb" rather than a forward-looking, optimizing framework to choose consumption and saving (as in Campbell and Mankiw, 1989). Such households may not react at all to changes in income risk. If low permanent income households are more likely to be rule-of-thumb consumers, this might generate an association between the level of permanent income and the precautionary response of wealth to risk.

Table 7 presents results for a second-stage equation that allows the precautionary response to vary with income by including the term $\Pr(\hat{u}_j) * \ln \hat{Y}_j^p$ as an independent variable. The estimates for the control variables in this specification are similar to those in the base specification, and so for exposition, we have not included them on the table.

Allowing for the interaction between unemployment risk and income produces a much different picture for the role of uncertainty: the estimated coefficients on $\Pr(\hat{u}_j)$ are uniformly negative and those on $\Pr(\hat{u}_j) * \ln \hat{Y}_j^p$ uniformly positive, with both significantly different from zero at the 5 percent level in the total samples for all three years and statistically significant at at least the 10 percent level in the sample excluding the self employed. Furthermore, the coefficients and t-statistics on $\ln \hat{Y}_j^p$ are much smaller than those in the base model.

To calculate $\partial(W_j/Y_j^p)/\partial\Pr(u_j)$, we now must account for both the role of $\ln \hat{Y}_j^p$ in the extended model and the inverse hyperbolic sine transformation; the appropriate statistic is

$$\frac{\partial W_j/Y_j^p}{\partial \Pr(\hat{u}_j)} = \left((\hat{\theta} W_j/Y_j^p)^2 + 1 \right)^{\frac{1}{2}} \left(b_u + b_{uy} \ln \hat{Y}_j^p \right) \quad (16)$$

where b_{uy} is the estimated coefficient $\Pr(\hat{u}_j) \cdot \ln \hat{Y}_j^p$. The bottom portion of table 7 presents estimates and (asymptotic) t-statistics for $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ that correspond to a 1 percentage point increase in $\Pr(\hat{u}_j)$ for the percentile of \hat{Y}_j^p listed in the first column.²⁶

For the samples including the self employed, at low levels of \hat{Y}_j^p , we see only small and statistically insignificant effects of unemployment risk on wealth in all three years. For example, in the 1983 sample including the self employed, at the 10th income percentile, a 1 percentage point increase in the probability of becoming unemployed is associated with a 0.03 increase in the ratio of wealth to income--or roughly a few days' income--and has a t-statistic of less than one. Thus, the results are consistent with the hypothesis that low-income households have little or no precautionary response.

However, the estimated precautionary responses become economically significant as fitted permanent income rises. By the 50th percentile of \hat{Y}_j^p , the 1983 total sample estimates imply that households respond to a 1 percentage point increase in the probability of becoming unemployed by increasing wealth-to-income ratio by 0.29, or about 3-1/2 (=0.29 x 12) month's income. These effects are large enough to be economically important, but not so large as to be intuitively implausible. Furthermore, for the 50th percentile of \hat{Y}_j^p and above, the estimates of $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ are statistically different from zero at the 5 percent level.

The estimates of $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ for the 1989 and 1992 SCF's including the self employed are of very similar magnitude as the 1983 results, with the effects differing at most by few percentage points of annual income. However, the point estimates are much less precise than those in 1983. In part, this might reflect the fact that the samples in those years are about 40 percent smaller than the 1983 sample. Once again, none of the OID tests reject the region exclusion restrictions.

The $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ estimates are somewhat smaller for the samples that exclude

²⁶Strictly speaking, we divided the data into 50 equally sized bins based on \hat{Y}_j^p and then calculated equation (16) using the median values of \hat{Y}_j^p and W_j/\hat{Y}_j^p for the bin containing the listed \hat{Y}_j^p percentile (so, for example, we actually are looking at the 49th percentile for \hat{Y}_j^p instead of the 50th percentile). We did this to avoid the event that the particular household at the chosen percentile of \hat{Y}_j^p had an outlier value for W_j . We also multiplied equation (16) by 0.01 so that the change in W_j/\hat{Y}_j^p corresponds to a 1 percentage point increase in the probability of becoming unemployed.

self-employed households. This result is consistent with the hypothesis that self-employed households may be more vulnerable to income losses when business conditions go bad and that the fact that they are less likely to be covered by unemployment insurance. The precautionary effects are substantially less precisely estimated than in the total sample; again, this is in part a function of reduced sample size.

IV.4. Second-Stage Results: Pooled SCF Samples

Recall that our sample restrictions left us with 1,689 households from the 1983 SCF, 1,025 households from the 1989 survey, and 1,032 households from 1992. It may be the case that these samples for the individual SCF years are too small to produce precise estimates of the relationship between wealth and unemployment risk. Yet the remarkable similarity in the sizes of the precautionary effects and the other coefficients across years suggests that we can pool the data from the three different SCF years for most of the second-stage wealth regressions in order to gain more precise estimates.

Table 8 presents results based on the pooled data for the base model (left-hand columns) and the model with the $\Pr(\hat{u}_j) \ln \hat{Y}_j^P$ term (right-hand columns). Note that although the coefficients on income, job-loss risk, and the C_j are fixed across SCF years in the second-stage regression, we continue to estimate the first-stage regressions for $\Pr(\hat{u}_j)$ and $\ln \hat{Y}_j^P$ separately for the three years to allow the relationships between unemployment risk, permanent income, and household characteristics to vary with the different macroeconomic conditions that prevailed in 1983, 1989, and 1992. We also add year dummies for 1989 and 1992 to allow the average level of wealth to vary over time with trend growth and with business cycle conditions.

The base model now shows a small precautionary effect--for the sample that includes the self-employed, the wealth-to-income ratio is increased by 0.06 (0.7 months of income) at the median income--but it is statistically significant only at the 20 percent level. In the model with the interaction, we still observe only small and statistically insignificant effects of unemployment risk on wealth at low levels of \hat{Y}_j^P , but statistically and economically significant precautionary wealth shows up for households at the 30th percentile of permanent income and higher.²⁷ At the median income, a one percentage point increase in the

probability of becoming unemployed is associated with an increase in precautionary balances equal to 0.17 times of annual income (2 months), and, with a t-statistic of over 3, this effect is highly statistically significant. The effects for the sample excluding the self-employed are a bit smaller than in the overall sample, but with the much larger sample size, the estimated effects are highly statistically significant for the upper half of the income distribution.

Note that these estimates of $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ appear much too large to be caused solely by the loss of expected permanent income associated with an increase in job-loss risk--a factor that would cause $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j) > 0$ even in a certainty equivalence world with no precautionary saving. More specifically, the certainty equivalence solution for consumption calculated from the model in section II (with $R = \beta$) is:

$$C_t = \frac{r}{1+r} X_t + \frac{r}{r-g} Y_t \quad (17)$$

Let $g = 0.02$ and $r = 0.05$ (so $r/(r - g) = 1.67$). Now, suppose that there is a one percent increase in the probability of losing a job for one-half year and that the household expects to subsequently find a new job that permanently pays 10 percent less. (Note that this unemployment spell and new-job income reduction are more severe than the averages found in most empirical work--see, e.g., Carrington (1993).) Equation (17) indicates that in a certainty equivalence world, consumption would fall--and wealth would rise--by

²⁷Note that although the point estimates for b_u and b_{uy} in the pooled sample fall within the range of estimates for three years, the point estimate of $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ does not: While equation (16) is monotone increasing in both b_u and b_{uy} , across years, larger values of b_u also are associated with more negative values of b_{uy} . Formal tests of the pooling restrictions are ambiguous. In tests that do not account for the fact that $\text{Pr}(u_j)$ and $\ln Y_j^p$ are estimated, an F-test easily fails to reject the hypothesis that the β_u , β_y , β_{uy} , and β_c are constant across years but a likelihood ratio test easily rejects the pooling restrictions (this test includes the restrictions on σ_e^2 and θ). However, both tests degenerate when they are adjusted for the first-stage estimates--they produce the (asymptotically) nonsensical result that the sum over the three SCF year regressions of adjusted estimates of $n_s \sigma_e^2$ exceeds the adjusted estimate from the pooled regression. This likely reflects asymptotic approximation error in the factors we use to blow the $\Sigma \hat{e}_j^2/n_s$ up to consistent estimates of $\sigma_e^2 = \text{plim} \Sigma e_j^2/n_s$. (In the income-interaction model, the adjustments boost σ_e by about 25 percent in 1983 and the pooled sample and about 50 and 70 percent in 1989 and 1992, respectively. See the appendix for details.)

$0.01*(0.5+0.1)*1.67 = 0.01$ of income. We estimate empirically that a 0.01 increase in job-loss risk raises the median consumer's wealth by 0.17 of income; thus, the increase in job-loss probability would have to be maintained for 17 years for our estimates to be in line with even a pessimistic scenario from a certainty-equivalence permanent income model.

We cannot precisely relate our results to other estimates of precautionary saving because of the difficulty of putting the studies' different measures of uncertainty on a comparable basis. That said, our findings appear to tell roughly the same story as Carroll and Samwick (1997) and Engen and Gruber, but probably indicate larger precautionary effects than those found in Lusardi (1997). Our median household increases W_j/\hat{Y}_j^p by 17 percent in response to a 1 percentage point--or roughly a 1/2 standard deviation--increase in $\Pr(\hat{u}_j)$. Carroll and Samwick (1997) estimate a 4 percent increase in the ratio of wealth to permanent income in response to a 1 percentage point increase in the variance of transitory income; but given that they estimate this variance to be between 2 and 10 percent, they are probably considering a smaller increase in uncertainty than we are. Engen and Gruber estimate that a 10 percent increase in the income replaced by unemployment insurance--about a 1/2 standard deviation change in the UI replacement rate (see Gruber 1994)--lowers the ratio of net financial assets to income by 16 percent. In contrast, using the Italian *Survey of Household Income and Wealth*, Lusardi finds that--roughly speaking--a 1 standard deviation increase in the variance of respondents' subjective expectations of nominal income changes is associated with just a 1-1/2 to 3 percent increase in W_j/\hat{Y}_j^p .

IV.5. Identification Issues

We noted earlier that region may be a better instrument to exclude from C_j in order to identify uncertainty than variables some other papers have used--such as occupation, education, or industry--which *a priori* are more likely to depend on unobserved taste parameters, such as risk aversion, that are also correlated with saving. And as in the year-by-year estimates, the p-values for the OID tests in table 8 are all quite large, suggesting that region is a valid excluded instrument. To consider the other instruments that we could have used--and that most previous authors have used--for identification, table 9 presents versions of the pooled SCF interaction model (for the sample including the self employed) reestimated

with occupation, education, and industry excluded one-by-one from the controls.²⁸ (For reference, we repeat the region-excluded regression in column 1.) When excluding occupation or education, $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ is reduced substantially, while the estimates when excluding industry are essentially the same as in the regions-exclusion case, suggesting that the effect of job-loss risk on wealth is the same whether that risk comes from living in a region that is temporarily undergoing a recession or from working in an industry where job-loss risk is currently high.²⁹ The OID tests reject the exclusion restrictions on occupation and industry at (close to) the five percent level, but do not reject education.

OID tests have been criticized because in some circumstances they have low power. The rejections of occupation and industry indicates that our specification is not one where OID tests are powerless, and the fact that region passes the test suggests that it is a valid instrument. However, the failure to reject education, which *a priori* seems likely to be endogenous with regard to unobserved traits determining saving, suggests that OID results may not have as much power as we might like. Alternatively, the problem may be that education is highly collinear with the included control variables. (Recall the earlier first-stage regressions for $\ln \hat{Y}_j^p$ suggested this.)

Because we estimate the first-stage equations independently for each SCF year, our instrument set in the pooled model effectively is comprised of the Z_j^{ys} interacted with SCF year dummies. In contrast, the coefficients on the C_j are restricted to be the same across SCF years. Thus, we can test if the excluded instruments should enter C_j without causing a singularity between C_j and the instrument set. The bottom row of table 9 presents p-values for Lagrange multiplier tests of the exclusion of the appropriate instrument from C_j . These

²⁸Our test for each of these variables parallels our treatment of region in Table 8: The variable in question is used as the *only* excluded instrument (in particular, when occupation, education, or industry is being tested, we include region in the set of controls).

²⁹In general, if the excluded instrument is correlated with risk aversion, then the estimates of precautionary effects probably are biased downward (although, of course, the bias cannot be signed directly in a multivariate model). For example, consider a case where occupation is the excluded instrument and people with high risk aversion both choose an occupation with low unemployment risk and hold higher precautionary balances; this behavior makes a negative contributes to the correlation between wealth and $\text{Pr}(u)$.

are soundly rejected in all specifications, including with our preferred regional exclusion.

The last column of table 9 presents estimates of the interaction model after adding region to the C_j . In this case, abstracting from functional form, *all* of the identification of the effect of unemployment risk on wealth in this specification comes from the *changes across years* in the regional job-loss risk variation. The estimates of $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ are smaller than those in our model excluding fixed-over-time region-specific intercepts. Still, by the 50th percentile of \hat{Y}_j^p , a one percentage point increase in $\text{Pr}(\hat{u}_j)$ is associated with a 0.09 increase in W_j/\hat{Y}_j^p , with the effect statistically significant at the 10 percent level. The effect rises in economic and statistical importance at higher levels of permanent income.

We view these estimates as a lower bound on the precautionary effects in this model. Given the relatively short time span of our data, particularly between 1989 and 1992, persistence in regional business conditions might mean that the time variation in the mix of regional job-loss probabilities is small enough that collinearity between the fixed-region effects and $\text{Pr}(\hat{u}_j)$ is reducing the estimate of $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$. To the extent that fixed-region effects are reducing $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ because they are correlated with elements of risk, the statistical significance of the regional variables actually is a sign of precautionary behavior. In contrast, to the extent that correlation between $\text{Pr}(\hat{u}_j)$ and fixed-region effects is due to nonprecautionary factors, then the results in table 9 indicate that precautionary motives are weaker than those estimated when region is excluded from C_j . Given that C_j includes income, industry, occupation, education, demographics, and other factors, we think this is unlikely. Note, though, that even if all the explanatory power of fixed-regional effects reflected nonprecautionary factors, the $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ estimates from the specification with region included in C_j still indicate the existence of at least some precautionary saving.³⁰

³⁰With regard to regional persistence, even though they find no permanent effect, Blanchard and Katz estimate that the transitory response of regional unemployment to a demand shock takes 6 years to disappear. With regard to nonprecautionary motives, one potential important correlation between wealth and region may be due to regional variation in house prices. However, this would likely work *against* finding $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j) > 0$ because regions with higher unemployment rates are likely to have depressed house prices. Finally, we also cannot rule out *a priori* that tastes differ systematically across regions. But Carroll, Rhee, and Rhee (1994) are unable to show that intercountry saving rate differentials are

IV.6. Sensitivity Checks

Table 10 presents a number of sensitivity checks on our results. All estimates are based on the model including the $\Pr(\hat{u}_j) \cdot \ln \hat{Y}_j^p$ interaction and the samples including the self employed and pooled across SCF years.

Reduced Instrument Set. First, we consider some potential omitted regressor biases that might be affecting the extended model. Recall that in the model with no interaction term, because all of the $Z_j^{\bar{u}S}$ (which are not used to estimate $\Pr(\hat{u}_j)$) are included in C_j , the consistency of b_u is unaffected by the possibility that the $Z_j^{\bar{u}S}$ might help predict $\Pr(\hat{u}_j)$. In the interaction model, strictly speaking, the coefficient on $\Pr(\hat{u}_j) \cdot \ln \hat{Y}_j^p$ would be unaffected by such unaccounted-for predictability only if we included the cross products of the $Z_j^{\bar{u}S}$ and the Z_j^{yS} in C_j . But, doing so would mean adding hundreds of variables to C_j . As a parsimonious check on the influence of some of these factors, we excluded the $Z_j^{\bar{u}S}$ from the first-stage estimate of $\ln \hat{Y}_j^p$: Because the resulting $\Pr(\hat{u}_j) \cdot \ln \hat{Y}_j^p$ no longer include any variation due to $Z_j^{\bar{u}S}$, b_{uy} should be less influenced by the (potential) appearance of the cross products of the $Z_j^{\bar{u}S}$ in C_j . Furthermore, note that a number of the control variables included in $Z_j^{\bar{u}S}$, such as the credit measures, are not usually used in regressions fitting permanent income proxies; thus, eliminating the $Z_j^{\bar{u}S}$ from the $\ln \hat{Y}_j^p$ regression also allows us to test the sensitivity of our results against a more traditional set of instruments for permanent income.

As can be seen in the first panel of table 10, use of this reduced instrument set has little qualitative effect on $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$. These results suggests that the cross products of the $Z_j^{\bar{u}S}$ are not highly correlated with the $\Pr(\hat{u}_j) \cdot \ln \hat{Y}_j^p$ used in the original interaction model. We can guess that the other higher-order terms excluded from C_j (the cross-products of the $Z_j^{\bar{u}S}$ and the Z_j^{uS}) probably do not greatly influence the estimated coefficients, either. And we can have more confidence that the inclusion of some nonstandard instruments in Z_j^{yS} does not greatly influence our empirical results.

Subaggregates of Wealth. We also consider how our results change when we use various subaggregates of household wealth as the numerator in our dependent variable.

attributable to taste or cultural differences, suggesting that it would be difficult to show that cultural factors *within* the United States strongly influence regional saving rates.

Engen and Gruber argue that households probably would allocate precautionary balances to assets that can be liquidated at little cost. Accordingly, we considered their "net financial assets" (NFA) grouping of checking accounts, savings accounts, certificates of deposit, stocks, bonds, and mutual funds less unsecured liabilities.³¹ As shown in the second column, the results change markedly: The values of $\partial(W_j/\hat{Y}_j^p)/\partial\text{Pr}(\hat{u}_j)$ are now close to zero at all income levels and are no longer statistically significant. Among the other independent variables (not shown), the largest changes are to the estimated coefficients on the homeowner dummy, which become much smaller and lose their statistical significance.

In the third column, we consider a wealth variable that excludes only real estate equity from our original measure of net worth. The estimated response of the ratio of wealth to permanent income to a 1 percentage point increase in $\text{Pr}(\hat{u}_j)$ is larger than that in the second column, but it remains statistically insignificant at all income levels. Next, we exclude only the equity in each households' principal residence from our original net worth variable. As shown in the fourth panel, these regressions also offer little evidence of a precautionary response of nonhousing wealth to an increased probability of becoming unemployed.³² Finally, we estimate the model using total net worth as the dependent variable but with the sample restricted to homeowners. The coefficient estimates (not shown) are about the same as those in table 8, but the t-statistics are all about 25 percent smaller, in line with the reduction in the sample size due to the exclusion of non-homeowners.

IV.7. Housing Wealth and Precautionary Saving

The results in table 10 point to home equity as a driving force behind the relationship between total net worth and employment risk presented in table 8. Our theoretical model--and most other theoretical analyses--yield little guidance as to what may be causing this

³¹NFA thus equals total net worth less the value of defined contribution pension plans, business assets and liabilities, real estate assets and liabilities, and assets and liabilities related to vehicles.

³²Using different data and methods, Carroll and Samwick (1997) also found that the relationship between total net worth and uncertainty was stronger than the relationship between uncertainty and the subaggregates of the balance sheet that they considered--namely, very liquid assets and non-housing, non-business wealth.

finding because they focus on a net worth measure comprised of only one type of liquid asset and one type of unsecured liability. Still, it seems counterintuitive that the precautionary response would be dependent on one of the seemingly least liquid of all household assets.

There are some possible theoretical justifications for consumers holding precautionary wealth in housing even though it is illiquid. One example is Laibson's (1997) model of consumers with hyperbolic time discount factors; such individuals will want to hold a buffer against income risk in the long run, but they are so impatient that they must force themselves to save by committing assets to instruments that are costly to liquidate. Alternatively, the illiquidity of housing may not be an important factor in determining its usefulness as precautionary saving. As Carroll and Samwick (1998) argue, less liquid assets may not be undesirable buffers if a household's chief concern is a high cost but low probability event such as job loss. Indeed, for such scenarios the transactions costs to tap any asset are small in expected value terms. The expected portfolio costs associated with holding housing wealth as a buffer against job loss also may not be large because state laws generally permit a household to keep its primary residence in the case of bankruptcy. Finally, it may simply be the case that housing wealth is more liquid than commonly supposed. For example, home equity lines of credit have made it quite easy for consumers to tap their housing wealth at rates well below those on credit cards. One problem with this explanation, however, is that when we repeated the table 9 sensitivity checks for each SCF year separately (not shown), we found that housing wealth was a driving factor even in 1983, when home equity lending was much less prevalent than today.

Our findings also may reflect a less direct routing of precautionary balances into housing wealth. The theoretical model in Carroll and Dunn (1997) shows that an increase in the probability of unemployment can cause some prudent homeowners to delay purchasing a new home because precautionary motives make them unwilling to incur the reduction in net worth caused by the large transactions costs associated with paying mortgage points and real estate commissions. If a significant portion of these costs were to be paid out of the equity in their existing house, then homeowners deferring purchases would be left with relatively more housing equity than those who went ahead with the upgrade. Thus, if consumers delay home

purchases when they face higher job-loss risk (as Dunn (1998) shows they do), this is likely to show up in our empirical results as a positive relationship between home equity and uncertainty.

Such behavior likely is reinforced by the supply side of the mortgage market. To compensate lenders for risk, households with high job-loss risk can face some combination of higher rates, larger downpayments, more strict mortgage insurance requirements, or higher origination charges than low-risk households.³³ These costs could cause a high-risk household to delay purchasing a home or to take out a smaller loan. Thus, all else equal, the pricing of risk in the mortgage market may force homeowners with higher unemployment probabilities to hold more housing equity than their low-risk counterparts.

V. Conclusion

This paper estimates the strength of the precautionary saving motive by relating the uncertainty associated with the risk of becoming unemployed to the cross-sectional distribution of household net worth. In doing so, we try to use best-practice techniques in choosing our uncertainty proxy, instrumental variables strategy, and an empirical specification that incorporates restrictions from a theoretical model. First, we consider the *ex ante* probability that a household head becomes unemployed because it may be a better measure of risk than the income or consumption variation proxies used in many previous papers. Second, we avoid relying on instruments which may be correlated with idiosyncratic tastes to identify variation in unemployment risk by using a variable (region) that presumably is uncorrelated with relevant household idiosyncratic characteristics. Finally, we use the inverse hyperbolic sine function to transform the wealth data, eliminating the need to exclude households with nonpositive net worth or to assume constant elasticities in the response of wealth to driving variables.

³³Because we have credit variables in C_j and limit our sample to household heads who have held the same job for the past three years, we probably have controlled for some of the important sources of heterogeneity in individual lending terms. On the other hand, lending terms show sufficient regional variation to suggest that more general regional economic conditions also affect the mortgage packages faced by the households in our sample. (Data on lending terms are published for many large metropolitan areas, but are not available broken down by the nine Census subdivisions we use as regional identifiers.)

Our empirical results suggest that households in the lowest permanent income groups do not engage in precautionary saving, but as income rises, precautionary behavior becomes significant in both economic and statistical terms. Our findings hold for samples that both include and exclude the self employed. However, the precautionary response generally does not show up for subaggregates of the household balance sheet that exclude housing wealth. While there are a number of reasons for precautionary saving behavior to be reflected in housing wealth--either directly through the potential to tap housing equity when needed or indirectly through the effect of delayed purchases of new homes--this dependence deserves closer scrutiny since housing is generally considered one of the least liquid assets that a family can hold.

Appendix

The Likelihood Function

The model we wish to estimate is:

$$g[W_j/Y_j^p, \theta] = \beta_0 + \beta_u \Pr(u_j) + \beta_y \ln Y_j^p + C_j \beta_c + \epsilon_j, \quad (19)$$

Assuming $\epsilon_j \sim N(0, \sigma^2)$, the log likelihood function of the W_j/Y_j^p corresponding to (19) is

$$L[\beta, \theta, \sigma^2] = K - (n_s/2) \ln \sigma^2 - (1/2\sigma^2) \sum_{j=1}^{n_s} \epsilon_j^2 - 1/2 \sum_{j=1}^{n_s} \ln[1 + \theta^2(W_j/Y_j^p)^2] \quad (20)$$

But, effectively, we maximize the likelihood function:

$$\begin{aligned} L[\beta, \theta, \sigma^2] &= K - (n_s/2) \ln \sigma^2 - (1/2\sigma^2) \sum_{j=1}^{n_s} e_j^2 - 1/2 \sum_{j=1}^{n_s} \ln[1 + \theta^2(W_j/\hat{Y}_j^p)^2] \\ e_j &= g[W_j/\hat{Y}_j^p, \theta] - \beta_0 - \beta_u \Pr(\hat{u}_j) - \beta_y \ln \hat{Y}_j^p - C_j \beta_c \\ &= \epsilon_j + \beta_u (\Pr(u_j) - \Pr(\hat{u}_j)) + \beta_y (\ln Y_j^p - \ln \hat{Y}_j^p) + g[W_j/\hat{Y}_j^p, \theta] - g[W_j/Y_j^p, \theta] \end{aligned} \quad (21)$$

First-Order Approximations

To estimate $\Pr(u_j)$ we posit the latent relationship $u_j^* = Z^U \alpha_u + v_j$, where v_j is distributed logistically and that 1) there exists a variable $u_{t+1|t}$ equal to 1 if a person is employed in period t and unemployed one year later and equal to 0 if a person is unemployed in both periods; 2) $u_{t+1|t} = 1$ if $u_j^* > 0$ and $u_{t+1|t} = 0$ if $u_j^* \leq 0$. This allows us to estimate $\Pr(\hat{u}_j) = \exp(Z_j^u \hat{\alpha}_u) / [1 + \exp(Z_j^u \hat{\alpha}_u)]$ using the $\hat{\alpha}_u$ estimated from a simple logistic regression.

Next, we assume that $\ln Y_j^p = Z_j^y \alpha_y$ and $\ln Y_j = \ln Y_j^p + \eta_j$, where Z_j^y and η_j are orthogonal. We regress the log of actual income on the Z_j^{yS} , and use the fitted values as our

estimates of $\ln \hat{Y}_j^p$.

Note that if P_Z is the projection matrix of the Z 's, $M_Z = I - P_Z$, and $\ln Y$ and Z^{yS} are the observations matrices on the $\ln Y_j$ and Z_j^{yS} , then $\ln Y - \ln \hat{Y}^p = \ln Y - \eta - P_Z \ln Y = M_Z \ln Y - \eta = \hat{\eta} - \eta$. Thus,

$$\ln Y_j^p = \ln \hat{Y}_j^p + (\hat{\eta}_j - \eta_j) \quad (22)$$

Next, taking first-order approximations of $\Pr(u_j)$ and $g[W_j/Y_j^p]$ about \hat{u}_j and $\ln \hat{Y}_j^p$, respectively, yields:

$$\begin{aligned} \Pr(u_j) &\approx \Pr(\hat{u}_j) + \frac{\partial \Pr(\hat{u}_j)'}{\partial \alpha_u} (\alpha_u - \hat{\alpha}_u) \\ &= \Pr(\hat{u}_j) (1 - \Pr(\hat{u}_j)) Z_j^{uS} (\alpha_u - \hat{\alpha}_u) \\ g[W_j/Y_j^p] &\approx g[W_j/\hat{Y}_j^p] + \frac{\partial g[W_j/\exp(\ln \hat{Y}_j^p)]}{\partial \ln Y_j^p} (\ln Y_j^p - \ln \hat{Y}_j^p) \\ &= g[W_j/\hat{Y}_j^p] - \frac{W_j/\hat{Y}_j^p}{[(\theta W_j/\hat{Y}_j^p)^2 + 1]^{1/2}} (\hat{\eta}_j - \eta_j) \end{aligned} \quad (23)$$

where $\partial \Pr(\hat{u}_j)/\partial \alpha_u$ is the vector of derivatives $\partial \Pr(\hat{u}_j)/\partial \alpha_{uk} = Z_{jk}^{uS} \Pr(\hat{u}_j)(1 - \Pr(\hat{u}_j))$, where α_{uk} is the k^{th} element of α . This means that e_j can be approximated by

$$\begin{aligned}
e_j \approx \epsilon_j + \beta_u [\Pr(\hat{u}_j) (1 - \Pr(\hat{u}_j))] Z_j^{uS'} (\alpha_u - \hat{\alpha}_u) + \\
\beta_y (\hat{\eta} - \eta) + \frac{W_j / \hat{Y}_j^p}{[(\theta W_j / \hat{Y}_j^p)^2 + 1]^{1/2}} (\hat{\eta}_j - \eta_j)
\end{aligned} \tag{24}$$

Consistency

Standard two-stage least squares assumptions insure that the probability limits of the products of $\ln \hat{Y}_j^p$ and C_j with the elements of e_j other than $(\Pr(u_j) - \Pr(\hat{u}_j))$ are zero. To show that the plims with $(\Pr(u_j) - \Pr(\hat{u}_j))$ are zero, the asymptotic approximations presented in equations (22) and (23) imply we need to show

$$\begin{aligned}
\text{plim } 1/n_s \sum_{j=1}^{n_s} Z_{rk}^{yS} Z_{jk}^{uS} \Pr(\hat{u}_j) (1 - \Pr(\hat{u}_j)) (\alpha_{uk} - \hat{\alpha}_{uk}) = \\
\text{plim } 1/n_s (\alpha_{uk} - \hat{\alpha}_{uk}) \sum_{j=1}^{n_s} Z_{rk}^{yS} Z_{jk}^{uS} \Pr(\hat{u}_j) (1 - \Pr(\hat{u}_j)) = 0
\end{aligned}$$

for all variables r in Z_j^{yS} and k in Z_j^{uS} , where n_s is the number of observations in the SCF sample and $\Pr(\hat{u}_j)$ is evaluated using Z_j^{uS} and $\hat{\alpha}$. Assuming $\sqrt{n_c}(\alpha - \hat{\alpha})$ converges in distribution to $N(0, \Omega)$, where n_c is the number of observations in the CPS sample, this condition will hold as long as $n_c \rightarrow \infty$ as $n_s \rightarrow \infty$ and $\text{plim } (1/n_s) \sum_j Z_{rk}^{yS} Z_{jk}^{uS} \Pr(\hat{u}_j) (1 - \Pr(\hat{u}_j))$ is bounded, which will be the case if the Z_j^S have finite second moments.

The other nonstandard consistency conditions to consider are the probability limits of the product of $\Pr(\hat{u}_j)$ with the $(\hat{\eta}_j - \eta_j)$ terms in e_j . Given the approximations in (22) and (23), it suffices to show $\text{plim } 1/n_s \sum_j (\hat{\eta}_j - \eta_j) Z_{jk}^{uS} \Pr(u_j) (1 - \Pr(u_j)) (\alpha_{uk} - \hat{\alpha}_{uk}) = 0$ and $\text{plim } 1/n_s \sum_j -W_j / \hat{Y}_j^p [(\theta W_j / \hat{Y}_j^p)^2 + 1]^{-1/2} (\hat{\eta}_j - \eta_j) Z_{jk}^{uS} \Pr(u_j) (1 - \Pr(u_j)) (\alpha_{uk} - \hat{\alpha}_{uk}) = 0$ for all k . Given that $\sqrt{n_c}(\alpha - \hat{\alpha})$ converges in distribution to $N(0, \Omega)$ and $\text{plim } 1/n_s \sum_j Z_{jk}^{uS} (\hat{\eta}_j - \eta_j) = 0$

(since $Z_{jk}^{uS} \subset Z_{jk}^{yS}$), these conditions will hold if $\text{plim } 1/n_s \sum_j \text{Pr}(\hat{u}_j) (1 - \text{Pr}(\hat{u}_j))$ and $\text{plim } 1/n_s \sum_j W_j/\hat{Y}_j^p [(\theta W_j/\hat{Y}_j^p)^2 + 1]^{-1/2} \text{Pr}(\hat{u}_j) (1 - \text{Pr}(\hat{u}_j))$ are bounded. The first condition is trivial, the second condition holds with θ bounded.

Standard Error and Test Statistic Adjustments

The results in Pagan (1984) indicate that the standard errors calculated from the Hessian of (20) will be too small. In particular, the correct standard errors must account for the estimated regressors $\text{Pr}(\hat{u}_j)$ and $\ln \hat{Y}_j^p$. To calculate the standard errors, we first add the auxiliary assumption that moments estimated from the CPS are independent of moments estimated from the SCF (see Angrist and Krueger 1992). Next, calculations along the lines of those in Pagan and Angrist and Krueger indicate that $\text{plim } 1/n_s \sum_j e_j^2 = \sigma^2 + (\beta_y - \text{dg})^2 \sigma_\eta^2 + \beta_u^2 \partial \text{Pr}(\hat{u}_j)/\partial \alpha' \Omega \partial \text{Pr}(\hat{u}_j)/\partial \alpha$, where $\text{dg} = \text{plim } 1/n_s \sum_j -W_j/\hat{Y}_j^p [(\theta W_j/\hat{Y}_j^p)^2 + 1]^{-1/2}$, but that $1/n_s \sum_j \hat{e}_j^2$ provides a consistent estimate of only σ^2 . Accordingly, the standard errors produced by the Hessian matrix need to be blown up by $[\text{plim } 1/n_s \sum_j e_j^2] / [1/n_s \sum_j \hat{e}_j^2]$ in order to calculate the correct standard errors; this adjustment is made to the statistics reported in the text. We estimate $\text{plim } 1/n_s \sum_j e_j^2$ using the obvious element-by-element estimates from the first and second stage regressions.

For the model with the interaction between $\text{Pr}(\hat{u}_j)$ and $\ln \hat{Y}_j^p$, the term $\beta_{uy}(\text{Pr}(u_j) \ln Y_j^p - \text{Pr}(\hat{u}_j) \ln \hat{Y}_j^p)$ must be added to e_j . The interaction term is accommodated in the standard error calculations by adding the estimate of $\beta_{uy}^2 (\sigma_\eta^2 \cdot \partial \text{Pr}(\hat{u}_j)/\partial \alpha' \Omega \partial \text{Pr}(\hat{u}_j)/\partial \alpha)$ to the blow-up factor. For example, in the model with the income interaction, in both 1983 and the pooled sample the adjustments boost the standard errors by about 25 percent. In 1989 and 1992, the adjustments boost the standard errors by about 50 and 70 percent, respectively.

Replicants

In order to more accurately represent the variance of imputed variables, the 1989 and 1992 SCFs each contain 5 replicants. That is, each observation is repeated five times in the data set, with noise added to any imputed series in accordance with the variance of the underlying series. As discussed in Montalto and Sung (1996) and Kennickell and Woodburn (1997), econometric estimates need to be corrected for the use of replicants. Their procedure requires estimating the models separately for each replicant, using the average of the

coefficient values for point estimates, and calculating standard errors according to a formula that accounts for both within-replicant and cross-replicant variation in the coefficients. As a first approximation to these more complicated corrections, researchers can simply use the coefficients from the model estimated using all of the data with the standard errors multiplied by the square root of the number of replicants--this essentially assumes that the true sample size equals that of one replicant. We make this simple adjustment to all of the standard errors and test statistics for 1989 and 1992 presented in the paper: In the test cases we ran, this approximation produced essentially the same coefficient estimates and standard errors as the more complicated corrections. When we ran the pooled models, we used only one replicant from the 1989 and 1992 SCFs.

References

- Angrist, Joshua D. and Alan 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 (June): 328-26.
- Blanchard, Olivier and Lawrence Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 1992:1: 1-61.
- Burbidge, John B., Lonnie Magee and A. Leslie Robb. 1988. "Alternative Transformations to Handle Extreme Values of the Dependent Variable." *Journal of the American Statistical Association* 83 (March): 123-27.
- Caballero, Ricardo J. 1990. "Consumption Puzzles and Precautionary Saving." *Journal of Monetary Economics* 25 (January): 113-36.
- Campbell, John Y. and N. Gregory Mankiw, 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence" in *NBER Macroeconomics Annual 1989*, edited by Olivier Jean Blanchard and Stanley Fischer (MIT Press: Cambridge, MA): 185-216.
- Carrington, William J. 1993. "Wage Losses for Displaced Workers: Is It Really the Firm that Matters?" *Journal of Human Resources* 28 (Summer): 435-62.
- Carroll, Christopher D. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 1992:2: 61-156.
- Carroll, Christopher D. 1994. "How Does Future Income Affect Current Consumption?" *Quarterly Journal of Economics* 109 (February): 111-48.
- Carroll, Christopher D. 1996. "Buffer Stock Saving: Some Theory." Manuscript. Department of Economics, Johns Hopkins University.
- Carroll, Christopher D., Rhee, Byung-Kun and Changyong Rhee. 1994. "Are There Cultural Effects on Saving? Some Cross-Sectional Evidence." *Quarterly Journal of Economics* 109: 685-99.
- Carroll, Christopher D. and Wendy E. Dunn. 1997. "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets" in *NBER Macroeconomics Annual 1997*, edited by Benjamin S. Bernanke and Julio Rotemberg (MIT Press: Cambridge, MA).
- Carroll, Christopher D. and Andrew Samwick. 1997. "The Nature of Precautionary Wealth."

Journal of Monetary Economics 40 (September): 41-71.

- Carroll, Christopher D. and Andrew Samwick. 1998. "How Important is Precautionary Saving?" *Review of Economics and Statistics* 80 (August): 410-19.
- Clark, Kim B. and Lawrence H. Summers. 1979. "Labor Market Dynamics and Unemployment: A Reconsideration." *Brookings Papers on Economic Activity: 1979:1*: 13-60.
- Curtin, Richard T., F. Thomas Juster, James N. Morgan. 1989. "Survey Estimates of Wealth: An Assessment of Quality" in *The Measurement of Saving, Investment, and Wealth*, edited by Robert E. Lipsey and Helen Stone Tice (University of Chicago Press: Chicago, IL).
- Diamond, P.A. and J.A. Hausman. 1984. "Individual Retirement and Savings Behavior." *Journal of Public Economics* 23: 81-114.
- Dunn, Wendy E. 1998. "Unemployment Risk, Precautionary Saving and Durable Goods Purchase Decisions." Manuscript. (October).
- Dynan, Karen E. 1993. "How Prudent Are Consumers?" *Journal of Political Economy* 101 (December): 1104-13.
- Dynan, Karen E., Jonathan Skinner and Stephen P. Zeldes. 1997. "Do the Rich Save More?" Manuscript (April).
- Engen, Eric M. and Jonathan Gruber. 1997. "Unemployment Insurance and Precautionary Saving." Manuscript. (March).
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. (Princeton University Press: Princeton, NJ).
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 1997. "Consumption Over the Life Cycle." Social Systems Research Institute working paper No. 9722 (November).
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1992. "Earnings Uncertainty and Precautionary Saving." *Journal of Monetary Economics* 30 (November): 307-37.
- Gruber, Jonathan. 1994. "The Consumption Smoothing Benefits of Unemployment Insurance." National Bureau of Economic Research Working Paper no. 4750 (May).
- Hubbard, R. Glenn, Jonathan Skinner and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40 (June): 59-126.

- Hubbard, R. Glenn, Jonathan Skinner and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (April): 360-99.
- Johnson, N. L. 1949. "Systems of Frequency Curves Generated by Methods of Translation." *Biometrika* 36: 149-76.
- Kennickell, Arthur B. and R. Louise Woodburn. 1997. "Consistent Weight Design for the 1989, 1992, and 1995 SCFs, and the Distribution of Wealth." Unpublished mimeo. (June).
- King, M.A. and L-D. L. Dicks-Mireaux. 1982. "Asset Holdings and the Life-Cycle." *The Economic Journal* 92 (June): 247-67.
- Kimball, Miles S. 1990. "Precautionary Saving in the Small and in the Large." *Econometrica* 58 (January): 53-73.
- Kuehlwein, Michael. 1991. "A Test for the Presence of Precautionary Saving." *Economics Letters* 37 (December): 471-75.
- Laibson, David I. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (May): 443-78.
- Lillard, Lee and Lynn Karoly. 1997. "Income and Wealth Accumulation over the Lifecycle." Manuscript.
- Lusardi, Annamaria. 1997. "Precautionary Saving and Subjective Earnings Variance." *Economic Letters* 57 (December): 319-26.
- Lusardi, Annamaria. 1998. "On the Importance of Precautionary Saving." *American Economic Review* 88 (May): 448-53.
- Montalto, Catherine and Jaimie Sung. 1996. "Multiple Imputation in the 1992 Survey of Consumer Finances." *Financial Counseling and Planning* 7: 133-46.
- Normandin, Michel. 1994. "Precautionary Saving: An Explanation for the Excess Sensitivity of Aggregate Consumption." *Journal of Business and Economics Statistics* 12 (April): 205-19.
- Pagan, Adrian. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25 (February): 221-47.
- Starr-McCluer, Martha. 1996. "Health Insurance and Precautionary Saving."

American Economic Review 86 (March): 285-95.

Welch, Finis R. 1993. "dm11: Matching the Current Population Survey." *Stata Technical Bulletin Reprints* 2: 4-5.

Zeldes, Stephen P. 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104 (May): 275-98.

Table 1
Steady State Distribution of Wealth

	Unemp. Rate Equals 2%	Unemp. Rate Equals 4%	Change Associated with Unemployment Rate Rising from 2% to 4%	
			Absolute	Percent
Fraction of Households with:				
$W_t/Y_t^P < 0$.133	.098	-.035	-25%
$W_t/Y_t^P = 0$.112	.075	-.037	-33%
$W_t/Y_t^P > 0$.755	.827	.072	10%
Mean W_t/Y_t^P of:				
All Households	.066	.119	.053	80%
20th-30th W_t/Y_t^P Percentile	.007	.055	.048	686%
40th-50th W_t/Y_t^P Percentile	.059	.119	.060	102%
60th-70th W_t/Y_t^P Percentile	.110	.186	.076	69%
80th-90th W_t/Y_t^P Percentile	.187	.272	.085	45%

Table 2
Wealth Holdings for Households with Different Employment Histories

	All Households	Unemployed in Year t	Unemployed in Year $t-1$	Unemployed in Year $t-2$	Not Unemp. in Last 3 Years
Fraction of Households with:					
$W_t/Y_t^p < 0$.133	1.000	.975	.525	.041
$W_t/Y_t^p = 0$.112	.000	.025	.275	.113
$W_t/Y_t^p > 0$.755	.000	.000	.200	.840
Mean W_t/Y_t^p	.066	-.331	-.116	-.017	.095

Note: All calculations assume an unemployment rate of 2 percent.

Table 3
SCF Summary Statistics

	1983	1989	1992
Mean net worth	\$151,363	\$143,393	\$100,798
Median net worth	\$68,494	\$67,545	\$54,392
Net worth, 95 th percentile	\$590,320	\$536,955	\$360,119
Median ratio of net worth to income	1.59	1.54	1.37
Percent with net worth of less than 1 month's income	8.0%	10.1%	9.1%
Percent with zero or negative net worth	3.7%	5.3%	5.8%
Median net worth of zero or negative net worth households	-\$1,338	-\$1,450	-\$2,475
Number of households	1689	1025	1032

Notes: Sample includes all households in SCF area with heads between the ages of 20 and 65 who have been employed at the same job for at least three years. Sample excludes top and bottom 0.1 percentiles of wealth or income. All calculations were done on a weighted basis. Dollar amounts are expressed in 1992 dollars.

Table 4
CPS First-Stage Estimation Results^a
Dependent Variable = 1 if Unemployed in $t+1$; 0 Otherwise

	1983	1989	1992
F-tests: ^b			
Occupation	0.007	0.134	0.008
Industry	0.000	0.000	0.000
Region	0.012	0.000	0.000
Education	0.000	0.000	0.001
White	0.003	0.001	0.007
Female Head	0.051	0.105	0.011
Age	0.168	0.798	0.980
Age * Occupation	0.011	0.107	0.097
Age * Education	0.069	0.042	0.061
Mean $\Pr(\hat{u})$	0.027	0.019	0.022
Std. Dev. $\Pr(\hat{u})$	0.021	0.015	0.017
Number of Observations	59,252	60,026	63,351

^a Sample includes all individuals who were employed in t , including the self-employed.

^b P-values.

Table 5
SCF First-Stage Estimation Results for Group Income^a
Dependent Variable = log(*Y*)

	1983	1989	1992
F-tests: ^b			
Occupation	0.684	0.361	0.847
Industry	0.009	0.111	0.084
Region	0.000	0.000	0.115
Education	0.417	0.742	0.016
White	0.667	0.061	0.463
Female Head	0.000	0.002	0.000
Home Owner	0.000	0.000	0.000
# Children	0.055	0.593	0.788
# Earners	0.000	0.000	0.000
Marital Status	0.000	0.008	0.058
Defined Benefit Coverage	0.000	0.000	0.009
Turned Down for Credit	0.331	0.076	0.897
Problems Servicing Debt	0.113	0.920	0.033
Age	0.000	0.000	0.016
Age * Occupation	0.008	0.059	0.759
Age * Education	0.376	0.553	0.248
Adjusted R ²	0.46	0.41	0.28
Number of households	1688	1025	1032

^a Sample includes all individuals who were employed in *t*, including the self-employed.

^b P-values.

Table 6
SCF Second-Stage Estimation Results

$$g[W/\hat{Y}^P, \theta] = \beta_0 + \beta_u \Pr(\hat{u}) + \beta_y \ln \hat{Y}^P + C \beta_c C + \epsilon$$

	1983		1989		1992	
	Including self-employed	Excluding self-employed	Including self-employed	Excluding self-employed	Including self-employed	Excluding self-employed
Pr(\hat{u})	0.52 (0.50)	0.23 (0.22)	-0.67 (-0.28)	-1.03 (-0.33)	0.03 (0.01)	-0.60 (-0.23)
ln \hat{Y}^P	0.09 (1.24)	0.12 (1.62)	0.23 (1.89)	0.44 (2.44)	0.40 (1.55)	0.26 (1.58)
Number of Earners	-0.04 (-2.24)	-0.05 (-2.58)	-0.08 (-2.05)	-0.11 (-1.94)	-0.18 (-2.19)	-0.11 (-1.99)
Owens Home	0.28 (9.00)	0.27 (8.61)	0.32 (5.80)	0.32 (4.89)	0.28 (3.46)	0.27 (4.55)
Defined Benefit Pension	-0.06 (-3.43)	-0.04 (-2.02)	-0.10 (-2.60)	-0.08 (-1.74)	-0.08 (-1.39)	-0.02 (-0.49)
Turned Down for Credit	-0.05 (-2.78)	-0.07 (-3.15)	-0.03 (-0.74)	-0.10 (-2.14)	-0.06 (-1.18)	-0.07 (-1.72)
Θ	4.56	4.56	3.86	3.3	3.31	3.26
Overidentify restr. p-value	0.38	0.57	0.32	0.54	0.83	0.44
$\frac{\partial W_j / \hat{Y}_j^P}{\partial \Pr(\hat{u}_j)}$ at median W/\hat{Y}^P	0.04 (0.50)	0.01 (0.22)	-0.05 (-0.28)	-0.06 (-0.32)	0.00 (0.01)	-0.02 (-0.23)
Number of Households	1689	1441	1025	835	1032	880

Note: Asymptotic t-statistics shown in parentheses. Regressions also include age, education, age*education, marital status, retirement status, race, sex, occupation, occupation*age, industry, and number of children.

Table 7
SCF Second-Stage Estimation Results

$$g[W/\hat{Y}^P, \theta] = \beta_0 + \beta_u \Pr(\hat{u}) + \beta_{uy} \Pr(\hat{u}) * \ln \hat{Y}^P + \beta_y \ln \hat{Y}^P + C \beta_c C + \epsilon$$

	1983		1989		1992	
	Including self-employed	Excluding self-employed	Including self-employed	Excluding self-employed	Including self-employed	Excluding self-employed
$\Pr(\hat{u})$	-28.32 (-2.53)	-26.75 (-2.24)	-57.22 (-1.99)	-66.09 (-1.92)	-67.77 (-1.99)	-50.46 (-1.69)
$\Pr(\hat{u}) * \ln \hat{Y}^P$	2.96 (2.58)	2.76 (2.27)	5.72 (1.98)	6.47 (1.90)	6.85 (1.99)	5.05 (1.68)
$\ln \hat{Y}^P$	-0.03 (-0.37)	0.01 (0.12)	0.09 (0.62)	0.25 (1.17)	0.17 (0.66)	0.11 (0.57)
$\frac{\partial W_j / \hat{Y}_j^P}{\partial \Pr(\hat{u}_j)}$ \hat{Y}^P percentile:						
	0.03 (0.95)	0.03 (0.51)	-0.01 (-0.18)	-0.04 (-0.60)	-0.01 (-0.15)	-0.02 (-.37)
	0.11 (1.84)	0.04 (1.34)	0.11 (0.74)	0.06 (0.30)	0.16 (0.75)	0.03 (0.38)
10th	0.29 (2.14)	0.24 (1.67)	0.27 (1.18)	0.15 (0.74)	0.17 (1.17)	0.10 (0.79)
30th	0.39 (2.34)	0.33 (1.87)	0.44 (1.40)	0.24 (1.03)	0.39 (1.46)	0.14 (1.05)
50th	0.52 (2.47)	0.42 (2.01)	0.71 (1.61)	0.55 (1.25)	0.77 (1.67)	0.41 (1.26)
70th						
Identify restr. p-value	0.50	0.71	0.32	0.61	0.89	0.46

Note: Asymptotic t-statistics shown in parentheses.

90th

Table 8
SCF Second-Stage Estimation Results: 1983, 1989, 1992 Pooled Data

	Model without $\Pr(\hat{u}) * \ln \hat{Y}^P$ interaction		Model with $\Pr(\hat{u}) * \ln \hat{Y}^P$ interaction	
	Including self-employed	Excluding self-employed	Including self-employed	Excluding self-employed
$\Pr(\hat{u})$	0.95 (1.34)	0.74 (1.00)	-30.64 (-3.34)	-27.15 (-2.80)
$\Pr(\hat{u}) * \ln \hat{Y}^P$	--	--	3.15 (3.45)	2.77 (2.88)
$\ln \hat{Y}^P$	0.07 (1.85)	0.13 (2.83)	0.00 (0.04)	0.05 (1.02)
$\frac{\partial W_j / \hat{Y}_j^P}{\partial \Pr(\hat{u}_j)}$ \hat{Y}^P percentile:				
10 th	0.02 (1.34)	0.01 (1.00)	0.01 (0.64)	0.00 (0.26)
30 th	0.03 (1.34)	0.02 (1.00)	0.06 (2.39)	0.05 (1.76)
50 th	0.06 (1.34)	0.04 (1.00)	0.17 (3.07)	0.12 (2.40)
70 th	0.08 (1.34)	0.05 (1.00)	0.30 (3.43)	0.21 (2.75)
90 th	0.09 (1.34)	0.07 (1.00)	0.43 (3.62)	0.36 (2.93)
Overidentify restr. p-value	0.31	0.59	0.34	0.57
Number of households	3746	3157	3746	3157

Note: Asymptotic t-statistics shown in parentheses.

Table 9
Identification Issues: Alternative Excluded Controls
1983, 1989, 1992 Pooled Data

	Instrument for $\ln \hat{Y}^P$ that is Excluded from Controls				No Excluded Instrument
	Region	Occupation	Education	Industry	
$\Pr(\hat{u})$	-30.64 (-3.34)	-20.15 (-2.39)	-26.20 (-2.81)	-28.58 (-3.49)	-23.85 (-2.82)
$\Pr(\hat{u}) * \ln \hat{Y}^P$	3.15 (3.45)	1.95 (2.33)	2.47 (2.65)	2.95 (3.67)	2.38 (2.82)
$\ln \hat{Y}^P$	0.00 (0.04)	-0.01 (-0.25)	0.06 (1.67)	-0.12 (-2.86)	-0.09 (-1.86)
$\frac{\partial W_j / \hat{Y}_j^P}{\partial \Pr(\hat{u}_j)}$ \hat{Y}^P percentile:					
10th	0.01 (0.64)	-0.02 (-1.34)	-0.03 (-2.69)	0.01 (1.01)	-0.00 (-0.40)
30th	0.06 (2.39)	-0.00 (-0.11)	-0.03 (-1.02)	0.06 (3.52)	0.03 (1.01)
50th	0.17 (3.07)	0.03 (0.60)	-0.00 (-0.06)	0.17 (4.38)	0.09 (1.64)
70th	0.30 (3.43)	0.09 (1.09)	0.05 (0.59)	0.29 (4.65)	0.17 (2.05)
90th	0.43 (3.62)	0.15 (1.45)	0.13 (1.09)	0.41 (4.66)	0.26 (2.33)
Overidentify restr. p-value	0.34	0.04	0.29	0.06	0.22
Exclusion p-value	0.01	0.00	0.00	0.00	n.a.

Note: Asymptotic t-statistics shown in parentheses.

Table 10
Sensitivity Checks: 1983, 1989, 1992 Pooled Data

	Reduced Instruments for $\ln \hat{Y}^P$	Alternative Dependent Variables			
		Net Financial Assets	NW excl. Real Estate	NW excl. Primary Res.	
$\Pr(\hat{u})$	-25.37 (-2.70)	0.14 (0.10)	-6.98 (-1.15)	-9.23 (-1.53)	
$\Pr(\hat{u}) * \ln \hat{Y}^P$	2.62 (2.81)	-0.02 (-0.15)	0.67 (1.11)	6.04 (0.90)	
$\ln \hat{Y}^P$	0.02 (0.47)	0.01 (1.65)	-0.00 (-0.01)	0.01 (0.18)	
$\frac{\partial W_j / \hat{Y}_j^P}{\partial \Pr(\hat{u})}$ \hat{Y}^P percentile:					
	10th	0.02 (0.97)	-0.00 (-0.63)	-0.01 (-0.65)	-0.01 (-0.79)
	30th	0.09 (2.39)	-0.00 (-0.68)	-0.00 (-0.07)	-0.00 (-0.00)
	50th	0.17 (2.86)	-0.00 (-0.64)	0.00 (0.27)	0.01 (0.44)
	70th	0.30 (3.05)	-0.00 (-0.58)	0.01 (0.49)	0.03 (0.72)
	90th	0.44 (3.15)	-0.01 (-0.51)	0.02 (0.66)	0.06 (0.94)
Overidentify restr. p-value	0.14	0.18	0.23	0.11	
Region exclusion p-value	0.00	0.06	0.87	0.21	

Note: Asymptotic t-statistics shown in parentheses.


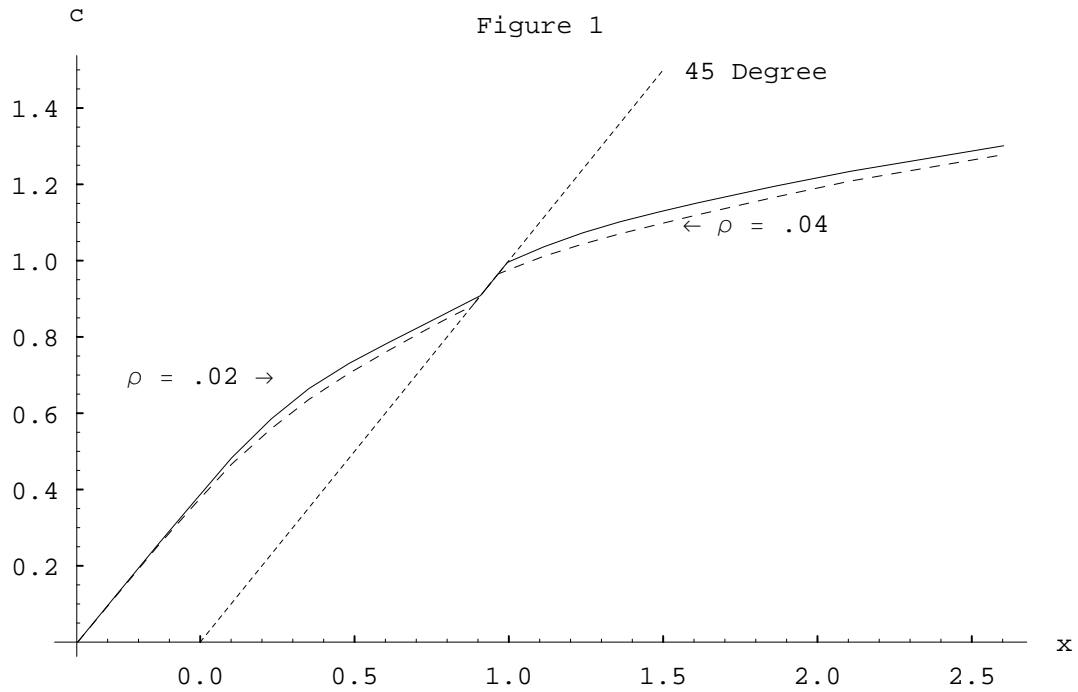


Figure 3

