

NBER WORKING PAPER SERIES

**HOW IMPORTANT IS
PRECAUTIONARY SAVING?**

**Christopher D. Carroll
Andrew A. Samwick**

Working Paper No. 5194

**NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1995**

We would like to thank Angus Deaton, Karen Dynan, Janice Eberly, Eric Engen, Jonathan Gruber, Jerry Hausman, Glenn Hubbard, Spencer Krane, Robin Lumsdaine, Joseph Matthey, Mark McClellan, James Poterba, Karl Scholz, Jonathan Skinner, Steve Zeldes, and seminar participants at the Federal Reserve Board and the North American Winter Meetings of the Econometric Society, 1993, for useful comments. The views expressed are those of the authors and do not necessarily reflect the views of the Board of Governors or the staff of the Federal Reserve System. Financial support from the National Institute of Aging and the Lynde and Harry Bradley Foundation is gratefully acknowledged. This paper is part of NBER's research programs in Aging, Economic Fluctuations and Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1995 by Christopher D. Carroll and Andrew A. Samwick. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

**HOW IMPORTANT IS
PRECAUTIONARY SAVING?**

ABSTRACT

We estimate the fraction of the wealth of a sample of PSID respondents that is held because some households face greater income uncertainty than others. We first derive an equation characterizing the theoretical relationship between wealth and uncertainty in a buffer-stock model of saving. Next, we estimate that equation using PSID data; we find strong evidence that households engage in precautionary saving. Finally, we simulate the wealth distribution that would prevail if all households had the same uncertainty as the lowest-uncertainty group. We find that between 39 and 46 percent of wealth in our sample is attributable to uncertainty differentials across groups.

Christopher D. Carroll
Mail Stop 80
Board of Governors of the Federal
Reserve System
20th and C Streets, NW
Washington, DC 20551

Andrew A. Samwick
Department of Economics
6106 Rockefeller Hall
Dartmouth College
Hanover, NH 03755-3514
and NBER

I. Introduction

Several recent papers have attempted to determine the proportion of either aggregate or household wealth attributable to precautionary saving. Unfortunately, theoretically plausible models of precautionary saving are difficult to solve and have been thought to imply no simple empirically estimable relationship between uncertainty and wealth.¹ The literature has therefore tended to adopt one of two strategies, which we will distinguish as the theoretical and the empirical approaches. Papers following the theoretical approach (e.g. Hubbard, Skinner, and Zeldes, 1994; henceforth, HSZ) assume that a particular theoretical model of saving is correct and calculate the amount of precautionary wealth implied by the model, without estimating any empirical relationship between uncertainty and saving or wealth.² Papers in the empirical literature (e.g. Guiso, Jappelli, and Terlizzese, 1992) estimate an empirical relationship but use theoretically implausible precautionary saving models whose chief virtue is their ability to generate a closed-form analytical expression which can serve as an econometric specification.³

Neither approach has achieved consensus. Among the theoretical papers, the estimates of the proportion of wealth attributable to precautionary motives range from as little as 0.7 percent (Krussell and Smith, 1994) to 50 percent (HSZ) or more. Even within a given paper, the answer depends greatly upon parametric choices--the HSZ estimates actually range from

¹By "theoretically plausible" we mean, at a minimum, models in which the consumer's utility function exhibits Decreasing Absolute Risk Aversion (DARA). See Kimball (1991b) for arguments that utility is of the DARA form. The most commonly used specific model in this class is one where the utility function is of the Constant Relative Risk Aversion (CRRA) form.

²HSZ do compare the aggregate amount of wealth generated by their model to aggregate household wealth in the U.S., but they do not present any formal statistical estimation or test of the model.

³The empirical papers generally assume that utility is of the Constant Absolute Risk Aversion (CARA) form. See our discussion in Section II, Kimball (1990b), or Deaton (1992), for arguments that CARA utility is implausible.

30 percent to 70 percent. Among empirical papers the range is similarly broad: Guiso, Jappelli, and Terlizzese state that "precautionary saving accounts for 2 percent of households' net worth," while Dardanoni (1991) claims that "more than 60 percent of savings ... arise as a precaution against future income risk."

This paper improves upon the existing literature by empirically estimating an econometric model whose specification is derived from a theoretically plausible model of precautionary saving. To be specific, Deaton (1991) and Carroll (1992, 1995) argue that important features of personal saving behavior are well captured by a version of the standard dynamic optimization model in which consumers engage in "buffer-stock" saving behavior over most of their working lifetimes. Carroll (1992, 1995) shows that one characteristic of buffer-stock savers is that they will aim to achieve a target wealth-to-income ratio, and the size of the target wealth stock is, in part, a function of the degree of income uncertainty.

One of the obstacles to empirical estimation of theoretically plausible models has been that in these models there is no single sufficient statistic to summarize the degree of uncertainty consumers face. In principle, optimal behavior depends upon the entire structure of the income distribution, so that, for example, two distributions which exhibit the same mean and variance might induce quite different saving behavior, depending on the higher moments of the distributions of income shocks. One of the contributions of this paper is to show that, in practice, if households behave according to a buffer-stock model of saving, there are at least two measures of uncertainty that do a good job of summarizing the optimal precautionary saving response to income distributions like those that households actually face. One of these is based on a simple theoretical measure of uncertainty derived by Kimball

(1990a), the Equivalent Precautionary Premium (EPP); the other is an atheoretical measure, the log of the variance of the log of income (LVARLY). We show that the buffer-stock model predicts a roughly linear relationship between the log of target wealth and either a normalized version of the EPP or LVARLY.

Armed with this specification, we next turn to empirical estimation of the slope coefficient in the relationship between wealth and uncertainty. To be specific, for each household in our PSID sample, we calculate the EPP and LVARLY, and then, using instrumental variables to correct for measurement error, we estimate the empirical relationship between log wealth and our measures of uncertainty. We find a statistically significant relationship between all tested measures of wealth and both measures of uncertainty.

In the final section of the paper we use our empirical estimates to answer the question posed in the title of the paper. We find that, abstracting from general equilibrium considerations, setting the uncertainty of every household in our sample to the smallest predicted uncertainty for *any* household in the sample would reduce total net worth in our sample by about 44 percent; net worth exclusive of housing and business equity by 46 percent; and very liquid asset holdings by 39 percent. We also show that the fraction of wealth attributable to income uncertainty is greater at lower deciles of the income distribution.

Before presenting our own results, we begin with a more detailed discussion of the existing literature.

II. Survey and Critique of the Literature

Although the pioneering theoretical work on precautionary saving was done over twenty-five years ago by Leland (1968), Sandmo (1970), and Drèze and Modigliani (1972),

the question of whether precautionary saving is an empirically important phenomenon was largely neglected until more recent theoretical research demonstrated qualitatively that the presence of a precautionary saving motive might resolve a variety of macroeconomic and microeconomic puzzles about consumption behavior and the distribution of household wealth.⁴ Determining whether precautionary behavior provides the *correct* resolution of these puzzles depends upon establishing the quantitative importance of precautionary saving. We will discuss first the theoretical literature and then the empirical literature.

In the typical paper that employs the theoretical approach, a model for aggregate or individual saving that includes a precautionary motive is parameterized, and the level of wealth is simulated under parameterizations in which precautionary saving is important and other parameterizations in which precautionary saving is unimportant (generally, by setting an uncertainty parameter to zero). The resulting difference in wealth is interpreted as the portion of wealth attributable to precautionary saving. The papers that have employed the theoretical approach have used a variety of models and have reached a variety of results. To our knowledge, the earliest example of the theoretical approach is Skinner (1988), who uses a Taylor approximation to a life cycle model of consumption for consumers with CRRA utility. He finds that under his preferred parameter values, precautionary saving accounts for 56 percent of total wealth. Caballero (1991) uses the analytical solution to a permanent income model with CARA utility to conclude that "reasonable parameters are consistent with levels of wealth originating from precautionary attitudes toward earnings uncertainty, in excess of 60

⁴See Caballero (1990), Carroll (1992), and Hubbard, Skinner, and Zeldes (1994) for macroeconomic puzzles and Kimball (1990a,b), Carroll (1995), Carroll and Samwick (1995), Samwick (1994) and Hubbard, Skinner, and Zeldes (1995) for microeconomic puzzles.

percent of the observed net U.S. wealth." Hubbard, Skinner, and Zeldes (1994) find that a vastly more sophisticated, and very elaborately specified and calibrated, life cycle model of consumption without uncertainty predicts an aggregate wealth-to-income ratio around half as large as the actual ratio in the U.S. but that, under their preferred parameter values, adding income uncertainty increases the model's predicted aggregate wealth-to-income ratio to close to the actual ratio.⁵

In contrast, theoretical general equilibrium models typically find small effects of precautionary saving. Aiyagari (1994) solves a standard growth model modified to incorporate precautionary saving and shows that for empirically plausible parameter values, uninsured idiosyncratic income risk accounts for a less than 3 percentage point increase in the aggregate saving rate but that, for sufficiently high variability and persistence in shocks to income, the aggregate saving rate can be 7-14 percentage points higher than in the model with no income uncertainty.⁶ A recent paper by Krusell and Smith (1994) finds an even smaller effect: in their general equilibrium model, aggregate wealth is increased by only 0.7 percent by precautionary saving arising from idiosyncratic income risk.

There are several reasons why the general equilibrium models may substantially underestimate the amount of precautionary saving engaged in by individual households. First, some general equilibrium models calibrate their measure of income uncertainty using data on

⁵Hubbard, Skinner, and Zeldes actually consider the effects of three separate sources of uncertainty in their papers: labor income, medical care expenditures, and mortality (length of life). However, they find that labor income uncertainty is by far the most important.

⁶In Aiyagari's general equilibrium model, the saving rate in the absence of income uncertainty (or in the presence of full insurance) is about 24 percent. Thus, these absolute increases in the saving rate correspond to increases in saving of no more than 3/24 or about 12 percent in the typical case and as much as 14/25 or 56 percent in the high variability and persistence case.

aggregate personal income, on the theory that idiosyncratic household-level risk is diversifiable in perfect capital markets. But, as shown by Cochrane (1991) and others, household-level risk is not diversified, or at least not fully diversified, so the appropriate measures of income uncertainty are those that apply to individual households -- which are much greater than those for aggregate income.

Second, the general equilibrium models were originally calibrated in such a way as to induce a capital stock large enough to match the U.S. capital stock *without* precautionary saving. When income uncertainty is added to such models, the effect is necessarily small, because the wealth accumulated for nonprecautionary reasons now serves a dual purpose: it is held primarily to achieve low-frequency consumption-smoothing goals, but if needed it is available to finance consumption in the event of bad income draws. A more realistic calibration would incorporate precautionary motives for saving initially.

A third reason the general equilibrium models may yield lower estimates of the importance of precautionary saving at the household level is through the effect that precautionary saving has on the interest rate. If households face a large amount of uncertainty, and as a result save a great deal, the aggregate capital stock rises. The higher capital stock depresses the real interest rate and thereby induces households to save less for life-cycle reasons. This is a true general equilibrium effect, and one not captured by our simulations of the effect of uncertainty on saving, largely because addressing it properly would require us to greatly extend the analysis and to commit to a particular general equilibrium model and a particular set of assumptions about tastes, the aggregate production function, and the relationship between the microeconomic and macroeconomic structure of

uncertainty.⁷ Additionally, we note in support of our methodology that, ultimately, the general equilibrium results should differ substantially from the partial equilibrium results only to the extent that the interest elasticity of saving is large; our reading of the literature is that there is little solid empirical evidence that the interest elasticity of saving is even positive, and virtually no evidence that it is large.⁸

The empirical approach to determining the magnitude of precautionary saving has been as inconclusive as the theoretical approach. The method of the empirical approach is to estimate the parameters of a model that allows for precautionary saving and then to use the resulting estimates to project the fraction of wealth due to uncertainty. The extremes of the empirical estimates are occupied by Dardanoni (1991) and Guiso, Jappelli, and Terlizzese (1992). Dardanoni (1991) examines a cross-section of U.K. households and shows that average consumption across occupational and industry groups is negatively related to the intragroup variance of income. As noted in the introduction, he finds that precautionary saving accounts for over 60 percent of all saving. Guiso, Jappelli, and Terlizzese (1992) use a cross-sectional dataset of Italian households in which consumers report an expected distribution of income for the next year and find that consumption is only marginally lower and wealth only slightly higher for consumers with a greater subjective variance of next year's income. They conclude that precautionary saving is negligible.

The disparity in the predictions of the empirical studies may be the result of two problems in implementation. First, both rely on only a single year of data to derive their

⁷See Poterba (1994) for a survey of the current state of research on the appropriate economic model for personal saving behavior when both life-cycle and precautionary motives are operative.

⁸See, for example, Hall (1988).

measures of income uncertainty. Dardanoni (1991) uses the cross-sectional dispersion of income within different occupation or industry groups as his measure of uncertainty. For instance, if the degree of cross-sectional income dispersion is greater across different kinds of service workers than across different farm laborers, then service workers are assumed to face greater income uncertainty. Using panel data on income from the PSID, Carroll and Samwick (1995) show that the correlation is actually negative between Dardanoni's intragroup income dispersion measure and a more theoretically appropriate measure of uncertainty, the variance of shocks to permanent income.

Guiso, Jappelli, and Terlizzese (1992) use a survey in which households were asked explicitly about the distribution of their expected income in the next year. However, the theoretically appropriate measure of uncertainty in their model is the variance of lifetime income, not the variance of next year's income. Strong assumptions about the income process are required in order for the two concepts to be identical (or linearly related). Additionally, it is clear that a substantial fraction of respondents did not fully understand the questions asked (for instance, many households responded with point expectations for variables which are clearly subject to some degree of uncertainty). These results may imply that people simply think about uncertainty in terms that were not well captured by the questions asked in the survey.

A second, and potentially more serious, problem is that the equations estimated in both Dardanoni and Guiso, Jappelli, and Terlizzese are derived from models in which consumers have Constant Absolute Risk Aversion (CARA) utility functions. Such models have a host of undesirable properties, including the prediction that, under plausible parameter values,

consumption will be negative for young households. Furthermore, the implausibility of holding absolute risk aversion constant can be illustrated with a simple example: CARA utility predicts that a person consuming a million dollars a year would react in exactly the same way as a person consuming \$10,000 a year to a 50 percent chance of having consumption reduced by \$5,000 next year. Our intuition suggests that absolute risk aversion should be decreasing, because the wealthier consumer will likely engage in less precautionary saving against a risk of a given dollar amount.⁹

III. Theoretical Framework

The Consumer's Intertemporal Optimization Problem

The model of precautionary saving that forms the basis of our empirical work is a variant of the "buffer-stock" models developed by Deaton (1991) and Carroll (1992, 1995). These models employ a standard intertemporal optimization framework in which households maximize expected discounted future utility given a stochastic process for income and a budget constraint. The models assume that consumers are both *prudent*, in Kimball's (1990a) sense of having a precautionary saving motive, and *impatient*, in the sense that, if income were perfectly certain, they would wish to borrow against future income to finance current consumption.¹⁰ In the presence of income uncertainty, these two assumptions can generate the implication that consumers will have a target wealth-to-income ratio. Carroll (1992) argues that this model is consistent with a variety of characteristics of macroeconomic data on

⁹See the sources cited in footnote 3 for more detailed analysis of CARA utility.

¹⁰Note that this does not require consumers to have a high (or even positive) rate of time preference. With positive expected income growth, consumers can satisfy the "impatience" condition even with no discounting of future utility.

consumption and saving, and Carroll (1994, 1995) and Carroll and Samwick (1995) find support for the model using microeconomic data.

The particular version of the buffer-stock model considered here imposes liquidity constraints directly, as in Deaton (1991), although similar results can be obtained in a version without liquidity constraints.¹¹ Specifically, the consumer is assumed to solve the following problem:

$$\begin{aligned}
 (1) \quad & \max_{\{C_t\}} E_s \sum_{t=s}^{\infty} \beta^{t-s} u(C_t) \\
 & \text{s.t. } X_{t+1} = R[X_t - C_t] \\
 & \quad Y_t = P V_t \\
 & \quad W_s \geq 0 \quad \forall s
 \end{aligned}$$

where $X_t = Y_t + W_t$ is the stock of physical resources available for spending in period t , $R = (1+r)$ is the gross interest rate, $\beta = 1/(1+\delta)$ is the discount factor (δ is the discount rate), Y_t is the total labor income of the household in period t , P is the permanent labor income of the household (that is, the income that would be earned if there were no transitory shocks), and V_t is the multiplicative transitory shock to income in period t . The utility function is of the CRRA form: $u(c) = c^{(1-\rho)}/(1-\rho)$, where ρ is the coefficient of relative risk aversion. We assume that the V_t are i.i.d. and will use empirical distribution functions calculated from the

¹¹We chose to impose liquidity constraints here because it simplifies the task of constructing an appropriate distribution function for the income shocks; see the discussion below of our kernel estimates of the distribution function.

PSID income data to estimate the distribution of the V_t .¹²

Our choice of a CRRA utility function guarantees that the consumer in this model will engage in precautionary saving; the coefficient of relative risk aversion ρ indexes the strength of both risk aversion and prudence. The model also satisfies our definition of theoretical plausibility by exhibiting decreasing absolute risk aversion.¹³ The drawback of assuming CRRA utility is that there is no closed-form, analytical solution for the level of consumption, wealth, or saving as a function of uncertainty. We therefore must solve the model numerically.

As Deaton (1991) and Carroll (1992, 1995) show, in the model with CRRA utility the optimal consumption rule can be written as a relationship between the ratio of consumption to permanent income, C/P , and the ratio of gross wealth X to permanent income, X/P . If lower case variables denote the upper case counterpart normalized by permanent income, this means that the optimal consumption rule can be written as $c_t(x_t)$ --the optimal consumption ratio is a function of the gross wealth ratio and nothing else (see Deaton (1991) for a proof).

The model and the method of solving it are described in Carroll (1995) and Deaton (1991), so we will confine the discussion here to an overview of our parametric assumptions and a broad description of the behavior of consumers in this model. The principal necessary condition for generating buffer-stock saving behavior is that, if income were certain, consumers with no assets would wish to borrow against future income in order to finance

¹²We also discuss the consequences for our results if there is a permanent as well as a transitory component to income shocks.

¹³The coefficient of absolute risk aversion at any consumption level, c , is given by: $ARA(c) = -u''(c)/u'(c) = \rho/c$. Taking the derivative with respect to c yields: $ARA'(c) = -\rho/c^2 < 0$, showing that absolute risk aversion is decreasing.

current consumption; the analytical condition which guarantees this in the continuous-time version of the model with only transitory shocks to income is $\rho^{-1}(r-\delta) < g$, where g is the expected growth rate of income (see Carroll (1995) for a derivation).¹⁴ The particular values we choose for solving the model are $\rho = 3$, $\delta = 0.04$, $r = 0$, and $g = 0.02$, but the results of the analysis in this section are qualitatively similar under a broad range of other parameter values provided consumers are prudent ($\rho > 0$) and impatient ($\rho^{-1}(r-\delta) < g$). The assumptions about the stochastic process for income are described below. The end result is an optimal rule which indicates how much should be consumed at each possible level of the wealth-to-income ratio. As described in Carroll (1992, 1995), this consumption rule implies a target wealth-to-income ratio such that if wealth is greater than the target, impatience will outweigh prudence and the consumer will plan to dissave, whereas if wealth is less than the target, prudence will outweigh impatience and the consumer will plan to save.

Measuring Income Uncertainty

In order to examine the model's predictions about the relationship between target wealth and income uncertainty, the model must be solved under an array of different assumptions about income uncertainty. We therefore turned to the PSID, which contains the kind of panel data on income necessary to calculate distributions of shocks to income.

The PSID extract used in this paper contains labor income data for the years 1981 through 1987 for a panel of 4,346 households who are not part of the poverty subsample of

¹⁴The discrete-time version is $(R\beta)^{-1/\rho} < G$, where $G = 1+g$. The economic logic behind these equations is as follows. A standard result for the continuous time CRRA utility model without income uncertainty or liquidity constraints is that the desired growth rate of consumption is $\rho^{-1}(r-\delta)$. Consider a consumer with zero assets. If the desired growth rate of consumption is below the growth rate of income, then if the intertemporal budget constraint is to be satisfied it must be true that the level of consumption is above the level of income, and so such a consumer would be running down his net worth.

the survey.¹⁵ We restrict the sample to households in which the same person was the head in each sample year, the marital status of the head did not change over the period, the head was between 25 and 50 years old over the whole period,¹⁶ for which there are valid data on occupation, education, wealth, and income, and for which measured wealth is positive (so that we can take its log).¹⁷ In addition to improving the average quality of data in the sample, these restrictions tend to remove households for whom the variation in reported income may be due to (possibly planned) changes in household composition or transitions into retirement. The households that remain in our sample therefore probably face less income uncertainty than the typical household, but the analysis below suggests that even in this sample, income uncertainty is responsible for a sizable portion of wealth holdings. Imposing these sample restrictions also narrows the sample size to about 900 households, with the exact number depending on the measure of wealth being considered. A more detailed breakdown of the sample restrictions is given in the Appendix.

As suggested by the literature review in Section II, using an appropriate measure of

¹⁵The main components of labor income as used in this study are the wage, salary, and self-employment income reported by all members of the household. Our measure also includes all welfare payments, such as those from Aid to Families with Dependent Children (AFDC) and Supplemental Security Income (SSI); all receipts from entitlement programs, such as Social Security, Veterans Administration or other pensions, unemployment insurance, and workers' compensation; and private assistance such as alimony, child support, and help from relatives. In light of this construction, a more descriptive term would be "noncapital" income; we use the term labor income for expositional convenience. The impact of social insurance and other transfer programs on precautionary saving is therefore manifested in lower estimates of income uncertainty for households that receive payments from such programs.

¹⁶We restrict the sample to people 50 or younger because Carroll (1995) and Carroll and Samwick (1995) argue on theoretical and empirical grounds that the buffer-stock model that we use to derive the econometric specification is most appropriate as a description of the behavior of households younger than about 50. When older households are added to the sample, the estimated coefficients on uncertainty are increased and their statistical significance is greater; these results are presented in the appendix. We discuss this issue further below.

¹⁷We also test whether our empirical results are robust to the exclusion of households with nonpositive wealth. We find little difference in our empirical estimates when we use alternative methods which correct for (or are not subject to) the sample selection restriction.

income uncertainty is critical to any empirical estimation of the importance of precautionary saving. Ideally, the consumer's optimization problem under uncertainty would generate a particular measure of uncertainty that would be exactly correct, in the sense that it would be a sufficient statistic for the amount of precautionary saving that would be induced by a given income distribution. Unfortunately, theory does not provide such explicit guidance about how to measure uncertainty. However, it is possible to derive a measure of uncertainty that is at least loosely based on theory: We construct a measure related to the Equivalent Precautionary Premium (EPP) defined by Kimball (1990a), who shows that the EPP is, in essence, a direct measure of the *intensity* of the precautionary saving motive *at the point of zero precautionary saving*.

Suppose that consumption is distributed randomly with a multiplicative shock V around a level \bar{c} , $c = \bar{c} X$. In this case Kimball's EPP would be defined by the amount ψ such that:

$$u'(\bar{c} - \psi) = E[u'(c)]$$

Under the CRRA utility function assumed above, $u'(c) = c^{-\rho}$, implying that we can solve for ψ :

$$\begin{aligned} (\bar{c} - \psi) &= E[c^{-\rho}]^{-\frac{1}{\rho}} = E[(\bar{c}X)^{-\rho}]^{-\frac{1}{\rho}} = \bar{c}E[X^{-\rho}]^{-\frac{1}{\rho}} \\ \psi &= \bar{c}(1 - E[X^{-\rho}]^{-\frac{1}{\rho}}) \end{aligned}$$

For our later empirical purposes, a scaleless measure of relative uncertainty is more useful; such a measure is given by $\psi/\bar{c} = 1 - E[X^{-\rho}]^{-1/\rho}$. We will call this measure the Relative Equivalent Precautionary Premium (REPP).

Because total consumption is not reported in the PSID, we cannot construct a measure of uncertainty that corresponds exactly to Kimball's EPP or the REPP.¹⁸ Instead, we follow Carroll (1994), who constructs a measure of income uncertainty that is analogous to the REPP. In the consumption model specified above, a household i in each period t receives income of $Y_{i,t}$ that equals permanent income P_i multiplied by the period- t value of an i.i.d. transitory shock $V_{i,t}$, so $Y_{i,t} = P_i V_{i,t}$. Substituting permanent and actual income for average and actual consumption (respectively) in the REPP formula gives $1 - E[V_i^{-\rho}]^{-1/\rho}$. Here, the shocks to consumption have been replaced by shocks to income; the income and consumption versions of the REPP would be identical only if the household consumed exactly its income in each year.¹⁹

In our sample we observe only actual realizations of income, so we must construct an estimator for the theoretical REPP. Our estimator is given by:

$$REPP_i = 1 - \left[\frac{1}{7} \sum_{t=1981}^{1987} \left(\frac{Y_{i,t}}{P_i} \right)^{-\rho} \right]^{-\frac{1}{\rho}}$$

If the level of permanent income is well approximated by the average income realization of a detrended income series, this amounts to replacing the expectation of $V_i^{-\rho}$ in the formula for

¹⁸Food consumption and selected other components of consumption are measured, but we consider the quality of these data to be too poor to produce credible results about the extent of uncertainty. Shapiro (1982) has estimated that 95 percent of the variation in the PSID food consumption data is noise; Runkle (1991) estimates that 70 percent is noise.

¹⁹Because prudent households build up a "buffer-stock" of assets precisely to insure consumption against shocks to income, the random element in consumption will be less variable than $Y_{i,t}$. The relationship between the two, however, is monotonic; increases in the variability of $V_{i,t}$ will increase the variability of consumption.

the theoretical REPP with the average of the ρ -exponentiated ratios of the seven income observations to the average income observation.

The theoretical reason to prefer the REPP over other measures of income uncertainty is that its magnitude is influenced by the convexity of the consumer's marginal utility function as indexed by ρ , the same characteristic that governs the size of the precautionary response to uncertainty in the consumer's first-order condition. Of course, under some circumstances, the REPP will not have an advantage over atheoretical measures of income uncertainty, such as the variance. For example, it has been common in previous work to assume that utility is CARA (i.e. $u(c) = -\exp(-\theta c)/\theta$), and that the shock to income is additive and distributed normally with a variance of σ_v^2 . Under these assumptions, the EPP has a theoretical value of $\theta\sigma_v^2/2$.²⁰ The normality assumption guarantees that the variance is a sufficient statistic for the dispersion in the distribution, and the CARA assumption gives the precise relationship this simple linear form. Thus, the specifications in this previous work eliminate by assumption any special role for the REPP. Indeed, it is exactly this simplification that allows closed-form analytical solutions to be derived between wealth and income variance in CARA models with additive normal shocks (See Caballero (1991) for an example).

An analogous relationship can be derived for the REPP if utility takes the more plausible CRRA form as we have assumed. If the shocks to income are lognormally distributed, then the REPP is approximately equal to a linear function of the variance of the

²⁰Because the shock is additive rather than multiplicative, there is no reason to scale the EPP by \bar{c} to get the REPP. Using Kimball's notation, the EPP (ψ) is determined by the equation:

$$e^{-\theta(\bar{c}-\psi)} = E[e^{-\theta c}] = e^{-\theta(\bar{c}-\theta\sigma^2/2)}$$

log of income. Define $v_{i,t} = \ln(V_{i,t}) \sim N(0, \sigma_{\ln v}^2)$. Because $\ln(X) \sim N(0, \sigma_x^2)$ implies that $\ln(E[X]) = \sigma_x^2/2$, it can be shown that $E[V_i^{-\rho}] = \exp(\rho^2 \sigma_{\ln v}^2 / 2)$, and, consequently that $(E[V_i^{-\rho}])^{-1/\rho} = \exp(-\rho \sigma_{\ln v}^2 / 2)$. Therefore, if V is lognormally distributed, then $REPP = 1 - \exp(-\rho \sigma_{\ln v}^2 / 2)$. Using the approximation that, if z is close to zero, $\exp(z) \approx 1 + z$, this expression can be further simplified to $REPP \approx 1 - (1 - \rho \sigma_{\ln v}^2 / 2) = \rho \sigma_{\ln v}^2 / 2$.

The crucial feature of the REPP is that it is disproportionately sensitive to large negative shocks when the distribution of shocks is asymmetric and thus not perfectly summarized by a mean and variance. For example, consider an income process in which three realizations of income are equally likely, $A = \{0.8, 1.0, 1.2\}$. The sample mean and variance of these shocks are 1.0 and 0.04, respectively. Another set of three equally likely realizations with the same mean and variance is $B = \{0.8697, 0.9, 1.2303\}$. Assuming a coefficient of relative risk aversion of 3, the REPP distinguishes between these two sets as follows:

$$REPP_A = 1 - \left(\frac{1}{3} (0.8^{-3} + 1.0^{-3} + 1.2^{-3}) \right)^{-\frac{1}{3}} = 0.0529$$

$$REPP_B = 1 - \left(\frac{1}{3} (0.8697^{-3} + 0.9^{-3} + 1.2303^{-3}) \right)^{-\frac{1}{3}} = 0.0436$$

$REPP_A$ is in this case 21 percent higher than $REPP_B$ because the 20 percent negative shock is weighted disproportionately relative to the two negative shocks of 13.03 and 10 percent.²¹

Target Wealth as a Function of Uncertainty

The foregoing analysis has suggested that when utility is of the theoretically plausible

²¹The magnitude of REPP, and the difference between these two cases, is increasing in the coefficient of relative risk aversion. For example, if $\rho = 6$, $REPP_A = 0.0760$, which is 30 percent higher than the value of $REPP_B = 0.0580$.

CRRA form and the distribution of shocks to income is potentially asymmetric, the REPP may more accurately reflect the strength of the precautionary saving motive than will the variance of income (the appropriate measure when utility is CARA). This constitutes our rationale for examining the REPP as a measure of income uncertainty in empirical analyses of precautionary saving. However, even the REPP does not have a closed-form analytical relationship with the target wealth-to-income ratio.

Our solution to this problem is to solve our model numerically for a variety of potential income distributions and to examine the relationship between the REPP calculated for each of those distributions and the model's predicted target wealth. The purpose of these simulation exercises is to suggest an appropriate econometric specification for the relationship between the REPP and target wealth. Given the approximate closed-form relationship derived above between the REPP and the variance of shocks to log income when income shocks are lognormally distributed, a natural way to proceed would be to assume that the distribution of the V 's is lognormal. However, statistical tests strongly reject lognormality for the V 's in our data. In particular, the empirical distribution exhibits considerable excess kurtosis, especially in the lower tail. This problem is potentially of great concern because theory predicts that the amount of precautionary saving is very sensitive to the possibility of large negative shocks to income.

Our approach was to construct nonparametric kernel estimates of the density function for the income shocks, thereby avoiding the potentially serious errors that might arise if we were to choose an improper functional form for the distribution of income shocks. Our method was as follows. We divided our sample into subsamples corresponding to the eight

occupation categories, twelve industry categories, and six education categories to which the head of household could belong, for a total of 26 different groups. For each household i , we calculated $\hat{V}_{i,t}$ as the ratio of observed income in year t to the mean (detrended) income over all seven years of the sample period. We treat each of these as an independent realization of V . Thus, if there are n households in a given group, this technique produces $7n$ observations on V for that group. The empirical CDFs of these \hat{V} 's for each of the 26 groups were then approximated by twenty-point kernel estimators. That is, the full distributions of the \hat{V} 's within each group were partitioned into twenty intervals of equal probability mass, and the kernel estimator for the distribution was constructed by replacing each actual value of $\hat{V}_{i,t}$ by the mean value of $\hat{V}_{i,t}$ in that interval. For the kernel estimates of the distribution of shocks to income for each of the 26 groups, we then solved the buffer-stock model described above and calculated the model's implied target wealth-to-income ratio. We also calculated the average value of the relative equivalent precautionary premium and the variance of the level of income and of the log of income for each of the categories.

Figure 1a plots the log of the model's predicted target wealth-to-income ratio for consumers facing each distribution of income against the REPP calculated for that distribution. Figure 1b plots the target log wealth-to-income ratio against the variance of income (henceforth, VARY), the theoretically appropriate measure of uncertainty under CARA utility. Comparison of the two figures shows that VARY is not as closely related to target wealth as is REPP, as expected from the foregoing discussion. Figure 1c plots the log target wealth-to-income ratio against the variance of the log of income (VARLY). Somewhat to our surprise, this measure performs a bit *better* than the REPP in explaining target wealth

holdings; we discuss implications of this finding below.

Table 1 formalizes the message of the figures using univariate regressions of the log of the model's target wealth-to-income ratio (W/P) on the various measures of income uncertainty. The first three lines of the table represent the regressions whose fitted lines are plotted in Figures 1a through 1c. The table confirms that the variance of the (detrended real) level of income is considerably less useful in explaining the model's predicted target wealth than either the REPP or the variance of the log of income. This is unsurprising, because the variance of the level of income (VARY) gives equal weight to positive and negative shocks, while REPP and VARLY considerably enhance the relative contribution of negative shocks, which have a strong impact on precautionary saving.

The next regression confirms another apparent result of Figures 1a-c: VARLY is more closely related than is REPP to the theoretical target wealth ratio. The second three regressions consider the relationship between the log of (W/P) and the log of the same three measures of uncertainty. Based on the goodness of fit measure, the REPP performs slightly worse and the VARY measure performs slightly better than in the first set of regressions, but their relative ranking is unchanged. The striking result is that when estimated in a constant elasticity specification, the log of VARLY performs even better than VARLY itself, with a t-statistic of 40.85 and an R^2 of 0.9852.

For simplicity, in our empirical work we wished to narrow the field of potential measures of uncertainty to two. We chose the REPP for its theoretical appeal and LVARLY because it performs better than any of the alternatives in Table 1.²²

²²We also performed the analysis using all the measures of uncertainty in Table 1. We found that results for REPP, LVARLY and LVARY were approximately equally good, and the results for all of these measures were

One possible objection to our procedure for constructing REPP and relating target wealth to it is that we used the same ρ in solving the model and in constructing our measure of REPP. When using REPP to analyze actual household wealth data, we do not know the "true" value of ρ ; our procedure would be problematic if the relationship between REPP calculated with a given ρ were a bad indicator of the target wealth that would result from assuming a different ρ . We checked whether this was a problem by calculating REPP under several different plausible assumptions about ρ , and regressing the target wealth generated for our baseline value of ρ on the REPP calculated for non-baseline values of ρ . In all cases we found R^2 's not much different from those in Table 1; in our later empirical results, we also experimented with different values of ρ and found little difference in the econometric results.

Another potential criticism of the procedure we used to produce Table 1 is that we have assumed that income shocks are transitory and independent over time. An alternative approach, pursued in Carroll and Samwick (1995), is to allow for both transitory and permanent shocks to income. They show that if both transitory and permanent shocks are i.i.d. and lognormally distributed, simple econometric formulas yield estimates of the variances of the transitory and permanent shocks. Carroll and Samwick (1995) estimate an empirical regression equation similar to ours but use as their measure of uncertainty the estimated transitory and permanent variances rather than the REPP or LVARLY.

This alternative approach has the advantage of treating the persistence of income shocks more realistically than we do, but it also has drawbacks. First, unlike the REPP, the variance measures used by Carroll and Samwick (1995) are not derived from the theory of

substantially better than the results for the other three measures.

precautionary saving. Second, because the assumption of lognormality of the shocks is key to the decomposition of shocks into transitory and permanent components, the Carroll and Samwick (1995) methodology does not allow us to make nonparametric estimates of the density function for shocks. As discussed above, if the income shocks are not lognormally distributed, variance measures like those used in Carroll and Samwick (1995) will be imperfect indicators of the true risk differences across groups. Finally, and most importantly, with two measures of uncertainty rather than one, it becomes substantially more difficult to decide what the appropriate experiment is for determining how much aggregate wealth is the result of differences in income uncertainty, particularly if (as Carroll and Samwick find) the true correlation between the two measures of uncertainty is difficult to determine.

One way to gauge how problematic our assumption of independent shocks may be is to examine the consequences if the actual structure of income shocks is like that in Carroll and Samwick (1995), but we nevertheless use the procedures outlined above to construct REPP and LVARLY and then regress the log of the wealth-to-income ratio on those measures of uncertainty. We solved a buffer-stock model of saving for a matrix of values for the standard deviation of transitory and permanent shocks ranging from 0.025 to 0.125 per year, generating a total of 25 possible configurations of uncertainty. We then calculated the REPP, VARY, and VARLY that would result from simulating each of those income processes for 7 years, and regressed the log of the target wealth-to-income ratio implied by the buffer-stock model on the calculated measures of uncertainty. The results are plotted in Figure 2 and summarized statistically in Table 2.

The results in Table 2 confirm the proposition that even when the income process is

very different from the process assumed in constructing REPP and LVARLY, there is still an approximately linear relationship between REPP or LVARLY and the log of the target wealth ratio. In these results, the uncertainty measures in levels outperform those in logs, with R^2 's of around 0.965 versus about 0.85. The relative performance for the REPP is slightly better than in Table 1, reflecting the nonlinearity that exists in the relationship when the shocks to income are larger.

Our conclusion from Tables 1 and 2 is that the buffer-stock model implies that the relationship between income uncertainty and the wealth-to-income ratio can be well represented econometrically by an equation based on:

$$(2) \quad \log\left(\frac{W}{P}\right) = a_0 + a_1\omega$$

where ω is a measure of uncertainty, either REPP or LVARLY. We also conclude from Table 2 (and further simulations not reported) that, even if the income process is very different from the i.i.d. process we assumed, our summary statistics will do a reasonably good job of capturing the degree of income uncertainty consumers face.²³

IV. Empirical Estimation of the Model

Basic Specification

Our econometric specification is derived from equation (3):

²³A final point that is logically necessary is to show that the coefficient on the uncertainty measures would not be positive in the absence of a precautionary saving motive. The proof for the standard certainty equivalence model is straightforward; we omit it for brevity.

$$(3) \quad \log\left(\frac{W}{P}\right) = a_0 + a_1 \omega + \nu$$

Adding $\log(P)$ to both sides gives:

$$(4) \quad \log(W) = a_0 + a_1 \omega + \log(P) + \nu$$

Our final specification is a slightly more general version of this equation:

$$(5) \quad \log(W) = a_0 + a_1 \omega + a_2 \log(P) + a_3'Z + \nu$$

where the Z variables are demographic controls for age, race, sex, marital status, and the number of children. Equation (5) does not constrain the coefficient on $\log(P)$ to be unity, thereby allowing the wealth-to-income ratio to vary with the level of permanent income.

The buffer-stock model of saving presented above assumes that there is only one, perfectly liquid, asset. The model's predictions about target wealth concern total net worth held in this single asset. We therefore estimate equation (5) using total net worth (NW) as the dependent variable. In reality, of course, consumers can invest in a wide range of assets which differ, among other ways, in their degree of liquidity.²⁴ Illiquid assets may be less useful as a safeguard against bad income shocks because of the extra time or money required to turn them into the cash needed to meet emergency expenses or to replace income. It would therefore not be surprising to find that holdings of more liquid assets are more sensitive to uncertainty. Consequently, we also estimate equation (5) for two progressively

²⁴Research on the effect of income uncertainty on the portfolio decisions of households is less developed than that on precautionary saving. Theoretical studies include Kimball (1991a) and Elmendorf and Kimball (1991), and some empirical work is presented by Guiso, Jappelli, and Terlizzese (1993). In addition to liquidity, these studies show that another important factor in portfolio allocation of precautionary saving is the covariance of asset returns with the shocks to income.

more liquid measures of wealth (the exact components of which are detailed in the Appendix): net worth excluding equity in the main home and in personally owned businesses (non-housing, non-business wealth, NHNBW), and very liquid assets (VLA) which can be liquidated on short notice with small transactions costs.²⁵

In the absence of a theoretical framework that explicitly incorporates liquidity, we simply note that the proposition that precautionary balances should be held in highly liquid forms is not as strong as it may at first appear. If the main motivation in holding precautionary assets is to self-insure against rare but large shocks to income (such as a prolonged spell of unemployment),²⁶ it may well be worthwhile to pay the transactions costs required to liquidate illiquid assets in the rare circumstance that the dreaded event actually occurs. Moreover, because our VLA measure does not subtract any debts from the measured assets, it is only weakly related to the concept of net worth in our model. As a result, the case for preferring one measure of wealth over the other two is not particularly strong.

An important assumption in the foregoing theoretical analysis is that consumers engage in buffer-stock saving behavior. The validity of our econometric specification is

²⁵In the Appendix, summary statistics of the three wealth measures in the PSID are compared to those from an analogous sample of households from the *Surveys of Consumer Finances, 1983*, a survey which oversampled high-income households and collected more detailed data on wealth than the PSID. The medians of the two samples are extremely close, typically within \$1,000 of each other. However, the PSID does not capture the extreme upper tail of the wealth distribution that is captured in the SCF and, consequently, the PSID data somewhat understates the mean. This finding echoes the results of Curtin, Juster, and Morgan (1989) who more carefully compare the SCF and PSID and conclude that the latter provides nearly as good a measure of wealth as the former for most of the population. One of our implicit justifications for using the log rather than the level of wealth in our empirical specifications is that, if regressions are performed in levels rather than logs, a small number of extremely wealthy households can greatly skew the empirical results.

²⁶Carroll (1992) solves a buffer stock model with lognormal shocks to annual transitory and permanent income and with a small probability that income goes to zero for the entire year, which he interprets as long spells of unemployment. Despite the assumption that such events are very rare, he finds that a large fraction of the buffer is attributable to the fear of these zero income events rather than to the annual transitory and permanent shocks.

therefore only assured for a sample in which households can reasonably be expected to engage in buffer-stock saving behavior. Carroll (1995), Samwick (1994), and Carroll and Samwick (1995) argue that a variety of empirical evidence is consistent with the view that households engage in buffer-stock saving behavior until roughly age 50, but begin to behave more like traditional life cycle savers in the years immediately preceding retirement.²⁷ During this immediate preretirement period, the amount of precautionary saving is still positively related to the degree of income uncertainty, but the consumer's expected wealth stock in each year is also a function of the timing and magnitude of the income drop at retirement as well as the precise structure of the age/income profile before and after retirement. We therefore limit our sample to households in which the head is no older than 50 during the sample period (although we also report the results obtained for the full sample including older households).

The final issue to be addressed before presenting the estimation results is the nature and justification of the instrumenting procedure used. Because both uncertainty and permanent income are unquestionably measured with error in our data, they must be instrumented if we are to obtain consistent parameter estimates. Our instrument set contains dummies for the occupation, education, and industry of the head of household in 1981, along with the demographic variables already contained in Z . We also interacted the occupation

²⁷The critical issue is the age at which the steady-state solution to the infinite-horizon problem in equation (1) ceases to be a good approximation to the optimal consumption rule in the analogous finite-horizon problem with a realistic life-cycle pattern of income growth rates. Recall that consumers must be impatient in order to engage in buffer-stock behavior, a requirement which is $(R\beta)^{(1/\phi)} < G$ in the infinite horizon, discrete-time problem. In the finite horizon problem, loosely speaking the right-hand side of the inequality is related to the expected annual growth rates in all remaining years of life. Because income growth slows as consumers get older and there is generally a discrete drop in income at retirement, at some point the condition no longer holds, and the consumer switches from buffer-stock type saving behavior to more traditional retirement saving.

and education variables with the age and age² terms in order to allow for different lifetime profiles of income and uncertainty for different occupation and educational groups. (The set of instruments is described fully in the Appendix).

A literature has recently developed on the potential biases in instrumental variables regressions that result when the instruments are only weakly correlated with the variables they are instrumenting.²⁸ Staiger and Stock (1994) conclude that "... if the F-statistic from the first stage is small, then second-stage statistics are nonnormal and can be badly biased and conventionally-constructed confidence intervals will have poor coverage rates." Bound, Jaeger, and Baker (1993) suggest that both the partial R² and F-statistic on the excluded instruments from the first stage regression be used as guides to the quality of IV estimates. For the first-stage regressions with REPP as the dependent variable, the partial R² is 0.12 and the F-statistic is 2.29. When LVARLY is the dependent variable, the statistics are 0.12 and 2.43, and when the log of income is the dependent variable, they are 0.24 and 7.44, respectively. The corresponding p-values on the F-statistics are all below 0.00001, suggesting a sufficient degree of predictive power of the instruments for income and uncertainty.

Results for our estimates of equation (5) for all three measures of wealth, and for both measures of uncertainty, are presented in Table 3.²⁹ The coefficients on the REPP and LVARLY terms are highly significant for all three measures of wealth, with VLA having a

²⁸In addition to the articles cited above, see Nelson and Startz (1990a,b).

²⁹In constructing REPP, we assumed a coefficient of relative risk aversion of three, for consistency with the simulation results in Tables 1 and 2. We also tried coefficients of 2 and 5; the empirical results were not materially different.

somewhat lower coefficient than NHNBW or NW.³⁰ The statistical significance of the uncertainty terms increases as the measure of wealth becomes more comprehensive, and for a given wealth measure, the coefficients on the REPP and LVARLY terms are of approximately equal significance. These variables are also highly correlated with each other, and regressions (not reported) in which both measures are included find that neither measure is individually significant. The coefficients on the uncertainty terms also generally lie well within the range of predictions in Tables 1 and 2. The coefficient on the logarithm of permanent income is always well in excess of (and statistically significantly different from) one, the value that would be expected if the wealth-to-income ratio were homothetic. Although this is, strictly speaking, a violation of our theoretical model, it accords with the common empirical finding that people with high permanent incomes hold disproportionately higher amounts of wealth.

The coefficients on the other covariates in the regressions are also plausible. Controlling for income and uncertainty, the partial effect of age on NHNBW and NW is positive for households until they reach their mid-40s and slightly negative thereafter. Married households accumulate more wealth, but hold less of it in very liquid form, than do unmarried households. Households with children tend to have less wealth, but the statistical significance of this effect declines as the measure of wealth becomes more comprehensive. Households in which the head is white have higher levels of wealth. Controlling for income and other demographic variables, the effect on wealth of having a female head of household is insignificantly different from zero.

Our instrumental variables specification is econometrically identified by the exclusion

³⁰Heteroskedasticity tests rejected the null hypothesis of homoskedasticity at the 5 percent level, so all standard errors are heteroskedasticity-robust.

of occupation, education, and industry variables from the regression of wealth on uncertainty. This exclusion restriction assumes that these instruments have no predictive power for wealth other than through their correlation with the level of permanent income and with uncertainty. We test this null hypothesis by performing the standard heteroskedasticity-robust overidentification test from Hansen (1982); results are reported in the last column of the table. In no case does the OID test reject the model at the 5 percent level, but when total net worth is the dependent variable the model is rejected at the 10 percent level for both measures of uncertainty.

Table 4 presents the results when the sample is extended to include households with heads in the 51-56 age group in 1981 (i.e., those still younger than 63 by the end of the sample period in 1987). The coefficient estimates on all of the measures of uncertainty increase and are more highly statistically significant. However, the OID tests perform considerably worse, rejecting at significance levels between 1 and 10 percent for both measures of uncertainty for both VLA and NW. Only when NHNBW is the dependent variable does the model pass the OID test at the ten percent level.

These results are exactly the pattern that would be expected based on the analyses of Carroll (1995) and Carroll and Samwick (1995). They show that for plausible parameter values, households will engage in buffer-stock saving behavior until retirement is approximately 15 years away and then will switch over to more traditional life cycle saving behavior in the remaining period before retirement. They also show that when the household is not in the buffer-stock phase, a given change in income uncertainty generates a larger saving response. Carroll (1995) shows that during the buffer-stock saving phase, target

wealth is determined primarily by the nature and degree of income uncertainty and the coefficient of relative risk aversion and is not very sensitive to the expected profile of income growth. Because the age profile of income varies by occupation, industry, and education (see, for example, Murphy and Welch (1990) and Carroll (1994)), the life cycle model that characterizes saving behavior in the later phase implies that saving ought to depend strongly on occupation, industry, and education in ways not captured by the relation of these variables to the level of permanent income or to income uncertainty. Thus, the analysis of Carroll (1995) implies that the OID test should not reject for households before the switchover from buffer-stock to traditional life-cycle saving but should reject for consumers in the life-cycle phase--exactly the pattern exhibited in Tables 3 and 4. As a further test, we also restricted the sample to households under age 45 and, as expected, the OID tests found even less evidence of misspecification than in Table 3 (all p-values were above 0.2 except for one which was 0.154) and the coefficient estimates were not much changed.³¹

Tests of Robustness

The estimates in Table 3 provide solid evidence for the existence of a precautionary saving motive; controlling for age and other demographics, households who face more income uncertainty hold more wealth. The empirical findings mirror the theoretical results in finding little difference between the predictive power of REPP and LVARLY for wealth.

As noted above, we restrict our sample to households with positive wealth so that we can use the log of wealth as our dependent variable. Although liquidity constraints were imposed in the theoretical model, they may not be so rigidly enforced in reality, so taking the

³¹When the sample is further restricted to those 40 and younger, the p-values were all above 0.3, and the coefficient estimates again remained about the same.

log of wealth eliminates from the sample the small number of households who report zero or negative wealth. There are two ways around such a truncation problem. The first is to estimate the regression in levels and rely on robust estimation methods to minimize the consequences of the inappropriate assumption that the error term is additive. To estimate such an equation, we ran median regressions of equation (5) in levels (using the predicted values of the level of income and uncertainty from a first stage least squares regression). The uncertainty terms were still significant at the 1 percent level in the NHNBW and NW regressions and at the 5 and 10 percent levels in the VLA regressions using REPP and LVARLY, respectively.

A second way to accommodate the truncation is to estimate a two-step model in which the effect of the truncation on the expected value of the error in the wealth equation is accounted for by predicting the probability of reporting nonpositive wealth. Using the instruments to predict whether wealth was positive and then instrumenting for the inverse Mills ratio yielded similar results to the median regressions. The uncertainty terms in the NHNBW and NW regressions were still significant and had almost identical coefficients to those in Table 3. The coefficients on uncertainty in the VLA regressions were reduced by about 25 percent relative to those in Table 3 and were now statistically significant at only the 10 percent level.

V. How Would Wealth Change if There Were Less Uncertainty?

The econometric estimates of the sensitivity of wealth to income uncertainty in Table 3 can now be used to determine the impact of income uncertainty on the aggregate wealth distribution. Most of the existing precautionary saving literature, in both the theoretical and

the empirical strands, has chosen to assess the importance of precautionary saving by comparing a model's predictions when income uncertainty is set to zero with the predictions when uncertainty is set to a value taken to roughly represent reality. We believe that this approach is inappropriate for both the empirical and the theoretical literatures.

Setting uncertainty to zero is dubious in the econometric setting because zero uncertainty is never observed, or even approximated, in the data sets. The farther from the estimating sample that a projection is made, the less precise it is. Even equations which fit very well in sample can make large errors when asked to describe a region of the relevant space far from the region in which the model was estimated.

In the theoretical literature, the problem with setting uncertainty to zero arises because the presence of even a small amount of uncertainty can lead to drastically different behavior than that which arises from a certainty or certainty-equivalent model. Another way to state this point is that there can be a discontinuity in behavior as uncertainty goes from zero to some small $\epsilon > 0$.³² Since in reality uncertainty can never be entirely eliminated, it seems more appropriate to compare the predictions of the model under the assumption of small (but in-sample) amounts of uncertainty to its predictions under typical amounts of uncertainty.

Our approach is therefore to use the empirical model to simulate the distribution of wealth that would prevail if all households faced the same, small, amount of uncertainty, instead of the amount of uncertainty they actually faced. The second-stage regressions

³²This point has been emphasized by Zeldes (1989), Carroll (1992, 1995) and others. A standard example is when a tiny probability of income going exactly to zero in every subsequent period is added to a perfect certainty model with CRRA utility and impatience. Consumers will borrow heavily under perfect certainty but will never borrow with this type of uncertainty.

presented in Table 3 using the REPP generate a wealth equation for each household:³³

$$(6) \quad \log(W_i) = a_0 + a_1 REPP_i + a_2 \log(P_i) + a_3' Z_i + v_i$$

The equation fits exactly because the estimation errors v_i are defined as the difference between the predicted value of log wealth and the actual value.³⁴ The procedure we employ merely substitutes a constant, small, value $REPP^*$ for $REPP_i$ while leaving the values of all other independent variables the same. This procedure indicates how much wealth this household would hold if its uncertainty were given by $REPP^*$ but it were otherwise identical (including the value of the error term v_i). That is, we construct a new measure of wealth, W_i^* , such that:

$$(7) \quad \log(W_i^*) = \log(W_i) - a_1 [REPP_i - REPP^*]$$

Implications about the aggregate wealth distribution can then be derived by aggregating up these simulated values for all the individual households.

One possible choice for $REPP^*$ would be the minimum value of $REPP_i$ observed in the sample. However, the fact that $REPP_i$ for any individual household is measured with error implies that the minimum actual value in the sample will probably correspond not to someone who actually had such a small amount of uncertainty but to someone for whom the *measurement error* for uncertainty was especially large and negative. Instead, we use the minimum *predicted* value of $REPP_i$ from the first-stage regression of REPP on the instrument set. By using the predicted value, the conceptual experiment is to set everyone's uncertainty

³³The simulation results were similar when we used LVARLY.

³⁴These v_i are the residuals from the second stage of the 2SLS regressions; they are not the IV residuals.

to the value that would be predicted for someone with a combination of characteristics such as a stable industry, high education, a secure occupation, etc. that predict a low value of income uncertainty.³⁵ As reference points, we also simulate the wealth distributions that result when REPP* is set to the 10th, 25th, and 50th percentiles of the predicted REPP distribution.

Figures 3a-d plot the simulated values of the log(NHNBW) measure of wealth against the true distribution of log(NHNBW) for each of our four values of REPP*. Specifically, two points are plotted for each household in the sample: (w_i, p_i) and (w_i^*, p_i) where p_i indicates the percentile ranking for household i in the true wealth distribution, w_i indicates the log of actual wealth W_i , and w_i^* indicates the simulated value of the log of wealth for that household when REPP_{*i*} is set to the chosen value of REPP* (i.e. $\log(W_i^*)$ in equation (7)).

In the first figure, where REPP* is set equal to the minimum predicted value of REPP_{*i*} in the sample, all simulated wealth points lie to the left of the actual wealth distribution. Everyone's uncertainty has been reduced (except the households that already had the minimum predicted value of REPP); consequently, everyone's wealth holdings are reduced. In the second figure, 10 percent of the points lie to the right of the original distribution; in the third figure, 25 percent are to the right; and in the final figure, half of the points are to the right and half are to the left.

The figures give a graphical depiction of how much the logarithm of wealth of individual households is changed by the reduction of uncertainty. The procedure for

³⁵This will presumably be an overestimate of the minimum amount of uncertainty faced by anyone in the sample, and for this reason the estimate of the effect setting everyone's uncertainty to the minimum predicted value is almost surely an *underestimate* of the true effect of setting uncertainty to the minimum *possible* value. In our sample, the minimum actual value is only 13 percent of the minimum predicted value of REPP.

calculating the effect on aggregate wealth given the changes at the household level is straightforward: if each household's simulated $\log(\text{wealth})$ is given by w_i^* , then its simulated level of wealth is given by $\exp(w_i^*)$. Simulated aggregate wealth is then given by the sum of the simulated wealth of the individual households. The results are shown in Table 4 for each of the three categories of wealth. The first row shows that reducing every household's uncertainty to the minimum predicted value reduces aggregate very liquid assets by 39 percent, aggregate non-housing, non-business wealth by 46 percent, and total net worth by 44 percent. The next three rows of the table present the results for the simulations in which REPP^* is set to the 10th, 25th, and 50th percentiles of REPP_i 's predicted value.

Our final simulations concern the effects of income uncertainty on wealth for households in different permanent income categories. The theoretical prediction for whether income uncertainty should be more important for higher or lower income households is ambiguous. Because wage and employment insecurity is something that prudent consumers would pay to avoid, we expect that households with more skills would tend to have lower income uncertainty, just as they have higher wages. As shown in the Appendix, for example, the average values of REPP and LVARLY decline with increasing educational attainment. On the other hand, it is also conceivable that workers who are willing to accept more income uncertainty are paid higher wages as a compensating differential relative to those of equal ability who are unwilling. Additionally, firms may attempt to induce greater effort by their employees by tying their compensation to output. If successful, such a scheme could result in higher wages and greater income uncertainty.

The first panel of Table 5 calculates the amount of wealth that is estimated (by the

procedure described above) to be due to precautionary saving for people in each decile of the permanent income distribution. Precautionary saving accounts for a larger portion of wealth for people in the lower income deciles, accounting for more than half of the wealth of the households in the bottom decile. Proceeding up the lower half of the permanent income distribution, the proportion of wealth due to precautionary motives steadily declines, suggesting that people with greater skills take their higher compensation in the form of both a higher level and a more certain distribution of income. As other factors become more important in the upper half of the income distribution, the precautionary component of wealth is smaller than in the lower half, but there is no clear pattern by decile or measure of wealth.

VI. Conclusion

The absence of a simple measure of uncertainty that embodies all of the relevant characteristics of a stochastic income distribution has plagued empirical analyses of whether precautionary saving is an economically important phenomenon. An important contribution of this paper is to show that if consumers behave according to a buffer-stock model of saving and face plausible distributions of income shocks, the relationship between the log of target wealth and income uncertainty should be approximately linear when the measure of income uncertainty is based on either Kimball's (1990a) Relative Equivalent Precautionary Premium (REPP) or the log of the variance of the log of income (LVARLY). Our principal econometric finding is that wealth holdings are indeed positively and significantly related to income uncertainty for various measures of wealth and both the REPP and the LVARLY measures of uncertainty.

Simulations of the econometric model suggest that approximately 45 percent of total

net worth and of non-housing, non-business wealth and 40 percent of very liquid assets of the households in our sample are held as a precaution against the systematically greater uncertainty that some households face as compared with others. Our finding that less of very liquid assets than of non-housing, non-business wealth is attributable to precautionary saving suggests that the bulk of precautionary saving exists to insure against relatively large shocks--perhaps substantial spells of unemployment--compared to which the cost of liquidating moderately illiquid assets is small. We also find that precautionary savings constitute a larger fraction of total savings for the households with the lowest permanent income.

An important limitation of our approach is that it does not directly address the question of the proportion of the wealth of consumers over age 50 (or of the entire population) that can be attributed to precautionary saving behavior. However, our empirical results (and those of Carroll (1995) and Carroll and Samwick (1995)) are consistent with a parameterization of the life cycle model under uncertainty which implies that consumers engage in buffer-stock saving behavior until around age 50 and switch over to traditional life cycle retirement saving thereafter. A natural extension would be to estimate the proportion of total wealth attributable to uncertainty, or to differentials in uncertainty across households, by performing simulations like those of Hubbard, Skinner, and Zeldes (1995) but under parameter values that generate buffer-stock saving behavior consistent with our empirical results for consumers under the age of 50. This would constitute a proper melding of the theoretical and the empirical approaches mentioned in the introduction by using a theoretical model whose empirical implications for the relationship between wealth and uncertainty had been explicitly tested.

More broadly, because our results provide direct evidence that precautionary saving is economically important, they bolster the growing consensus that uncertainty and precautionary saving behavior are fundamental to understanding consumption and saving in both macroeconomic and microeconomic areas. The certainty equivalent model increasingly appears to be an inadequate description of how consumers actually behave. While the mathematical apparatus of models with realistic treatments of uncertainty is more cumbersome than the analytically tractable certainty equivalent model, it appears that it is not possible to seriously address important questions--such as why the personal saving rate in the U.S. has declined or what effect unemployment insurance has on consumer behavior--without considering the role of precautionary saving.

References

- Aiyagari, S. Rao. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*. 109 (1994): 659-684
- Bound, John, David A. Jaeger, and Regina Baker. "The Cure Can Be Worse than the Disease: A Cautionary Tale Regarding Instrumental Variables." Technical Working Paper no. 137. Cambridge, Mass.: NBER, 1993.
- Caballero, Ricardo J. "Consumption Puzzles and Precautionary Savings." *Journal of Monetary Economics*. 25 (1990): 113-136.
- . "Earnings Uncertainty and Aggregate Wealth Accumulation." *American Economic Review*. 81 (1991): 859-871.
- Carroll, Christopher D. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity*, no. 2 (1992), 61-135.
- . "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." Manuscript. Federal Reserve Board, June 1995.
- . "How Does Future Income Affect Current Consumption?" *Quarterly Journal of Economics*. 109 (1994): 111-148.
- , and Andrew A. Samwick "The Nature of Precautionary Wealth." Manuscript. Federal Reserve Board, April 1995.
- Cochrane, John H. "A Simple Test of Consumption Insurance." *Journal of Political Economy*. 99 (1991): 957-976.
- Curtin, Richard T.; Juster, F. Thomas; and Morgan, James N. "Survey Estimates of Wealth: An Assessment of Quality." In *The Measurement of Saving, Investment, and Wealth*, edited by Robert E. Lipsey and Helen S. Tice. Chicago: University of Chicago Press, 1989.
- Dardanoni, Valentino. "Precautionary Savings Under Income Uncertainty: a Cross-Sectional Analysis." *Applied Economics*. 23 (1991): 153-160.
- Deaton, Angus. "Saving and Liquidity Constraints." *Econometrica*. 59 (1991): 1121-1142.
- . *Understanding Consumption*. Oxford: Oxford University Press, 1992.
- Drèze, J., and Modigliani, Franco. "Consumption Under Uncertainty." *Journal of Economic Theory*. 5 (1972): 308-335.

- Elmendorf, Douglas W., and Kimball, Miles S. "Taxation of Labor Income and the Demand for Risky Assets." Working Paper no. 3204. Cambridge, Mass.: NBER, 1991.
- Guiso, Luigi, Tullio Jappelli and Daniele Terlizzese. "Earnings Uncertainty and Precautionary Saving." *Journal of Monetary Economics* 30 (1992): 307-337.
- . "Earnings Uncertainty and Portfolio Choices." Manuscript. Rome: Bank of Italy, 1993.
- Hall, Robert E. "Intertemporal Substitution in Consumption." *Journal of Political Economy*. 96 (1988): 339-357.
- Hansen, Lars P. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*. 50 (1982): 1029-1054.
- Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes. "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving." In Allan H. Meltzer and Charles I. Plosser (eds.) *Carnegie-Rochester Conference Series on Public Policy* 40 (1994): 59-126.
- . "Precautionary Saving and Social Insurance," *Journal of Political Economy*. 103 (1995): 360-399.
- Kimball, Miles S. "Precautionary Saving in the Small and in the Large." *Econometrica*. 58 (1990a): 53-73.
- . "Precautionary Saving and the Marginal Propensity to Consume," Working Paper no. 3403. Cambridge, Mass.: NBER, 1990b.
- . "Precautionary Motives for Holding Assets" In *The New Palgrave Dictionary of Money and Finance*, edited by J. Eatwell, M. Milgate, and P. Newman. New York: W. W. Norton, 1991.
- . "Standard Risk Aversion." Technical Working Paper no. 99. Cambridge, Mass.: NBER, 1991.
- Krussell, Per and Anthony Smith. "The Stochastic Growth Model with Heterogeneous Agents and Fixed Costs of Adjustment." Manuscript, The University of Pennsylvania, 1994.
- Murphy, Kevin M. and Finis Welch. "Empirical Age-Earnings Profiles." *Journal of Labor Economics* 8 (1990): 202-229.
- Leland, Hayne E. "Saving and Uncertainty: The Precautionary Demand for Saving." *Quarterly Journal of Economics*. 82 (1968): 465-473.

- Nelson, Charles R. and Richard Startz. "Some Further Results on the Exact Small Sample Properties of the IV Estimator." *Econometrica* 58 (1990a): 967-976.
- Nelson, Charles R. and Richard Startz. "The Distribution of the Instrumental Variables Estimator and Its t-ratio When the Instrument Is a Poor One." *Journal of Business* 63 (1990b): S125-S140.
- Poterba, James M. "Personal Saving Behavior and Retirement Income Modelling: A Research Assessment" Manuscript, Massachusetts Institute of Technology, 1994.
- Runkle, David E. "Liquidity Constraints and the Permanent Income Hypothesis." *Journal of Monetary Economics*. 27 (1991): 73-98.
- Samwick, Andrew A. "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving." Manuscript. Dartmouth College, 1994.
- Sandmo, Agnar. "The Effect of Uncertainty on Saving Decisions." *Review of Economic Studies*. 37 (1970): 82-114.
- Shapiro, Matthew D. "A Note on Tests of the Permanent Income Hypothesis in Panel Data." Manuscript. University of Michigan, 1982.
- Skinner, Jonathan. "Risky Income, Life Cycle Consumption, and Precautionary Savings." *Journal of Monetary Economics*. 22 (1988): 237-255.
- Staiger, Douglas and James H. Stock. "Instrumental Variables Regression with Weak Instruments." Technical Working Paper no. 151. Cambridge, Mass: NBER, 1994.
- Zeldes, Stephen P. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics*. 104 (1989): 275-298.

Appendix

1) Sample restrictions imposed

Several restrictions are imposed on the sample drawn from the PSID to ensure that observed fluctuations in income over the sample period 1981 - 1987 are not unduly influenced by (possibly planned) demographic transitions. The number of observations eliminated with each such restriction or by missing values for key variables are given in the following table:

Sample Restriction	Number of Households Eliminated	Number of Households Remaining
Full Sample		8129
No households from the supplemental Survey of Economic Opportunity (poverty subsample)	3783	4346
No households with the head younger than 26 in 1981	596	3750
No households with the head older than 62 in 1987	796	2954
No households with the head older than 50 in 1981	572	2382
Head of household in 1981 is head for all years 1981-1987.	894	1488
Marital status of head of household is the same in all years 1981-1987.	213	1275
Head of household is in the labor force in 1981 (occupation is reported).	81	1194
Nonmissing industry and education reported in 1981.	33	1161
Nonmissing labor income data reported in all years 1981-1987.	3	1158

All tables below and analyses within the text that depend only on income data use this sample of 1158 households. The number of observations in each regression in Table 3 is also affected by the availability of wealth data as defined in section (4) of this Appendix.

2) *The set of instrumental variables*

Throughout the empirical section, we make use of a set of variables to instrument for income uncertainty and permanent income. They are:

A) Household composition

Variable	Description	Comment
Age	Age in years	Sample mean is 35.6 years.
Age ²	Square of age	
Married	Marital status indicator	83.1% of household heads are married.
White	Race indicator	91.8% of household heads are white.
Female	Sex indicator	11.5% of household heads are female.
Kids	# of children under age 18	Sample mean is 1.4 children.

B) Indicator variables for occupation

Occupation group	Mean REPP	Mean LVARLY	Percent of sample
Professional and Technical Workers	0.086	-3.828	23.66
Managers (not self-employed)	0.092	-3.623	11.57
Managers (self-employed)	0.161	-2.601	4.32
Clerical and Sales Workers	0.100	-3.486	14.16
Craftsmen	0.093	-3.527	20.47
Operatives and Laborers	0.132	-3.154	17.17
Farmers and Farm Laborers	0.301	-1.919	2.68
Service Workers	0.115	-3.301	6.04

The occupation variables are also interacted with Age and Age² to allow for occupation-specific age-income and age-uncertainty profiles.

C) Indicator variables for education

Education group	Mean REPP	Mean LVARLY	Percent of sample
0-8 Grades	0.151	-2.851	3.71
9-12 Grades	0.130	-3.086	9.50
High School Diploma	0.127	-3.318	18.57
Some College, No Degree	0.098	-3.499	39.38
College Degree	0.097	-3.617	19.86
Some Advanced Education	0.104	-3.704	8.98

The education variables are also interacted with Age and Age² to allow for education-specific age-income and age-uncertainty profiles.

D) Indicator variables for industry

Industry Group	Mean REPP	Mean LVARLY	Percent of sample
Agriculture, Forestry, Fishing	0.243	-2.241	3.97
Mining	0.067	-3.861	1.12
Construction	0.149	-2.925	6.65
Manufacturing	0.099	-3.479	28.58
Transportation, Communications, and Utilities	0.088	-3.658	10.62
Wholesale and Retail Trade	0.121	-3.252	14.68
Finance, Insurance, and Real Estate	0.100	-3.475	4.75
Business and Repair Services	0.103	-3.544	3.71
Personal Services	0.165	-2.597	1.55
Entertainment and Recreation Services	0.098	-3.387	0.50
Professional and Related Services	0.101	-3.642	15.98
Public Administration	0.063	-4.069	7.77

3) *The wealth regressions*

The econometric specifications in Tables 3 and 4 have the household composition variables (A) as independent variables, along with permanent income and income uncertainty. The instrument set consists of all the variables given in (A) - (D) above.

4) *Wealth measures*

We use three measures of wealth constructed from the wealth supplement to the 1984 wave of the PSID. They are:

A) Very Liquid Assets (VLA)

Includes balances in checking accounts, savings accounts, money market funds, certificates of deposit, government savings bonds, Treasury bills, shares of stock in publicly held corporations, mutual funds, and investment trusts. Any such assets in IRA's are also included because they are not reported separately.

B) Non-housing, Non-business Wealth (NHNBW)

Includes VLA plus the net value of real estate other than the main home, including a second home, land, rental real estate, and money owed on a land contract; and the net value of vehicles, including cars, trucks, motor homes, trailers, and boats. Outstanding balances on credit cards, student loans, medical or legal bills, and loans from friends are subtracted.

C) Total Net Worth (NW)

Includes NHNBW plus the net equity in the main home and the net value of farms and businesses.

Sample statistics on these measures of wealth, taken from the sample of 1158 households for whom income uncertainty measures can be constructed, are given in the following table, with comparisons to an analogously constructed sample in the *Survey of Consumer Finances*, 1983.

Wealth in the PSID and SCF, Households 50 Years Old and Younger				
Wealth Measure	Number of Observations	Mean	Median	Standard Deviation
VLA				
PSID				
Full Sample	1020	14,748	3,500	50,924
VLA > 0	896	16,789	5,000	54,021
SCF				
Full Sample	1306	22,996	3,369	471,086
VLA > 0	1247	24,118	4,068	482,419
NHNBW				
PSID				
Full Sample	958	33,532	10,075	116,901
NHNBW > 0	860	38,987	12,700	120,554
SCF				
Full Sample	1306	31,853	10,599	103,710
NHNBW > 0	1163	36,665	13,384	107,847
NW				
PSID				
Full Sample	922	82,452	43,325	171,448
NW > 0	874	87,473	47,020	174,545
SCF				
Full Sample	1306	101,374	45,487	338,617
NW > 0	1243	107,164	49,430	346,886
<p>1) All dollar amounts are in real 1984 dollars.</p> <p>2) Dollar amounts for the SCF are weighted to account for the oversampling of high-income households. Dollar amounts for the PSID are unweighted.</p> <p>3) Sample restrictions in the SCF are analogous to those in the PSID described above.</p>				

Table 1					
Regressions of Simulated Log of Target Wealth Ratios on Uncertainty Measures Univariate Nonparametric Distributions of Shocks to Income Estimated from PSID					
Uncertainty Specification	Constant	REPP	VARY	VARLY	Adjusted R ²
Level	-1.377 (15.49)	2.782 (13.40)			0.8772
Level	-0.564 (8.45)		1.105 (6.137)		0.5946
Level	-1.022 (33.21)			3.019 (27.67)	0.9683
Log	0.746 (8.76)	1.068 (11.65)			0.8434
Log	-0.361 (3.87)		0.448 (6.60)		0.6298
Log	0.947 (32.45)			0.8477 (40.85)	0.9852

Notes:

- 1) REPP: Relative Equivalent Precautionary Premium
VARY: Variance of Income
VARLY: Variance of Log(Income)
- 2) Each regression has 26 observations (8 occupation, 6 education, 12 industry groups)
- 3) Level (log) specifications use the level (log) of the uncertainty measures as regressors.
- 4) T-statistics are reported in parentheses.

Table 2					
Regressions of Simulated Log of Target Wealth Ratios on Uncertainty Measures Bivariate Normal Distributions of Permanent and Transitory Shocks to Income					
Uncertainty Specification	Constant	REPP	VARY	VARLY	Adjusted R ²
Level	-2.948 (144.7)	10.48 (26.50)			0.9669
Level	-2.939 (140.2)		1.105 (25.36)		0.9640
Level	-2.938 (141.1)			26.02 (25.50)	0.9644
Log	-1.157 (10.26)	0.410 (11.61)			0.8480
Log	-0.786 (5.49)		0.407 (11.71)		0.8501
Log	-0.791 (5.54)			0.406 (11.72)	0.8503

Notes:

- 1) REPP: Relative Equivalent Precautionary Premium
VARY: Variance of Income
VARLY: Variance of Log(Income)
- 2) Each regression has 25 observations, one for each possible pairing of standard deviations of permanent and transitory shocks from {0.025, 0.05, 0.075, 0.10, 0.125}.
- 3) Level (log) specifications use the level (log) of the uncertainty measures as regressors.
- 4) T-statistics are reported in parentheses.

Table 3

Instrumental Variables Regressions of Wealth on Income Uncertainty -- All Households 50 and Under

Uncertainty Wealth	Constant	Uncertainty	Income	Age	Age ² *10 ⁻³	Married	White	Female	Kids	OID Test
REPP										
VLA (896 obs)	-24.697** (3.120)	3.980* (1.837)	2.995** (0.292)	0.089 (0.103)	-0.090 (0.137)	-0.404 (0.244)	0.352 (0.218)	0.085 (0.174)	-0.233** (0.054)	0.264
NHNBW (860 obs)	-16.282** (2.838)	5.352** (1.508)	2.057** (0.255)	0.196* (0.085)	-2.400* (1.131)	0.128 (0.207)	0.268 (0.193)	0.090 (0.176)	-0.111* (0.047)	0.293
NW (874 obs)	-13.634** (2.607)	5.344** (1.480)	1.790** (0.231)	0.212* (0.089)	-2.371* (1.167)	0.587** (0.219)	0.388 (0.202)	0.071 (0.150)	-0.042 (0.050)	0.097
LVARLY										
VLA (896 obs)	-24.095** (3.051)	0.368* (0.161)	3.048** (0.304)	0.127 (0.106)	-1.420 (1.423)	-0.493* (0.238)	0.322 (0.219)	0.081 (0.180)	-0.245** (0.055)	0.188
NHNBW (860 obs)	-16.557** (2.946)	0.574** (0.166)	2.248** (0.303)	0.256** (0.091)	-3.242** (1.228)	-0.017 (0.210)	0.280 (0.193)	0.142 (0.185)	-0.130** (0.047)	0.277
NW (874 obs)	-13.363** (2.449)	0.510** (0.129)	1.922** (0.247)	0.259** (0.089)	-3.023* (1.181)	0.460* (0.210)	0.418* (0.198)	0.097 (0.152)	-0.060 (0.051)	0.053

- 1) VLA - Very Liquid Assets; NHNBW - Non-housing, Non-business Wealth; NW - Total Net Worth
REPP - Equivalent Precautionary Premium; LVARLY - Log of Variance of Log Income Shocks
- 2) The first (second) 3 equations use REPP (LVARLY) as the measure of income uncertainty.
- 3) ** indicates significance at the 1 percent level, * indicates significance at the 5 percent level
- 4) All wealth and income values are in logs.
- 5) All observations with negative reported wealth are excluded.
- 6) OID test is the p-value from the heteroskedasticity-robust test of the overidentifying restrictions.
- 7) Heteroskedasticity-robust standard errors are reported in parentheses.
- 8) The set of instrumental variables is described in the Appendix.

Table 4

Instrumental Variables Regressions of Wealth on Income Uncertainty -- All Households

Uncertainty Wealth	Constant	Uncertainty	Income	Age	Age ² *10 ⁻³	Married	White	Female	Kids	OID Test
REPP										
VLA (1056 obs)	-23.096** (2.934)	4.855* (1.933)	2.950** (0.282)	0.007 (0.070)	0.029 (0.088)	-0.367 (0.231)	0.504* (0.234)	0.043 (0.174)	-0.246** (0.051)	0.011
NHNBW (1008 obs)	-14.326** (2.608)	6.557** (1.599)	2.073** (0.241)	0.064 (0.059)	-0.053 (0.074)	0.169 (0.209)	0.274 (0.193)	0.098 (0.176)	-0.136** (0.045)	0.160
NW (1025 obs)	-13.600** (2.541)	6.838** (1.696)	1.891** (0.222)	0.148* (0.062)	-1.460 (0.078)	0.572** (0.220)	0.316 (0.192)	0.069 (0.152)	-0.069 (0.048)	0.003
LVARLY										
VLA (1056 obs)	-22.250** (2.720)	0.475** (0.172)	3.036** (0.298)	0.039 (0.071)	-0.015 (0.091)	-0.485* (0.228)	0.501* (0.238)	0.052 (0.183)	-0.255** (0.052)	0.007
NHNBW (1008 obs)	-13.305** (2.451)	0.648** (0.158)	2.197** (0.269)	0.107 (0.091)	-1.123 (0.077)	0.009 (0.201)	0.343 (0.188)	0.165 (0.180)	-0.148** (0.045)	0.129
NW (1025 obs)	-12.464** (2.262)	0.650** (0.141)	2.010** (0.237)	0.189** (0.063)	-2.016* (0.079)	0.408* (0.207)	0.390* (0.185)	0.120 (0.157)	-0.082 (0.048)	0.002

- 1) VLA - Very Liquid Assets; NHNBW - Non-housing, Non-business Wealth; NW - Total Net Worth
REPP - Equivalent Precautionary Premium; LVARLY - Log of Variance of Log Income Shocks
- 2) The first (second) 3 equations use REPP (LVARLY) as the measure of income uncertainty.
- 3) ** indicates significance at the 1 percent level, * indicates significance at the 5 percent level
- 4) All wealth and income values are in logs.
- 5) All observations with negative reported wealth are excluded.
- 6) OID test is the p-value from the heteroskedasticity-robust test of the overidentifying restrictions.
- 7) Heteroskedasticity-robust standard errors are reported in parentheses.
- 8) The set of instrumental variables is described in the Appendix.

Table 5

Simulated Reduction in Household Wealth for Fixed Relative Equivalent Precautionary Premiums			
REPP Fixed at Percentile:	Very Liquid Assets	Non-housing, Non-business Wealth	Total Net Worth
0	0.39	0.46	0.44
10	0.16	0.24	0.23
25	0.12	0.19	0.16
50	0.04	0.06	0.05


Each cell contains the percent reduction in wealth holdings were each household to face a REPP equal to the value of the REPP at the specified percentile of the predicted REPP distribution instead of its actual REPP.

Table 6			
Simulated Reductions in Household Wealth by Permanent Income Decile			
Reduction in Wealth level by decile when REPP is fixed at its minimum predicted value:			
Percentile Range	VLA	NHNBW	NW
0 - 10	0.53	0.64	0.62
10 - 20	0.46	0.55	0.58
20 - 30	0.45	0.51	0.50
30 - 40	0.37	0.47	0.47
40 - 50	0.36	0.47	0.45
50 - 60	0.37	0.42	0.42
60 - 70	0.35	0.54	0.39
70 - 80	0.37	0.44	0.38
80 - 90	0.33	0.40	0.46
90 - 100	0.42	0.42	0.39


Each cell contains the percent reduction in wealth holdings by decile were each household to face the minimum predicted REPP instead of its actual REPP.

Figure 1

Log W Target



Log W Target



Log W Target




Figure 2






Figure 3

