UNEMPLOYMENT RISK AND PRECAUTIONARY WEALTH: EVIDENCE FROM HOUSEHOLDS' BALANCE SHEETS

Christopher D. Carroll, Karen E. Dynan, and Spencer D. Krane*

Abstract—This paper examines precautionary behavior by relating job-loss risk to household net worth. We use existing best practice and some new strategies to deal with some problematic issues inherent in this literature regarding proxying uncertainty, instrumentation, and incorporating theoretical restrictions. We do not find precautionary variation in the wealth holdings of households with low permanent income, but do find precautionary effects for moderate and higher-income households. When the dependent variable is total net worth, these findings are robust to several alternative specifications. But we do not find precautionary responses in subaggregates of wealth that exclude home equity.

I. Introduction

A. Overview

Many studies have noted the potential economic importance of precautionary saving. Caballero (1990) and Normandin (1994) have pointed out that it may be able to explain certain stylized facts about aggregate consumption such as excess sensitivity to movements in income. Carroll (1992) and Carroll and Dunn (1997) have argued that precautionary saving is an important driving force for consumption-led business cycles. And simulations in Hubbard, Skinner, and Zeldes (1994) suggest that it 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 (2001), and Lusardi (1997, 1998) find evidence of a statistically significant and economically important precautionary motive. The coefficient of relative risk aversion estimated by Gourinchas and Parker (2002) also implies substantial precautionary wealth.

The mixed findings may reflect a number of inherent 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 from theoretical models. This paper adds to the work on precautionary saving by building on best-practice techniques and bringing to bear new strategies to address problems in each of these categories.

B. 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 of a household's income (Carroll, 1994; Carroll and Samwick, 1997, 1998) or the variability of its expenditures (Dynan, 1993; Kuehlwein, 1991). But, as Lusardi (1997, 1998) and Guiso, Jappelli, and Terlizzese (1992) have pointed out, variability measures may be poor uncertainty proxies because they can contain large controllable elements. For example, a tenured college professor who, by choice, teaches or consults every other summer may have 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 variability of quarterly expenditures between households may simply reflect differences in preferences towards regular seasonal outlays such as vacations or school tuition.

Our measure of uncertainty is the probability of job loss—specifically, the estimated probability that a consumer who currently is employed will be unemployed one year hence.[1] This represents a potential major interruption to income over which households generally have little influence, and thus should provide a much cleaner signal of the uncertainty faced by a household than variability of income or expenditures.

Because no single source contains high-quality information on household-level income, wealth, and job loss, we use a source of good data on employment and unemployment, the Current Population Survey (CPS), to estimate unemployment risk based on observable household characteristics. We use these results to predict job-loss risk for households in a data set with good information on income and wealth, the Survey of Consumer Finances (SCF). We then relate this predicted job-loss risk to household net worth in the SCF.

Received for publication September 11, 2001. Revision accepted for publication May 20, 2002.

* The Johns Hopkins University and NBER; Federal Reserve Board; and Federal Reserve Bank of Chicago, respectively.

We are grateful to Dan Bergstresser, Peter Brady, Martin Browning, Eric Engen, Arthur Kennickell, Steve Lumpkin, Annamaria Lusardi, Martha Starr-McCluer, Valerie Ramey, and seminar participants at the American Economic Association Annual Meetings, the Johns Hopkins University, the NBER Summer Institute, Georgetown University, MIT, and the Federal Reserve Bank of Chicago for helpful discussions. We also thank Dan Bergstresser, David Brown, 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 System or its staff.

[1] Lusardi (1998) and Engen and Gruber (2001) also consider measures of job loss: 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.

The Review of Economics and Statistics, August 2003, 85(3): 586–604

C. Instrumental Variables Strategy

Because uncertainty is measured with significant error, most studies instrument their uncertainty proxy using variables such as 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 uncertainty.

Finding an appropriate instrument to exclude is difficult. For example, suppose that more risk-averse consumers both hold more precautionary wealth and choose occupations with lower job-loss risk.[2] 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 proxy will be biased because of correlation between occupation-instrumented job-loss risk and the unmeasured risk aversion in the error term. Similar arguments can be made regarding excluding educational achievement or industry of work. And we find some empirical evidence that these concerns may be warranted.[3]

To avoid this identification problem, we include most variables that have in the past been used as excluded instruments for uncertainty (for example, education, occupation, and industry) as independent controls in our econometric model of precautionary wealth. This requires us to find some other instrument that is correlated with job-loss risk to exclude from the control variables. We use the region in which the household resides. The large variation in regional economic conditions suggests that region will be significantly correlated with job-loss risk. In addition, if we assume 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 determinants of wealth.[4]

D. Insights from a Structural Model

We first solve a theoretical model of precautionary saving. It shows that in the steady state, permanently high-risk households hold higher average precautionary wealth than permanently low-risk households. But, following an increase in unemployment risk, the precautionary balances of those who remain employed build up only gradually while the wealth of households actually entering unemployment falls sharply. This means that empirical tests may be affected by recent job loss; even if all consumers had a precautionary motive, a regression of wealth on job-loss risk might find a negative coefficient if shocks that boost both actual unemployment and perceived risk occurred shortly before the period of observation. The model also suggests a simple sample restriction to avoid this problem, because it unambiguously predicts a positive relationship between current unemployment risk and wealth among households that have not recently experienced a spell of unemployment.

In addition, the theoretical model provides guidance on the functional form for our empirical specification. Many past studies have dealt with the extreme skewness of the wealth distribution by using the logarithm rather than the level of wealth as the dependent variable. This transformation requires dropping or making ad hoc adjustments to the nonnegligible share of households with nonpositive wealth, and it imposes constant elasticities. However, our model shows that low or negative wealth may be consistent with optimal behavior and that precautionary responses may vary with wealth and income. As an alternative, we transform wealth with the inverse hyperbolic sine function; this still downweights large values, but can be applied to positive, zero, and negative numbers and does not impose constant elasticities.[5]

E. 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 low permanent income to boost their net worth significantly, 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 estimate a significant precautionary effect in broad measures of wealth, but not in narrower subaggregates that exclude home equity. 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 consider a stylized model of household behavior in order to provide 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. To isolate precautionary responses, we consider a model in which there is no saving for retirement or other nonprecautionary purposes.

[2] In the 1989 Italian *Survey of Household Income and Wealth*, Lusardi (1997) finds that one-half of households mention job security as a reason for choosing their jobs.

[3] A related problem occurs when proxying uncertainty with insurance coverage: Risk-averse households may both save more and obtain more insurance. Starr-McCluer addresses this problem by instrumenting health insurance coverage with the percentage of the local workforce employed by large firms. Engen and Gruber argue that because unemployment insurance coverage is determined by state policy, it is probably largely exogenous to the individual household's saving behavior.

[4] Engen and Gruber provide a detailed discussion of issues concerning instrument validity. They also argue that regional variables likely satisfy exogeneity requirements.

[5] We thank Martin Browning for suggesting this transformation.

A. *The Household's Problem*

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

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

subject to

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

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 1 plus the rate of time preference, and R_{t+1} is 1 plus the rate of return, r_{t+1}.

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

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

where γ is the coefficient of relative risk aversion. Since this utility function has a positive third derivative, a mean-preserving spread in consumption uncertainty raises expected marginal utility, so 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, 1990).

The process generating household income, Y_t, is

$$Y_t = Y_t^P V_t, \tag{4}$$

$$Y_{t+1}^P = G Y_t^P N_{t+1}. \tag{5}$$

Y_t^P is a permanent component, which is independent of V_t, a transitory shock to income.[6] N_{t+1} is a serially uncorrelated shock log-normally distributed with variance σ_n^2, so log Y_t^P is a random walk with drift $g = \log G$. Our focus is on 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. With probability $1 - \rho$ the household is employed, in which case V_t is log-normally distributed with variance σ_v^2 and mean $(1 - \rho V_{\min})/(1 - \rho)$. (This ensures $E_t V_{t+1} = 1$, so that changes in ρ affect the variance but not the expected value of income.)

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

$$W_t \equiv X_t - C_t. \tag{6}$$

In the real world, many households' net worth essentially is zero, and many others hold negative levels of wealth (see section IIIB below). If $V_{\min} = 0$, no optimizing household would hold $W_t \leq 0$, because of the possibility that $C_{t+1} = 0$ and $U'(C_{t+1}) = \infty$ (see Zeldes, 1989). Thus, in our model, $V_{\min} > 0$ is necessary to characterize the distribution of wealth. To generate a clustering of net worth around zero, our consumers face different interest rates for lending and borrowing: Consumers ending period t with positive wealth earn a return that applies to lenders, R^{lend}, whereas those who borrow (end the period with $W_t < 0$) pay a higher rate, R^{borrow}. The Euler equations

$$u'(C_t) = \beta R^{\text{lend}} E_t[u'(C_{t+1})], \tag{7}$$

$$u'(C_t) = \beta R^{\text{borrow}} \mathrm{E}_t[u'(C_{t+1})] \tag{8}$$

describe the behavior of lenders and borrowers, respectively. Because $R^{\text{borrow}} > R^{\text{lend}}$, 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 $W_t = 0$.

B. *Parameterization*

The time period is one year. We set $\sigma_n^2 = \sigma_v^2 = 0.01$ and $g = 0.03$, consistent with evidence from the Panel Study on Income Dynamics (Carroll, 1992). We set $\beta = 0.03$ and $\gamma = 0.2$. The baseline probability of job loss, ρ, is 0.02.[7] We set $V_{\min} = 0.2$, so that the social safety net ensures resources equal to 20% of normal income. Finally, we assume $R^{\text{borrow}} = 0.15$ and $R^{\text{lend}} = 0.03$, close, respectively, to the interest rate on many credit cards and the after-tax rate of return on Treasury bills. These parameter values ensure that households will not desire to accumulate assets without bound and that the model has a steady state.

C. *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, 2001). The solid line in figure 1 depicts the relationship between x_t and c_t under our baseline parameters. 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 consumers have suffered 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

[6] Consistent with Friedman (1957), Y_t^P is the path around which earnings exhibit transitory fluctuations, as opposed to the discounted value of future cash flows. V_t affects cash flow, but not Y_t^P.

[7] We chose a probability lower than the actual unemployment rate to offset the fact that we have normalized unemployment spells to last one year, as compared with an actual average length of several months (see, for example, Clark and Summers, 1979).

FIGURE 1.—CONSUMPTION FUNCTIONS FOR $\rho = 0.02$ AND $\rho = 0.04$

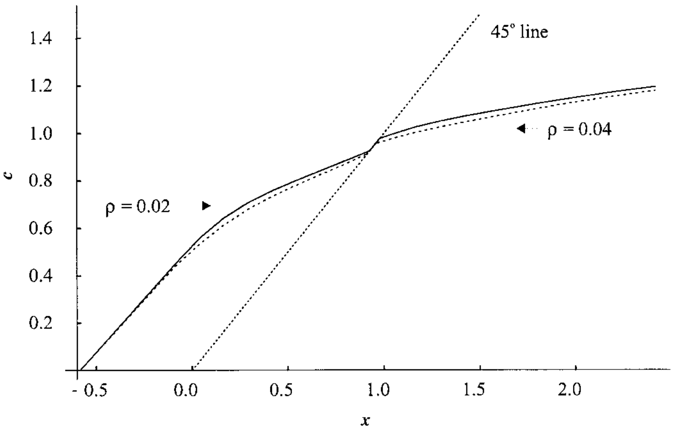

The solid line is the consumption policy rule for the baseline model when the probability of job loss is 2% per year. The dashed line is the policy rule when the probability of job loss is increased to 4% per year.

FIGURE 2.—STEADY-STATE DISTRIBUTION OF WEALTH-TO-PERMANENT-INCOME RATIO FOR $\rho = 0.02$ AND $\rho = 0.04$

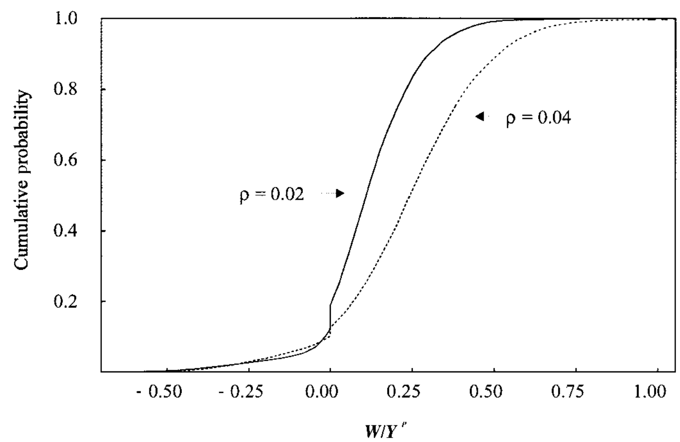

The solid line is the cumulative distribution function (CDF) for the ergodic distribution of the end-of-period wealth-to-permanent-income ratio for a population of consumers behaving according to the baseline model with a probability of job loss of 2% per year. The dashed line is the CDF for the ergodic distribution with a job loss probability of 4% per year.

line satisfy equation (7) and spend less than their current resources, leaving some wealth both to 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 gap 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, 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% annually, rather than our baseline assumption of 2%. With the exception of a portion where $c_t = x_t$ in both regimes, the new consumption function lies below the old one, indicating that at most levels of wealth a consumer facing a 4% unemployment risk would consume less than one facing a 2% risk.[8] For households consuming less than their current resources, the precautionary response shows up as increased saving, whereas for households consuming more than their current resources, the precautionary response shows up as a reduction in borrowing.

D. *The Distribution of Net Worth*

How do these precautionary responses affect the observable variable in our data, the cross-sectional distribution of net worth? With no income uncertainty, the gap between R^{borrow} and R^{lend} would cause households to hold wealth of exactly zero. With income uncertainty, on average, households carry some precautionary assets. The solid line in figure 2 and the first column of the top panel of table 1 show the steady-state distribution of the ratio of net worth to permanent income under $\rho = 0.02$.[9] At the mean of the distribution, W_t/Y_t^P is 0.115—the level of wealth is roughly $1\frac{1}{3}$ months' (0.115 × 12) worth of permanent income. Because wealth would be zero if income were certain, this can be interpreted as the average level of precautionary balances associated with the income uncertainty in our

TABLE 1.—STEADY-STATE DISTRIBUTION OF WEALTH-TO-PERMANENT-INCOME RATIOS

	$V_{\min} = 0.20$		$V_{\min} = 0.40$	
	Unemployment Rate:		Unemployment Rate:	
	2%	4%	2%	4%
Mean W_t/Y_t^P of:				
All households	0.115	0.237	0.081	0.108
W_t/Y_t^P percentile:				
20th to 30th	0.026	0.109	0.003	0.009
40th to 50th	0.093	0.218	0.056	0.085
60th to 70th	0.160	0.318	0.120	0.165
80th to 90th	0.258	0.452	0.209	0.274
Fraction of households with:				
$W_t/Y_t^P < 0$	0.116	0.103	0.162	0.167
$W_t/Y_t^P = 0$	0.068	0.024	0.097	0.069
$W_t/Y_t^P > 0$	0.817	0.874	0.741	0.764

Change in mean W_t/Y_t^P with unemp. rise from 2% to 4%:	$V_{\min} = 0.20$		$V_{\min} = 0.40$	
	Absolute	Percent	Absolute	Percent
All households	0.122	106	0.027	33
W_t/Y_t^P percentile:				
20th to 30th	0.082	311	0.007	256
40th to 50th	0.125	135	0.028	51
60th to 70th	0.158	99	0.045	37
80th to 90th	0.195	76	0.065	31

[8] The entire shift in the consumption function reflects precautionary behavior, because we adjust the mean of the V_t so that the change in ρ does not affect the expected level of income.

[9] We simulate the model for a large number of households for 10 periods—sufficient time for the W/Y^p distribution to stabilize. We refer to this stable or ergodic distribution as the steady-state distribution.

baseline parameterization. Precautionary balances rise to about 3 months' worth of income for households in the 80th to 90th percentile of the wealth-to-income distribution. Still, many households hold very little wealth: W_t/Y_t^P is just 0.026, or about $1\frac{1}{2}$ weeks' income, for households in the 20th to 30th percentile. Indeed, as shown in the middle panel of table 1, a noticeable proportion of households hold $W_t \leq 0$ as a result of recent negative shocks to income: Given the support to consumption provided by social insurance, even with a 15% borrowing rate, optimal behavior implies that about 12% of households are in debt; and the difference between borrowing and lending rates is large enough that almost 7% of households have wealth of exactly zero.

The dashed line in figure 2 and the second column of table 1 show the steady-state distribution of wealth when the probability of job loss is 4%. Reflecting the precautionary response, this distribution lies to the right of the base-case distribution. As shown in the first column in the bottom panel, the increase in the average wealth ratio is 0.122—about $1\frac{1}{2}$ months' worth of permanent income; this additional precautionary saving represents a 106% increase in the average ratio of net worth to permanent income (column 2). As can be seen reading down the bottom panel, the absolute change in W_t/Y_t^P increases with the ratio of wealth to income. But, because the bottom part of the distribution is at such a low level of net worth, even a small absolute shift represents a much larger percentage increase in W_t/Y_t^P at the lower percentiles than it does at higher levels of wealth.

Our specification of $V_{\min}$ is an extremely simple form of social insurance that lacks real-world features such as means testing or other forms of progressivity. Such factors might generate a systematic relationship between precautionary behavior, permanent income, and wealth that could be important for empirical work. However, incorporating a more realistic specification that, for example, makes $V_{\min}$ a decreasing function of wealth and income would greatly increase the computational complexity of the model. Still, we can provide some crude idea of the effects of such features by comparing the steady-state distributions of wealth from simulations that use different values of $V_{\min}$. As shown in the right-hand columns of table 1, raising $V_{\min}$ to 40% (increasing the generosity of social insurance) lowers wealth holdings throughout the W_t/Y_t^P distribution. Furthermore, as can be seen in the lower panel, the precautionary response of W_t/Y_t^P to an increase in unemployment risk is substantially smaller (in both absolute and percentage terms) than under our base case $V_{\min}$. These results suggest that making $V_{\min}$ a decreasing function of W_t and Y_t^p would tend to induce positive correlations between permanent income and both the level of precautionary wealth and the precautionary response to a change in unemployment risk. Indeed, such effects are generated by the more elaborate social insurance programs embedded in the model of Hubbard, Skinner, and Zeldes (1995). In this model, social insurance tends to be available only once a household has nearly exhausted its own resources. As a result, the marginal utility of saving is quite small for very low-income households, so they do not raise saving much in response to an increase in uncertainty.

TABLE 2.—WEALTH-TO-PERMANENT-INCOME RATIOS FOR HOUSEHOLDS WITH DIFFERENT EMPLOYMENT HISTORIES

	Not Unemp. in Last 3 Years	Unemployed in Year t	Unemployed in Year $t-1$	Unemployed in Year $t-2$
Fraction of households with:				
$W_t/Y_t^P < 0$	.063	.995	.975	.810
$W_t/Y_t^P = 0$	.074	.000	.015	.055
$W_t/Y_t^P > 0$	.863	.005	.010	.135
Mean W_t/Y_t^P	.136	−.390	−.210	−.094

All calculations assume an unemployment rate of 2%.

Finally, our model implies that shocks to income have persistent effects on household balance sheets that could distort comparisons of precautionary wealth holdings based on the current status of households' unemployment risk. As indicated by the consumption functions in figure 1, at any given level of cash on hand, an increase in unemployment risk will raise precautionary saving (with the exception of some households where $c_t = x_t$). This raises the aggregate wealth of households whose members remain employed because, on average, their income does not fall. However, as seen in table 2, virtually all households whose members 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 39% of their permanent income. Furthermore, job loss leaves lasting scars: Two years after the fact, households that experienced a spell of unemployment (column 4) hold substantially less wealth and are significantly more likely to have negative net worth than their counterparts who were continuously employed over the past three years (column 1)—even though both groups face the same current risk of becoming unemployed. These results point out a complication for empirical work: Outside the steady state, the effect of higher unemployment risk on aggregate wealth may be ambiguous because the balance between the decline in wealth for the unemployed and the rise in wealth for the employed is unclear.

In sum, the model has several important implications for empirical work. A substantial fraction of optimizing households may have low or negative net worth. For households with similar recent employment experience, those facing a greater probability of job loss should hold more wealth. However, a household may hold less net worth than a household with the same or less unemployment risk if the first household has recently experienced a spell of unemployment. Finally, precautionary responses may vary with wealth-to-permanent-income ratio and with household income. Of course, the model is only meant to highlight certain considerations that are relevant for testing the relationship between wealth and unemployment risk. Other important empirical issues such as controlling for other motives for saving and for differences in preferences are discussed in sections III and IV.

III. Data and Econometric Methodology

Our empirical goal is to estimate the relationship between W/Y^p and unemployment risk, taking into account the salient features of the wealth distributions discussed in section II. We use a two-step procedure. The first step constructs estimates of unemployment risk and permanent income for each household. The second step estimates the relationship between W/Y^p and unemployment risk, controlling for other factors that may influence wealth accumulation.

A. Data Sources

Like Starr-McCluer, Engen and Gruber, and Carroll and Samwick, we focus on the cross-sectional relationship between household wealth and uncertainty. We use wealth data from the Survey of Consumer Finances, which has the best available information on households' balance sheets.[10] However, the SCF has at most only about 4000 households in each wave—too few to accurately estimate the probability of job loss for the regional breakdown that we use. Consequently, we estimate unemployment risk with data from the much larger outgoing rotation groups of the Current Population Survey, using the quasi-panel structure to obtain records of individuals' employment status taken one year apart. Our analysis uses the 1983, 1989, and 1992 waves of the SCF and corresponding CPS samples. We do not use data from more recent years because data before and after the redesign of the CPS in 1994 are not completely comparable.

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 focuses on households' total net worth.

The simulations in section IID indicate that, on average, households that have recently experienced a spell of unemployment or other major earnings disruption will have lower net worth than other households. Thus, because bad draws to income are more likely among consumers with high current unemployment risk, we could see a negative correlation between wealth and the ex ante probability of unemployment even in the presence of strong precautionary saving motives. Ideally, we would like to include explanatory variables to control for earlier employment and income disruptions. Unfortunately, the SCF does not record income or unemployment history or reasons why the respondent may have switched jobs; the only question regarding employment history simply asks "How many years have you been employed at your current job?" Given the absence of direct controls, we implemented the simple indirect control suggested by our model. That is, we excluded those who may have recently experienced unemployment from the analysis by limiting our sample to households whose head had been working at the same job for at least three years. As we discussed earlier, our model unambiguously predicts a positive relationship between current unemployment risk and wealth for such continuously employed households.

Other more standard sample restrictions and other aspects of the data are discussed in the appendix. Altogether, the SCF sample contains 1,689 households from the 1983 wave, 1,025 households from the 1989 wave, and 1,032 households from 1992 wave. The CPS samples corresponding to the three SCF waves contain 59,252 households, 60,026 households, and 63,351 households, respectively.

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, 95th 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

Sample includes all households in the SCF area probability sample 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.

B. Wealth Transformation

Table 3 shows summary statistics for net worth, assets, and liabilities. The distribution of wealth is highly skewed at the top end, with the net worth of the median household only about half the size of the average in each year. At the bottom end of the distribution, between 4% and 6% of households have zero or negative net worth, and between 8% and 10% have net worth less than one month's income.

Clearly, some transformation of wealth is necessary in order to avoid undue influence from the very small number of extremely wealthy households. The distribution is sufficiently skewed that the potential solution provided by the

[10] Curtin, Juster, and Morgan (1989) compare the SCF wealth data with those 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).

FIGURE 3.—DISTRIBUTIONS OF RESIDUALS AND DATA TRANSFORMATIONS

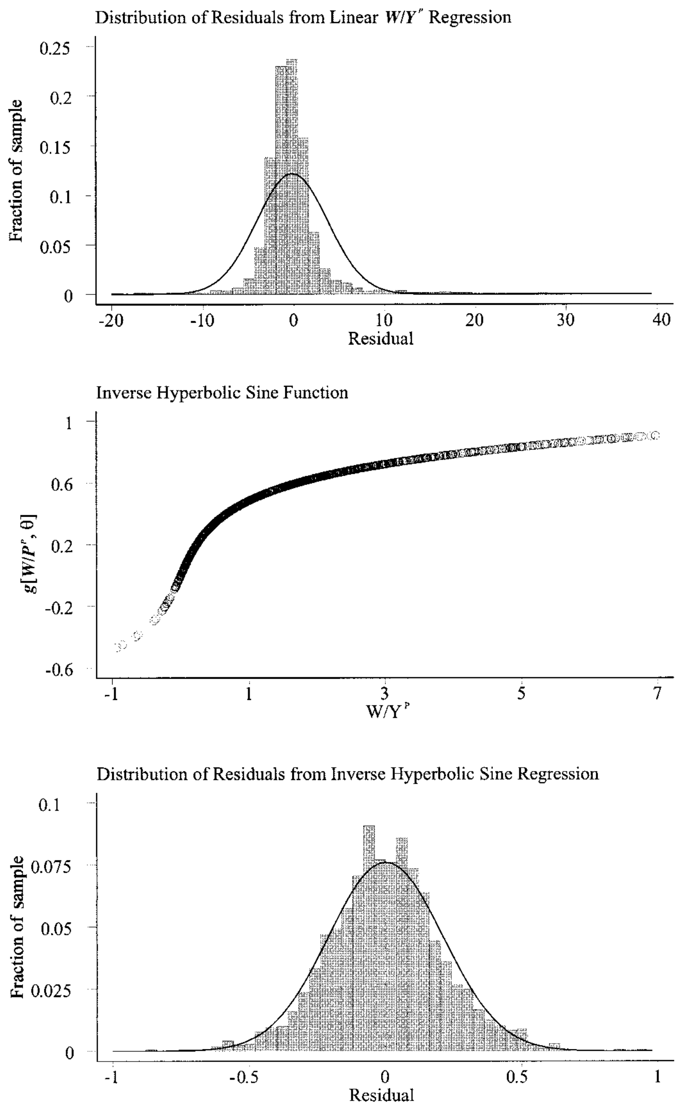

model in section II—using the ratio of net worth to some proxy for permanent income as the dependent variable—does not work. The top panel of figure 3 plots a histogram of the residuals from a linear regression of a W/Y^p proxy on the explanatory variables used in the model estimated in table 6 for 1983 (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.[11]

Most previous studies have transformed wealth by taking logarithms. They also discard households with wealth below some threshold (at least zero) or use a different transformation for those with $W \leq 0$.[12] However, discarding such households would reduce the size of our samples considerably. It also would limit our ability to explore the potentially important differences in precautionary behavior that may exist across the wealth and income distributions. The estimation technique also would need to allow for truncation or alternative data transformations. Furthermore, the log transformation assumes a constant elasticity of W/Y^p with respect to changes in explanatory variables. In contrast, our simulations imply that at the low end of the distribution, small increases in dollar terms in precautionary wealth can represent extremely large increases in percentage terms. This suggests that average precautionary effects estimated using logs could give undue weight to responses at the lower end of the wealth distribution.

As an alternative, we transform net worth with the inverse hyperbolic sine function, suggested by Burbidge, Magee, and Robb (1988). The inverse hyperbolic sine of z is

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

where θ is an estimated damping parameter. Like the logarithm, $g[z, \theta]$ downweights large values of z. However, $g[z, \theta]$ has several advantages over the use of logs: it admits 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 $\partial z/\partial x = (\partial g[z, \theta]/\partial x)(\partial z/\partial g[z, \theta])$, where the first factor on the right-hand side equals the regression coefficient on x—call it β_x—and the second equals $(\theta^2 z^2 + 1)^{1/2}$. Thus, $(\partial z/\partial x)(x/z) = \beta_x x(\theta^2 + 1/z^2)^{1/2}$, so that elasticities are decreasing functions of z.

The middle panel of figure 3 shows $g[z, \theta]$ with z equal to our W/Y^P proxy and $\theta = 3.87$ (the value we estimate for the 1983 total sample in table 6). The bottom panel shows the residuals from a regression of $g[W/Y^p, \theta]$ on the 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.

C. The Likelihood Function

According to our theoretical model, a household's wealth holdings are a function of its job-loss risk, its permanent income (through the interaction with $V_{\min}$), the income-generating process, and factors determining rates of time preference and risk aversion. A more complete model would

[11] For 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.

[12] Some examples of transformations and restrictions found in the literature are Diamond and Hausman (1984): ln W using only $W >$ \$4000; King and Dicks-Mireaux (1982): $\ln(W/Y^p)$ using only $W >$ \$2500; Starr-McCluer (1996): ln W if $W > 0$ and 0 if $W \leq 0$; Carroll and Samwick (1997, 1998): $\ln[W - \min(W, 0) + 1]$.

also include life cycle factors. In terms of observable proxies, the simplest empirical model that we will estimate is

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

where j indexes households, $g[\ ]$ is the inverse hyperbolic sine function, 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 described below that are meant to capture other determinants of wealth.

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

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

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

D. First-Stage Regressions for Unemployment Risk

For currently employed individual j, we assume there exists a latent variable $u_j^* = Z_j^u\alpha_u + \nu_j$ 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. ν_j is a logistically distributed idiosyncratic shock that is uncorrelated with Z_j^u, a row vector of observable characteristics for individual j at time t. Thus, $\Pr(u_j|e_j)$, the probability of a currently employed person becoming unemployed, is

$$\Pr(u_j|e_j) = \frac{\exp(Z_j^u\alpha_u)}{1 - \exp(Z_j^u\alpha_u)}. \tag{12}$$

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 (based on the first of the two readings from the outgoing rotation group of the CPS) and unemployed in month $t + 12$ (based on the second CPS reading), and takes on a value of 0 if individual j is employed in both periods.

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. This $\Pr(\hat{u}_j)$ is then used in our SCF wealth regressions. Parameter consistency of such two-sample estimators requires that the samples be randomly drawn from the same population (see Angrist and Krueger, 1992). Because both the CPS and the SCF area-probability samples are random draws from the noninstitutional U.S. population, this procedure produces consistent estimates of the probability that a household in the SCF area sample will become unemployed.[13]

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

Note that we are exploiting the quasi-panel structure of the CPS to estimate the conditional probability of being unemployed in the future rather than simply the unconditional probability of currently being unemployed. $\Pr(u_j)$ can be thought of as a rational expectation of the odds of a currently employed household being unemployed one year from now conditional on Z_j^u and is therefore an unbiased estimate of this measure of unemployment risk. This risk is particularly relevant for the households in our SCF samples, because they all are currently employed and have been so for at least three years. Still, our proxy has clear shortcomings. Importantly, because the CPS outgoing rotation groups only provide snapshots of individuals at one-year intervals, we cannot observe spells of unemployment that are started and completed between interviews and we cannot calculate the length of an observed spell of unemployment.

E. 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}:

$$\ln Y_j^p = Z_j^{yS}\alpha_y. \tag{13}$$

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 the number of children in the household, the number of earners, the log of any retirement income, and dummies to indicate home ownership, retirement status, whether the head or spouse has a defined benefit pension, and 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, begin-

[13] We did not use data from the SCF's nonrandom "list" samples; these were designed to overrepresent wealthy households and thus cover a very different population than the area-probability samples.

ning with Friedman. Note, though, that such measures do not capture any unobservable individual-specific components of permanent income.

F. *Control Variables*

The control variables C_j are meant to capture factors that may affect wealth through some channel other than changing unemployment risk or permanent income, such as life cycle considerations or rates of time preference and risk aversion.

To identify the effect of uncertainty on wealth, some instrument for the uncertainty proxy must be excluded from C_j. Many studies have excluded a variable such as occupation or industry in specifications similar to equation (10). However, because such variables can be related to the expected life cycle profile of income or to discount rates or 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 the Census subregion in which the household resides, which we believe is likely to be uncorrelated with income profiles or preference parameters related to saving. Furthermore, Blanchard and Katz's (1992) finding of no permanent differences in unemployment rates across regions suggests that most households do not choose 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 $\Pr(u_j)$ on wealth than a measure based solely on the other variables in Z_j^u.[14]

G. *Other Estimation Issues*

As a result of the assumptions on Z_j^u, C_j, and Z_j^{yS}, the coefficient β_y is identified from region, whereas 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 parameters 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 $\Pr(u_j)$, 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 $\Pr(u_j)$.

[14] Relative regional unemployment risk does vary over time. In simple top-to-bottom orderings, the average change in the rankings of the nine Census subregions' unemployment rates between adjacent survey years is about 1¾ positions. If regional effects were fixed, there would be no change; if positions were completely random, the average change in rankings would be 4½ positions. Thus, using three surveys covering 9 years should reduce the odds that we are just capturing regional fixed effects on wealth.

TABLE 4.—FIRST-STAGE ESTIMATION FOR UNEMPLOYMENT RISK (CPS) AND GROUP INCOME (SCF)

	CPS $\Pr(\hat{u})$			SCF[a] $\ln \hat{Y}^P$		
	1983	1989	1992	1983	1989	1992
F-tests (*p*-values):						
Occupation	0.007	0.134	0.008	0.684	0.361	0.847
Industry	0.000	0.000	0.000	0.009	0.111	0.084
Region	0.012	0.000	0.000	0.000	0.000	0.115
Education	0.000	0.000	0.001	0.417	0.742	0.016
White	0.003	0.001	0.007	0.667	0.061	0.463
Female head	0.051	0.105	0.011	0.000	0.002	0.000
Age	0.168	0.798	0.980	0.000	0.000	0.016
Age × occupation	0.011	0.107	0.097	0.008	0.059	0.759
Age × education	0.069	0.042	0.061	0.376	0.553	0.248
Marital status	0.000	0.000	0.000	0.000	0.008	0.058
Head	0.025	0.917	0.000			
Head × age	0.237	0.774	0.001			
Homeowner				0.000	0.000	0.000
No. of children				0.055	0.593	0.788
No. of earners				0.000	0.000	0.000
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
Retired				0.161	0.274	0.001
log(retirement income)				0.000	0.129	0.071
Mean $\Pr(\hat{u})$	0.027	0.019	0.022			
Std. dev. $\Pr(\hat{u})$	0.021	0.015	0.017			
Adjusted R^2				0.46	0.41	0.28
No. of households	59,252	60,026	63,351	1688	1025	1032

[a] Sample includes the self-employed.

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

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

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

Clearly, using first-stage equations to estimate variables in equation (14) introduces complications. A discussion of parameter consistency and adjustments we make to summary statistics to (asymptotically) account for the first-stage estimation can be found in an earlier version of this paper.[15]

IV. Empirical Results

A. *First-Stage Results*

The first three columns of table 4 present results from the first-stage CPS logit equations for $\Pr(\hat{u}_j)$. The *p*-values for

[15] See Finance and Economics Discussion Series working paper 1999-15, available on the Federal Reserve Board's Web site. The most important adjustments blow up the $\sum \hat{e}_j^2/n_s$ to consistent estimates of $\sigma_e^2 = \text{plim} \sum e_j^2/n_s$. In the income-interaction model described below, this boosts σ_e by about 25% in the 1983 and the pooled sample, and by between 50% and 70% in the 1989 and 1992 runs.

TABLE 5.—SCF SECOND-STAGE ESTIMATION RESULTS: $g[W/\hat{Y}^P, \theta] = \beta_0 + \beta_u \Pr(\hat{u}) + \beta_y \ln \hat{Y}^P + C\beta_c + \epsilon$

	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	0.23	−0.67	−1.03	0.03	−0.60
	(0.50)	(0.22)	(−0.28)	(−0.33)	(0.01)	(−0.23)
$\ln \hat{Y}^P$	0.09	0.12	0.23	0.44	0.40	0.26
	(1.24)	(1.62)	(1.89)	(2.44)	(1.55)	(1.58)
Number of earners	−0.04	−0.05	−0.08	−0.11	−0.18	−0.11
	(−2.24)	(−2.58)	(−2.05)	(−1.94)	(−2.19)	(−1.99)
Owns home	0.28	0.27	0.32	0.32	0.28	0.27
	(9.00)	(8.61)	(5.80)	(4.89)	(3.46)	(4.55)
Defined benefit pension	−0.06	−0.04	−0.10	−0.08	−0.08	−0.02
	(−3.43)	(−2.02)	(−2.60)	(−1.74)	(−1.39)	(−0.49)
Turned down for credit	−0.05	−0.07	−0.03	−0.10	−0.06	−0.07
	(−2.78)	(−3.15)	(−0.74)	(−2.14)	(−1.18)	(−1.72)
Θ	4.56	4.56	3.86	3.30	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	0.01	−0.05	−0.06	0.00	−0.02
	(0.50)	(0.22)	(−0.28)	(−0.32)	(0.01)	(−0.23)
Number of households	1689	1441	1025	835	1032	880

Asymptotic t-statistics shown in parentheses.

F-tests of the joint significance of industry, region, education, and race dummies (as groups) are all highly statistically significant. 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, and age interacted with education is significant at the 7% level or better in all three years. The lower part of the table provides summary statistics for the predicted probability of becoming unemployed. In order to obtain precise estimates of β_u, we need the variation in the $\Pr(\hat{u}_j)$ to be large; fortunately, it is, with their standard deviations running between 70% and 80% of the means.

The second three columns of table 4 present results from the first-stage income regressions. The Z_j^{yS} explain close to half of the variation in the log of reported income in 1983, and between 30% and 40% in 1989 and 1992. Most of the explanatory variables are statistically significant. A notable exception is education in 1983 and 1989; but a simple regression of income on education dummies alone produces significant coefficients, so that the insignificance in the overall equation likely results from collinearity and not the irrelevance of education to permanent income.

B. Second-Stage Results: The Basic Model

The results for our second-stage equation are found in table 5; we show estimated coefficients and (asymptotic) t-statistics only for selected variables. The first column for each year presents results from the total area samples, and the second shows estimates for samples that exclude households whose head is self-employed.[16] We omit these households because their balance sheets can be heavily influenced by business holdings and because they may have different attitudes toward risk.

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

$$\frac{\partial W_j/\hat{Y}_j^P}{\partial \Pr(\hat{u}_j)} = \left[\left(\frac{\hat{\theta} W_j}{\hat{Y}_j^P}\right)^2 + 1\right]^{1/2} b_u, \tag{15}$$

where b_u is the estimated coefficient on $\Pr(\hat{u}_j)$. The table presents $\partial(W_j/\hat{Y}_j^P)/\partial \Pr(\hat{u}_j)$ and asymptotic t-statistics that correspond to a 1 percentage point increase in $\Pr(\hat{u}_j)$ at the median W_j/Y_j^P. [Because $\Pr(\hat{u}_j)$ averages between 2% and $2\frac{3}{4}$% percent with a standard deviation of $1\frac{1}{2}$ to 2 percentage points, a 1 percentage point increase in $\Pr(\hat{u}_j)$ is a fairly large change in our metric of employment risk.] In the 1983 sample including the self-employed, a 1 point increase in $\Pr(u_j)$ increases W_j/Y_j^P by 0.04, or about $\frac{1}{2}$ 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

[16] The runs that exclude self-employed households also exclude them in the first-stage regressions.

TABLE 6.—SCF SECOND-STAGE ESTIMATION RESULTS: $g[W/\hat{Y}^P, \theta] = \beta_0 + \beta_u \Pr(\hat{u}) + \beta_{uy} \Pr(\hat{u}) \ln \hat{Y}^P + \beta_y \ln \hat{Y}^P + C\beta_c + \epsilon$

	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	−26.75	−57.22	−66.09	−67.77	−50.46
	(−2.53)	(−2.24)	(−1.99)	(−1.92)	(−1.99)	(−1.69)
$\Pr(\hat{u}) \ln \hat{Y}^P$	2.96	2.76	5.72	6.47	6.85	5.05
	(2.58)	(2.27)	(1.98)	(1.90)	(1.99)	(1.68)
$\ln \hat{Y}^P$	−0.03	0.01	0.09	0.25	0.17	0.11
	(−0.37)	(0.12)	(0.62)	(1.17)	(0.66)	(0.57)
$\frac{\partial W_j/\hat{Y}_j^P}{\partial \Pr(\hat{u}_j)} \times \hat{Y}^P$ percentile:						
10th	0.03	0.03	−0.01	−0.04	−0.01	−0.02
	(0.95)	(0.51)	(−0.18)	(−0.60)	(−0.15)	(−0.37)
30th	0.11	0.04	0.11	0.06	0.16	0.03
	(1.84)	(1.34)	(0.74)	(0.30)	(0.75)	(0.38)
50th	0.29	0.24	0.27	0.15	0.17	0.10
	(2.14)	(1.67)	(1.18)	(0.74)	(1.17)	(0.79)
70th	0.39	0.33	0.44	0.24	0.39	0.14
	(2.34)	(1.87)	(1.40)	(1.03)	(1.46)	(1.05)
90th	0.52	0.42	0.71	0.55	0.77	0.41
	(2.47)	(2.01)	(1.61)	(1.25)	(1.67)	(1.26)
Overidentify restr. p-value	0.50	0.71	0.32	0.61	0.89	0.46

Asymptotic t-statistics shown in parentheses.

not statistically different from zero. The samples excluding the self-employed produce similar small and statistically insignificant estimates of b_u and $\partial(W_j/Y_j^p)/\partial \Pr(u_j)$.

The signs of the coefficients on the other independent variables generally make sense and are reasonably stable across SCF years and samples. With regard to those possibly related to precautionary saving, we find W_j/Y_j^p is decreasing in the number of earners per household, consistent with the notion that two (or more) earners lower precautionary reserves because both earners are unlikely to become unemployed at the same time.[17] The coefficient on ln Y^p is positive for all years and samples and is of at least marginal statistical significance in all but the 1983 sample including the self-employed. These results are consistent with the possibility that higher social insurance replacement rates reduce precautionary wealth held by lower-income households. With regard to factors related to saving more generally, homeowners have higher net worth than nonhomeowners relative to income, and having a defined benefit pension plan has a negative effect on net worth, as predicted by a standard forward-looking model of consumer behavior.

Identification of β_y and, as a practical matter, of β_u requires that the error term be orthogonal to the excluded instruments, the regional dummies. In econometric terms, the key assumption is that, conditioned on the C_j, region is correlated with the dependent variable only via its correlation with $\Pr(\hat{u}_j)$ and $\ln \hat{Y}_j^P$. We test this assumption formally with overidentifying restrictions (OID) tests; the p-values shown in table 5 are not close to statistical significance, suggesting that region is a valid excluded instrument. Of course, this finding is subject to the usual caveat that these tests may have low power.

C. Second-Stage Results: An Extended Model

A possible explanation for the weak results in table 5 is that the empirical model is too stylized to fully characterize actual household behavior. As discussed earlier, means testing and progressivity in social insurance programs imply that, all else equal, higher-income households might exhibit a larger precautionary response to a change in unemployment risk. Furthermore, a forward-looking optimizing framework may not apply to all households. Instead, some households may use a rule of thumb to determine consumption (as in Campbell and Mankiw, 1989), simply setting current consumption equal to some fraction of current income and not reacting at all to changes in income risk. Such behavior could contribute to a positive relationship between permanent income and the observed precautionary response of wealth to risk if high permanent income households were less likely to be rule-of-thumb consumers.

These considerations suggest that a second-stage equation that allows the precautionary response to vary with permanent income—say by including a term $\Pr(\hat{u}_j) \times \ln \hat{Y}_j^p$ as an independent variable—might give a more accurate representation of $\partial(W_j/Y_j^p)/\partial \Pr(u_j)$. Table 6 presents this model. The estimates for the control variables in this specification are not presented, as they are similar to those in table 5.

[17] Guiso and Jappelli (1994) found that households saved less if both spouses worked. Summers and Carroll (1987) suggested that that the rise in two-earner families may help to explain the secular decline in saving in the United States. See Browning and Lusardi (1996) for more discussion of this point.

Allowing for the interaction between unemployment risk and income suggests a much different 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% level for each wave's full sample and statistically significant at the 10% level or better in the samples excluding the self-employed. Also, the coefficients and t-statistics on $\ln\ \hat{Y}_j^p$ are much smaller than those in the base model.

Letting b_{uy} be the estimated coefficient on $\Pr(\hat{u}_j)\ \ln\ \hat{Y}_j^P$. Then in this extended model

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

The bottom portion of table 6 presents $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ and (asymptotic) t-statistics corresponding to a 1 percentage point increase in $\Pr(\hat{u}_j)$ for the percentile of $\hat{Y}_j^p$ listed in the first column.[18] For the three 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. For example, in 1983, 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 slightly over a week's worth of income, and has a t-statistic of less than 1. 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 $\hat{Y}_j^p$ rises. In the 1983 full sample, at the median of the distribution a household responds to a 1 percentage point increase in the probability of becoming unemployed by increasing $W_j/\hat{Y}_j^p$ by 0.29, or about $3\frac{1}{2}$ months' income. These effects are large enough to be economically important, but not so large as to be intuitively implausible. For the 50th percentile of $\hat{Y}_j^p$ and above, these 1983 estimates of $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ are statistically different from zero at the 5% level. The $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ for the 1989 and 1992 SCFs are of very similar magnitude to those in 1983, but are estimated much less precisely. In part, this might reflect the fact that the samples in those years are about 40% 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 self-employed households. This is consistent with the idea that self-employed households may be more vulnerable to income losses when business conditions go bad and the fact that they are less likely to be covered by unemployment insurance. The precautionary effects are somewhat less precisely estimated than in the total sample; again, this is in part a result of the smaller sample size.

[18] Strictly speaking, we divided the data into 50 equal-size 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. We did so to avoid undue influence from outlier values for W_j.

D. Second-Stage Results: Pooled SCF Samples

Our sample restrictions left us with between 1,025 and 1,689 households in the individual SCF samples. These samples may be too small to produce precise estimates. Yet the similarity across years in the sizes of the precautionary effects and the other coefficients suggests that we can gain precision by pooling the data from the three different SCF years in the second-stage wealth regressions.

Table 7 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). 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 each year to allow the relationships between these variables 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 other nation-wide factors.

The model with no interaction between $\Pr(\hat{u}_j)$ and $\hat{Y}_j^P$ now shows a small precautionary effect—for the sample that includes the self-employed, increasing job-loss risk by 1 percentage point raises the wealth-to-income ratio by 0.06 (0.7 months of income) at the median income. However, the estimate is statistically significant at just the 20% level. In the model with the $\Pr(\hat{u}_j)$-$\hat{Y}_j^P$ interaction, we still observe only small and statistically insignificant effects of job-loss risk on wealth at low levels of $\hat{Y}_j^p$, but find statistically and economically significant influences at the 30th percentile of permanent income and higher. At the median income, a 1 percentage point increase in $\Pr(\hat{u}_j)$ is associated with an increase in precautionary balances of 0.17 times annual income (2 months), and the t-statistic on this effect is over 3. The estimated effects for the sample excluding the self-employed are a bit smaller, but with the larger sample size, they still are highly statistically significant for the upper half of the income distribution.[19]

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

$$C_t = \frac{r}{1+r} X_t + \frac{r}{r-g} Y_t. \tag{17}$$

[19] Tests of the pooling restrictions are ambiguous. In ones ignoring the fact that $\Pr(u_j)$ and $\ln\ Y_j^p$ are estimated, an F-test easily fails to reject that the β's are constant across years while a likelihood ratio test (which includes restrictions on σ_e^2 and θ) easily rejects the pooling restrictions. However, both tests degenerate when adjusted for first-stage estimation; this likely reflects approximation error in our estimates of σ_e^2.

TABLE 7.—SCF SECOND-STAGE ESTIMATION RESULTS: 1983, 1989, 1992 POOLED DATA

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

Asymptotic t-statistics shown in parentheses.

Let $g = 0.02$, $r = 0.05$ $[r/(r - g) = 1.67]$, and suppose there is a 1% increase in the odds of a severe episode (see Carrington, 1993) in which the breadwinner loses their job for one-half year and then finds a new job that permanently pays 10% less. From equation (17), in a certainty equivalence world, C_t would fall and W_t would rise by $0.01 \times (0.5 + 0.1) \times 1.67 = 0.01$ of income, well below our estimate of 0.17 at the median $\hat{Y}_j$. In other words, the increase in job-loss probability would have to be maintained for 17 years for our estimates to be in line with this pessimistic scenario from a model without precautionary saving.

Our findings appear to tell roughly the same story as Carroll and Samwick (1997) and Engen and Gruber (2001), but probably indicate larger precautionary effects than in Lusardi (1997). Our median household increases wealth by 17% of $\hat{Y}_j^p$ in response to a 1 percentage point—or roughly a $\frac{1}{2}$ standard deviation—increase in $\Pr(\hat{u}_j)$. Carroll and Samwick (1997) estimate a 4% increase in wealth (relative 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%, they are probably considering a smaller increase in uncertainty than we are. Engen and Gruber estimate that a 10% increase in the income replaced by unemployment insurance—about a $\frac{1}{2}$ standard deviation change in the UI replacement rate (see Gruber, 1997)—lowers the ratio of net financial assets to income by 16%. 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 $1\frac{1}{2}$% to 3% increase in $W_j/\hat{Y}_j^p$.

E. *Identification Using Region*

We noted earlier that region may be a better instrument to exclude from C_j in order to identify uncertainty than variables most other papers have used, such as occupation, education, or industry, which a priori are more likely to depend on unobserved taste parameters that are also correlated with saving. This section considers the sensitivity of our results to this assumption. Alternative estimates of $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ are shown in the top panel of table 8, based on modifications to the pooled SCF interaction model (for the sample including the self-employed). For reference, we repeat the base model results in column 1.

Alternative Excluded Instruments: Columns 2, 3, and 4 present results using occupation, education, or industry as the excluded instrument.[20] When excluding occupation or education, $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ is reduced substantially.[21] The estimates when excluding industry are essentially the same as in the base model, suggesting that the effect of job-loss

[20] The variable in question is the only excluded instrument (we include region in C_j).

[21] Correlation between the excluded instrument and risk aversion probably would bias down estimated precautionary effects (of course, the bias cannot be signed directly in a multivariate model). For example, suppose 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 would make a negative contribution to the correlation between wealth and $\Pr(u)$.

TABLE 8.—IDENTIFICATION ISSUES AND SENSITIVITY CHECKS

	Base Model	Alternative Excluded Instruments: Occupation	Alternative Excluded Instruments: Education	Alternative Excluded Instruments: Industry	No Excluded Instrument
$\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ by $\hat{Y}^P$ percentile:					
10th	0.01	−0.02	−0.03	0.01	−0.00
	(0.64)	(−1.34)	(−2.69)	(1.01)	(−0.40)
30th	0.06	−0.00	−0.03	0.06	0.03
	(2.39)	(−0.11)	(−1.02)	(3.52)	(1.01)
50th	0.17	0.03	−0.00	0.17	0.09
	(3.07)	(0.60)	(−0.06)	(4.38)	(1.64)
70th	0.30	0.09	0.05	0.29	0.17
	(3.43)	(1.09)	(0.59)	(4.65)	(2.05)
90th	0.43	0.15	0.13	0.41	0.26
	(3.62)	(1.45)	(1.09)	(4.66)	(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.

	Base Model	Reduced Instruments for ln $\hat{Y}^P$	Alternative Pr($\hat{u}$): a1	Alternative Pr($\hat{u}$): a2	Including Unemployed and Job Changers
$\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ by $\hat{Y}^P$ percentile:					
10th	0.01	0.02	−0.01	−0.00	0.01
	(0.64)	(0.97)	(−1.24)	(−0.70)	(1.52)
30th	0.06	0.09	−0.01	0.00	0.03
	(2.39)	(2.39)	(−0.61)	(0.04)	(2.78)
50th	0.17	0.17	−0.00	0.01	0.11
	(3.07)	(2.86)	(−0.05)	(0.49)	(3.19)
70th	0.30	0.30	0.02	0.02	0.14
	(3.43)	(3.05)	(0.39)	(0.79)	(3.31)
90th	0.43	0.44	0.04	0.04	0.18
	(3.62)	(3.15)	(0.72)	(0.99)	(3.33)
Overidentify restr. p-value	0.34	0.14	0.15	0.11	0.01
Exclusion p-value	0.01	0.00	0.00	0.00	0.00

Asymptotic t-statistics shown in parentheses. All estimates based on pooled 1983, 1989, and 1992 data.

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. If, ex ante, we believe that region is a preferable instrument, this suggests that we may also want to have more confidence in the findings of studies that use industry as an instrument to identify precautionary behavior.

The OID tests reject the exclusion restrictions on occupation and industry at (close to) the 5% level, but do not reject education. These rejections indicate that our specification is not one where OID tests are powerless. Indeed, the fact that region passes the test suggests that it is a valid instrument; and the rejection of industry indicates that despite similar point estimates, industry probably is a worse instrument than region. Still, the failure to reject education, which a priori seems likely to be endogenous with regard to unobserved traits determining saving, suggests that the OID tests may not have as much power as we might like.[22]

Including Region As a Control: If an excluded variable is really proxying for some nonprecautionary factors affecting wealth, then it should be included as a control variable. Because we estimate the first-stage equations independently for each SCF year, our instrument set in the pooled model effectively is 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 the top panel of table 8 presents p-values for Lagrange multiplier tests of the exclusion of the appropriate instrument from C_j. These are soundly rejected in all specifications, including with our preferred regional exclusion.

The last column presents estimates of the pooled interaction model after adding region to the C_j. Here, abstracting from functional form, all of the identification of the effect of unemployment risk comes from the changes across years in regional job-loss risk variation. The estimates of $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ in this specification are smaller than those in base our model. Nonetheless, by the 50th percentile of $\hat{Y}_j^p$, a 1 percentage point increase in $\Pr(\hat{u}_j)$ is still associated with a 0.09 increase in $W_j/\hat{Y}_j^p$, with the effect statistically significant at the 10% level. The effect rises in economic and statistical importance at higher levels of permanent income.

These results indicate that even if the fixed-region effects reflect nonprecautionary factors, significant precautionary motives exist, albeit somewhat weaker than those in the

[22] Alternatively, the problem may be that education is highly collinear with the included control variables, as suggested by the first-stage regressions for ln $\hat{Y}_j^p$.

base model. However, given that C_j includes income, industry, occupation, education, demographics, and other factors, we think it unlikely that region is picking up nonprecautionary effects. Furthermore, because of the short time span of our data, persistence in regional business conditions might mean that collinearity between fixed-region effects and $\Pr(\hat{u}_j)$ could be reducing the estimate of $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$. And to the extent that fixed-region effects are reducing $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$ because they are correlated with risk, they should be added back to any calculation of precautionary behavior.

F. Other Sensitivity Checks

The bottom panel of table 8 presents a number of other sensitivity checks on our results. All estimates are based on the model including the $\Pr(\hat{u}_j) \times \ln \hat{Y}_j^p$ interaction, the samples including the self-employed, and the data pooled across SCF years, and with region excluded from C_j.

Reduced Instrument Set: Potential omitted regressor biases might be affecting the extended model. In the model with no interaction, 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 $\Pr(u_j)$. But for this to be true for b_{uy} in the interaction model, we would need to add the hundreds of cross products of the $Z_j^{\bar{u}S}$ and Z_j^{yS} to C_j. This clearly is impractical.

As a parsimonious check, we excluded the $Z_j^{\bar{u}S}$ from the first-stage estimate of $\ln \hat{Y}_j^p$: Because the resulting $\Pr(\hat{u}_j) \times \ln \hat{Y}_j^p$ no longer includes any variation due to $Z_j^{\bar{u}S}$, b_{uy} should be less influenced by the (potential) appearance of the $Z_j^{\bar{u}S}$ cross products in ϵ_j. Furthermore, a number of the $Z_j^{\bar{u}S}$, such as the credit measures, are not usually used in permanent income proxies; this exercise thus also tests the sensitivity of our results against a more traditional set of instruments for $\ln \hat{Y}_j$. As can be seen in the second column, moving to this reduced instrument set has little qualitative effect on $\partial(W_j/\hat{Y}_j^p)/\partial\Pr(\hat{u}_j)$. This suggests that omitting the cross products of the $Z_j^{\bar{u}S}$ from C_j does not bias our results; we can guess that leaving out cross products of the $Z_j^{\bar{u}S}$ and the Z_j^{uS} probably does not greatly influence the results, either. And the results indicate that the inclusion of nonstandard instruments in Z_j^{ys} likely does not greatly affect our analysis.

Alternative Measures of Unemployment Risk: Our measure for $\Pr(u_j|e_j)$, $\Pr(\hat{u}_j)$, is calculated using the odds based on current information (the Z_j^u for period t) that an individual who is employed in period t will be unemployed in period $t + 12$. This measure might be inappropriate if households' actual expectations were formed in a less forward-looking manner. The most obvious alternatives consider the odds that people with similar characteristics are currently unemployed: A conditional alternative (a1) can be constructed by rerunning equation (12) with the dependent variable marking the period-t employment status of people who were employed in month $t - 12$ (using Z_j^u from $t - 12$); an unconditional alternative (a2) can be formed based on the period-t employment status of all individuals (using Z_j^u from t). For the SCF regressions, $\Pr^{a1}(\hat{u}_j)$ and $\Pr^{a2}(\hat{u}_j)$ are calculated using the first-stage coefficients, β_u^{a1} and β_u^{a2}.

The third and fourth columns of the bottom panel of table 8 show the resulting $\partial(W_j/\hat{Y}_j^p)/\partial\Pr^a(\hat{u}_j)$; these are minuscule compared to those estimated using our preferred $\Pr(\hat{u}_j)$ and are statistically insignificant. In retrospect, however, these results are not surprising; they probably reflect biases in $\Pr^{a1}(\hat{u}_j)$ and $\Pr^{a2}(\hat{u}_j)$ when applied to the SCF. All household heads in our SCF samples are currently employed and have been so for at least three years. Such households likely are much less of an unemployment risk than the households used to estimate β_u^{a1} and β_u^{a2}, because these latter groups include individuals who actually are experiencing unemployment in period t. This lower risk is most likely recognized by currently employed households, so that our SCF-sample households will hold lower precautionary balances than a household whose true unemployment risk is $\Pr^{a1}(\hat{u}_j)$ or $\Pr^{a2}(\hat{u}_j)$. This would bias down the precautionary effects.[23]

Including Unemployed Households and Those at Their Current Jobs for Less than Three Years: The potential biases in the estimates using $\Pr^{a1}(\hat{u}_j)$ and $\Pr^{a2}(\hat{u}_j)$ highlight the importance of accounting for the relationship between employment status and unemployment risk. The theoretical model highlighted another issue related to employment status—the importance of accounting for reductions in precautionary balances that may have been caused by recent spells of unemployment.

We controlled for this influence in our empirical work by removing the currently unemployed and those who have may have recently been unemployed from our sample. To quantify the effect of this stratification, we reestimated the interaction pooled model augmenting the baseline sample with households whose heads currently were unemployed (152 observations) or had been employed at their current job for less than three years (1339 observations). $\Pr(\hat{u}_j)$ for the unemployed was estimated using $\Pr^{a2}(\hat{u}_j)$, and dummy variables were added to identify the new groups. All else equal, the unemployed and job-switchers hold 0.07 and 0.03 less $W_j/\hat{Y}_j^p$, respectively, than other households, with the differences statistically significant. As shown in the lower right-hand column, the augmented sample produces coefficients on $\Pr(\hat{u}_j)$ and $\Pr(\hat{u}_j) \times \ln \hat{Y}_j^P$ that are about one-third the

[23] Although $\Pr^{a1}(\hat{u}_j)$ and $\Pr^{a2}(\hat{u}_j)$ are fairly highly correlated with $\Pr(\hat{u}_j)$, they contain important differences. Because $\Pr(\hat{u}_j)$ is a consistent estimate of period-$(t + 1)$ $\Pr(u_j)$ for people employed in period t, we can examine potential biases in the alternatives by running the regressions $\Pr(\hat{u}_j) = a + b \Pr^a(\hat{u}_j)$ using only currently employed individuals and testing $a = 0$ and $b = 1$. For each alternative and year, we easily rejected the null hypotheses. And with only one exception, b was smaller than one (averaging 0.82 for a1 and 0.29 for a2), and the mean of $\Pr(\hat{u}_j) - \Pr^a(\hat{u}_j)$ was negative (averaging -0.005 for a1 and -0.45 for a2); thus, the $\Pr^a(\hat{u}_j)$ may overstate the level and variation in $\Pr(u_j)$ for the currently employed.

TABLE 9.—HOUSING WEALTH CONSIDERATIONS

		Alternative Wealth Measures That Exclude Housing		Alternative $\Pr(\hat{u})$	
	Base Model	Net Financial Assets	Net Worth Excluding Residence	Excluding House-Price Effects	Excluding Loan-Term Effects
$\Pr(\hat{u})$	−30.64	0.14	−9.23	−24.46	−22.49
	(−3.34)	(0.10)	(−1.53)	(−2.83)	(−2.62)
$\Pr(\hat{u}) \times \ln \hat{Y}^P$	3.15	−0.02	6.04	2.46	2.24
	(3.45)	(−0.15)	(0.90)	(2.87)	(2.64)
$\ln \hat{Y}^P$	0.00	0.01	0.01	0.02	0.03
	(0.04)	(1.65)	(0.18)	(0.56)	(0.84)
$\partial(W_j/\hat{Y}_j^P)/\partial \Pr(\hat{u}_j) \times \hat{Y}^P$ percentile:					
10th	0.01	−0.00	−0.01	−0.00	−0.01
	(0.64)	(−0.63)	(−0.79)	(−0.17)	(−0.50)
30th	0.06	−0.00	−0.00	0.03	0.02
	(2.39)	(−0.68)	(−0.00)	(1.38)	(0.91)
50th	0.17	−0.00	0.01	0.10	0.08
	(3.07)	(−0.64)	(0.44)	(2.03)	(1.58)
70th	0.30	−0.00	0.03	0.20	0.16
	(3.43)	(−0.58)	(0.72)	(2.43)	(2.02)
90th	0.43	−0.01	0.06	0.29	0.24
	(3.62)	(−0.51)	(0.94)	(2.66)	(2.28)
Overidentify restr. p-value	0.34	0.18	0.11	0.18	0.19
Exclusion p-value	0.01	0.06	0.21	0.00	0.00

Asymptotic t-statistics shown in parentheses. All estimates based on pooled 1983, 1989, and 1992 data.

size of those in our base case. $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ remains an increasing function of $\hat{Y}_j^p$, with effects similar to the baseline results at low $\hat{Y}_j^p$, but between one-half and one-third smaller at the 50th percentile of the income distribution and higher. The t-statistics on all the precautionary measures are similar to those in our base case. Together, these results support the view that precautionary motives can be masked to some degree by households that are more likely to have recently experienced bad draws in the income or employment lottery. Furthermore, the OID test now rejects the region exclusion restriction. This suggests that some wealth-related behavior of the households added to our sample may be correlated with regional employment prospects; for example, in high-job-loss regions household heads may be more likely to have recently switched jobs and to have drawn down precautionary reserves.

G. *Liquid Assets, Housing Wealth, and Precautionary Saving*

Engen and Gruber argue that households probably allocate precautionary balances to assets that can be liquidated at little cost. The second column of table 9 shows results when we define wealth as their liquid net financial assets (NFA) grouping of checking and savings accounts, certificates of deposit, stocks, bonds, and mutual funds less unsecured liabilities. The results change markedly: The $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ are now close to zero at all income levels and are no longer statistically significant. We also reestimated the model using a wealth variable that just excludes the equity in a households' principal residence from our original measure of net worth (column 3). The $\partial(W_j/\hat{Y}_j^p)/\partial \Pr(\hat{u}_j)$ are larger than in the NFA model, but still statistically insignificant at all income levels.[24]

These results point to home equity as a driving force behind the relationship between total net worth and employment risk presented in table 7. Our theoretical model, and most other theoretical analyses, yield little guidance as to what may be causing this finding, because they focus on a net-worth measure comprising 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.

Are the Results Spurious? Our results could be spurious—by chance, regions with high house prices may have had higher unemployment rates during the periods we examine, generating a positive correlation between $\Pr(\hat{u}_j)$ and net-worth measures that include housing. Indeed, when we reran the logistic regression (12) for $\Pr(\hat{u}_j)$ replacing regional dummies with the levels (H_r) and three-year growth rates (G_r) of regional median house prices (both in 1992 dollars), the levels entered with positive and statistically significant coefficients for 1983 and 1992. Still, house prices are not the whole regional story: Likelihood ratio tests soundly reject the restrictions on the regional effects imposed by these regressions in 1989 and 1992.

[24] Using 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 some subaggregates of the balance sheet—namely, very liquid assets and nonhousing, nonbusiness wealth.

To quantify this potential effect, we attempted to strip the influence of regional house price variation out of $\Pr(\hat{u}_j)$ by replacing the coefficient on each regional dummy, β_r, with $\beta_r - \beta_H H_r - \beta_G G_r$ (β_H and β_G are the logit coefficients on H_r and G_r) and using the resulting adjusted $\Pr^H(\hat{u}_j)$ in our pooled model. The results, shown in the fourth column of table 9, continue to indicate significant precautionary behavior, although somewhat weaker than in the base model. At the 30th percentile of the $\hat{Y}^p_j$ distribution, $\partial(W_j/\hat{Y}^p_j)/\partial \Pr(\hat{u}_j)$ is about half the size of the base case, and its t-statistic drops to 1.4; the values for higher percentiles are 35% to 40% smaller than in the base case, but continue to be statistically significant. These results suggest that our findings may be inflated somewhat by spurious correlation between house prices and regional unemployment, but they probably are not driven solely by such effects.

Why Precautionary Wealth May Be Reflected in Housing: There are a number of reasons why consumers may hold 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. Or, as Carroll and Samwick (1998) argue, illiquid assets may not be undesirable buffers if a household's chief concern is a high cost but low probability event such as job loss. For such scenarios the transaction costs to tap an illiquid 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.

It may also 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.

There also are less direct routes for precautionary balances to show up in housing wealth. Dunn (1998) shows that an increase in the probability of unemployment can cause consumers to delay purchasing a new home. This may reflect precautionary motives that make them unwilling to incur the nontrivial reduction in net worth associated with paying mortgage points and real estate commissions (see Carroll and Dunn, 1997). If a large portion of these costs were to be paid out of equity in an existing house, then homeowners deferring purchases would be left with relatively more housing equity than those who upgraded. Similarly, a household may forgo taking out a home equity loan in response to a rise in unemployment risk. Either case would show up as a positive relationship between home equity and uncertainty.

Potential Influences from the Credit Supply Side of Mortgage Markets. Supply-side considerations may also influence the link between housing wealth and unemployment risk. For example, lenders may require a lower loan-to-value ratio for homeowners facing high job-loss risk. Alternatively, higher-risk households might choose a lower loan-to-value ratio because lenders might otherwise impose higher interest rates, stricter mortgage insurance requirements, or higher origination charges. Either set of circumstances would boost the housing wealth held by higher-risk households. Of course, if higher-risk households did not feel a precautionary need to maintain higher total wealth, they could offset the extra home equity by adjusting some other item on their balance sheet, such as credit card debt. Our base-case estimates indicate that they do not do so; instead, any loan-term effects from the supply side are at most only partially offset, so that total net worth still rises with unemployment risk. That said, especially stringent requirements on loan-to-value ratios could lead a higher-risk household to hold more total wealth than would be predicted by our simple theoretical model. In that case, the precautionary effects estimated in our base model would reflect not only the conventional type of precautionary behavior, but also a separate supply-side channel through which unemployment risk affects net worth.

We attempted to purge the effect of loan terms to gauge how important these channels might be. We first gathered data on loan-to-value ratios and initial fees and charges levied by major mortgage lenders delineated by Census subregion.[25] We then removed these variables' influence from our unemployment risk measure by estimating the CPS logistic regression (12) using the lending terms instead of the regional dummies, creating an adjusted $\Pr^{LT}(u_j)$ by subtracting the lending-term effects from the regional dummy coefficients in the original $\Pr(u_j)$ regression, and rerunning the pooled model using $\Pr^{LT}(u_j)$ calculated in the SCF sample. As shown in the right-hand column of table 9, $\partial(W_j/\hat{Y}^p_j)/\partial \Pr(\hat{u}_j)$ remains an increasing function of permanent income, with the estimates economically significant for higher portions of the income distribution. However, the effects are only about half as large as in our base case and only marginally statistically significant at the 50th percentile of $\hat{Y}^p$. Note that the difference between these results and those in the base case likely is an upper bound on a separate supply-side link between housing wealth and unemployment risk. This is because the loan-term-purging exercise

[25] Because 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 individual-specific heterogeneity in lending terms. The concern is whether our regional dummies are capturing region-specific supply-side lending factors.

removes both the voluntary reallocation of wealth (including precautionary balances) toward housing and any additional pure supply-side boost to net worth. Accordingly, the purged results can serve as a lower bound on desired precautionary behavior—which still is of economic and statistical significance for many households in the income distribution.

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 that may be correlated with idiosyncratic tastes to identify variation in unemployment risk by using a variable—the region in which a household resides—that presumably is uncorrelated with unobserved determinants of wealth. 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. In turn, this likely aids us in distinguishing how precautionary behavior might vary with household wealth and income.

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 include and that exclude the self-employed. However, the precautionary response 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, because housing is generally considered one of the least liquid assets that a family can hold.

REFERENCES

Angrist, Joshua D., and Alan B. Krueger, "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 1992), 328–336.

Blanchard, Olivier, and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity* 1 (1992), 1–61.

Browning, Martin, and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature* 34 (December 1996), 1797–1855.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb, "Alternative Transformations to Handle Extreme Values of the Dependent Variable," *Journal of the American Statistical Association* 83 (March 1988), 123–127.

Caballero, Ricardo J., "Consumption Puzzles and Precautionary Saving," *Journal of Monetary Economics* 25 (January 1990), 113–136.

Campbell, John Y., and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence" (pp. 185–216), in Olivier Jean Blanchard and Stanley Fischer (Eds.), *NBER Macroeconomics Annual 1989* (Cambridge, MA: The MIT Press, 1989).

Carrington, William J., "Wage Losses for Displaced Workers: Is It Really the Firm that Matters?" *Journal of Human Resources* 28 (Summer 1993), 435–462.

Carroll, Christopher D., "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* 2 (1992), 61–156.

——— "How Does Future Income Affect Current Consumption?" *Quarterly Journal of Economics* 109 (February 1994), 111–148.

——— "Buffer Stock Saving: Some Theory." Department of Economics, Johns Hopkins University, manuscript (2001).

Carroll, Christopher D., and Wendy E. Dunn, "Unemployment Expectations, Jumping (*S*, *s*) Triggers, and Household Balance Sheets," in Benjamin S. Bernanke and Julio Rotemberg (Eds.), *NBER Macroeconomics Annual 1997* (Cambridge, MA: The MIT Press, 1997).

Carroll, Christopher D., and Andrew Samwick, "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40 (September 1997), 41–71.

——— "How Important is Precautionary Saving?" *Review of Economics and Statistics* 80 (August 1998), 410–419.

Clark, Kim B., and Lawrence H. Summers, "Labor Market Dynamics and Unemployment: A Reconsideration," *Brookings Papers on Economic Activity* 1 (1979), 13–60.

Curtin, Richard T., F. Thomas Juster, and James N. Morgan, "Survey Estimates of Wealth: An Assessment of Quality," in Robert E. Lipsey and Helen Stone Tice (Eds.), *The Measurement of Saving, Investment, and Wealth* (Chicago: University of Chicago Press, 1989).

Diamond, P. A., and J. A. Hausman, "Individual Retirement and Savings Behavior," *Journal of Public Economics* 23 (1984), 81–114.

Dunn, Wendy E., "Unemployment Risk, Precautionary Saving and Durable Goods Purchase Decisions," Federal Reserve Board Finance and Economics Discussion Series working paper no. 1998-49 (November 1998).

Dynan, Karen E., "How Prudent Are Consumers?" *Journal of Political Economy* 101 (December 1993), 1104–1113.

Engen, Eric M., and Jonathan Gruber, "Unemployment Insurance and Precautionary Saving," *Journal of Monetary Economics* 47 (June 2001), 545–579.

Friedman, Milton, *A Theory of the Consumption Function* (Princeton, NJ: Princeton University Press, 1957).

Gourinchas, Pierre-Olivier, and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica* 70 (January 2002), 47–91.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review* 87 (March 1997), 192–205.

Guiso, Luigi, and Tullio Jappelli, "Risk Sharing and Precautionary Saving" (pp. 246–270), in Albert Ando, Luigi Guiso, and Ignazio Visco (Eds.), *Saving and the Accumulation of Wealth* (Cambridge, U.K.: Cambridge University Press, 1994).

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, "Earnings Uncertainty and Precautionary Saving," *Journal of Monetary Economics* 30 (November 1992), 307–337.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* 40 (June 1994), 59–126.

——— "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103 (April 1995), 360–399.

King, M. A., and L.-D. L. Dicks-Mireaux, "Asset Holdings and the Life-Cycle," *The Economic Journal* 92 (June 1982), 247–267.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica* 58 (January 1990), 53–73.

Kuehlwein, Michael, "A Test for the Presence of Precautionary Saving," *Economics Letters* 37 (December 1991), 471–475.
Laibson, David I., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112 (May 1997), 443–478.
Lusardi, Annamaria, "Precautionary Saving and Subjective Earnings Variance," *Economic Letters* 57 (December 1997), 319–326.
——— "On the Importance of Precautionary Saving," *American Economic Review* 88 (May 1998), 448–453.
Normandin, Michel, "Precautionary Saving: An Explanation for the Excess Sensitivity of Aggregate Consumption," *Journal of Business and Economics Statistics* 12 (April 1994), 205–219.
Starr-McCluer, Martha, "Health Insurance and Precautionary Saving," *American Economic Review* 86 (March 1996), 285–295.
Summers, Lawrence H., and Christopher D. Carroll, "Why is U.S. National Saving So Low?" *Brookings Papers on Economic Activity* 2 (1987), 607–636.
Welch, Finis R., "dm11: Matching the Current Population Survey," *Stata Technical Bulletin Reprints* 2 (1993), 4–5.
Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics* 104 (May 1989), 275–298.

APPENDIX

Data Sources and Methods

1. Survey of Consumer Finances

1.a Definitions of selected variables

Net worth: 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 market value; other assets are valued at face value. Net worth includes the value of defined contribution pension plans, but excludes the value of defined benefit pension plans and durable goods other than vehicles.

Net financial assets: The value of all checking accounts, savings accounts, certificates of deposit, stocks, bonds, and mutual funds minus the value of all unsecured liabilities.

Income: Total before-tax household income for the calendar year preceding the SCF interview. (The SCF does not record households' tax payments or liabilities.)

Self-employment: Dummy for whether respondent was self-employed, as opposed to being employee of a private business or government (1983 wave) or working for someone else (1989, 1992 waves).

Occupation: Dummies for (1) managerial and professional, (2) technical, sales, and support, (3) services, (4) precision, craft, and repair, (5) operators and laborers, and (6) farming, forestry, and fisheries.

Industry: Dummies for (1) agricultural, (2) mining and construction, (3) manufacturing, (4) wholesale and retail trade, (5) FIRE and business services, (6) other private services, and (7) public administration.

Region: Dummies for the nine Census subregions.

Education: Dummies for (1) no high school degree, (2) high school degree only, (3) some college, and (4) college degree.

Defined benefit pension: Dummy for whether respondent or spouse is enrolled in a plan with some form of defined benefit features.

Retirement status: Dummy for household head reported being partly or fully retired.

1.b Deflation

Net worth and other nominal variables were divided by the deflator for personal consumption expenditures from the National Income and Product Accounts.

1.c Non-public-use data

The 1989 public-use data set has limited information on region and industry, and neither the 1989 nor the 1992 public-use data set has indicators for the area probability sample. The SCF staff helped us by running our regressions on their internal data sets, which have this detail.

1.d Sample selection

In addition to the restrictions discussed in the text, we exclude households whose heads are younger than 20 or older than 65, since employment risk likely is irrelevant for them. To remove some obviously extreme outliers, we also dropped households in the lowest and highest 0.1 percentiles of wealth and income.

2. Current Population Survey

2.a Linking CPS observations

The CPS interviews a household for four consecutive months, rotates it out of the sample for eight months, and then rotates it back for four more months. We used the outgoing rotation groups (the fourth and eighth interviews) to create two readings on employment and demographics taken one year apart. To match records, we used programs provided by Welch (1993).

2.b Sample selection

SCF interviewing took place over periods of 6 to 8 months. To avoid potential problems with seasonality in employment and unemployment, our baseline $\Pr(u_j)$ reflects information from households whose fourth interview fell sometime during the twelve months preceding the end of the SCF sampling periods. The backward-looking alternative $\Pr(u_j)$ uses households whose eighth interview fell in that twelve-month period; the unconditional $\Pr(u_j)$ uses households whose fourth or eighth interviews fell in the period. To avoid the considerable reduction in sample size associated with limiting the CPS sample to household heads, we included dummies in the CPS regressions for whether the record was for the head of household (and interacted with age, race, and gender). Persons not in the labor force at the time of the fourth or eighth interview were excluded from the sample, and we applied the same age restrictions as for the SCF.

2.c Employment status

Usual CPS definitions of employed, unemployed, and self-employed were used.

2.d Demographic variables

These variables were defined to be consistent with those constructed for the SCF. In our base $\Pr(u_j)$, they reflect the profile of an individual as of their fourth interview. The backward alternative also uses demographic information from the fourth interview, which occurred sometime preceding the SCF sampling window. The unconditional alternative uses data from either the fourth or eighth interview during the 12 month window.

3. Housing Data

3.a House prices

We calculated state-level median home prices in 1980 using the 1980 Census. We then created state-level time series by extrapolating the 1980 prices using state-level indices of the change in the price of existing homes from Freddie Mac. Finally, we averaged the state series together, weighting by population, to obtain time series of home prices by region.

3.b Loan terms

We used loan-to-price ratios and initial fees and charges collected by the Federal Housing Finance Board and covering conventional home mortgages issued by major mortgage lenders. The data are recorded at the state level; we used state population as weights to aggregate them to the nine Census subregions.