ELSEVIER Journal of Monetary Economics 47 (2001) 545–579

Journal of MONETARY ECONOMICS

www.elsevier.nl/locate/econbase

Unemployment insurance and precautionary saving ☆

Eric M. Engen[a], Jonathan Gruber[b,*]

[a] Board of Governors of the Federal Reserve System, Washington, DC 20551, USA

[b] Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, E52-355, Cambridge, MA 02142, USA

Received 26 June 1997; received in revised form 19 February 2001; accepted 19 February 2001

Abstract

Models of precautionary saving imply that households will hold more assets when faced with greater income uncertainty. However, previous empirical studies of income uncertainty have produced somewhat mixed support for the precautionary saving hypothesis. In this paper, we note that differences in the state-contingent income stream available to workers through the unemployment insurance (UI) program provides an excellent source of variation for testing the presence of a precautionary savings motive. Simulations of a stochastic life cycle model suggest that a UI system similar to the type currently in place in the U.S. can lead to a significant reduction in the assets accumulated by a median worker. Moreover, there is considerable variation in the UI benefit schedules for workers living in different states in the U.S., which provides an exogenous source of variation for empirically testing the precautionary saving hypothesis. We carry out this test using data on expected UI benefit replacement rates and financial assets held by households in the Survey of Income and Program

☆ The views expressed are those of the authors and do not necessarily reflect those of the Board of Governors or the other members of its staff. We are grateful to Julie Berry, Jeff Campione, and Raaj Chitaley for exceptional research assistance. We appreciate the helpful comments offered by Doug Bernheim, David Cutler, Jan Eberly, Jerry Hausman, Glenn Hubbard, Bruce Meyer, Steve Pischke, Jim Poterba, Andrew Samwick, Louise Sheiner, Jon Skinner, Martha Starr-McCluer, Doug Staiger, and seminar participants at the University of Chicago, the Federal Reserve Board, Harvard University, Michigan State University, the University of Minnesota, NBER, and the Econometric Society Meetings. Gruber acknowledges financial support from the National Institute on Aging.

*Corresponding author. Tel.: +1-617-253-8892; fax: +1-617-253-1330.

E-mail addresses: gruberj@mit.edu (J. Gruber).

PII: S0304-3932(01)00051-4

Participation. Our empirical results are consistent with the predictions of the model and suggest that reducing the UI benefit replacement rate by 50 percent would increase gross financial asset holdings by 14 percent, or \$241, for the average worker. We also find empirical evidence that this "crowd out" effect of UI on household saving is stronger for those facing higher unemployment risk and weaker for older workers, both of which are implications from our precautionary saving model.

JEL: D91; E21; H31

Keywords: Precautionary saving; Unemployment insurance; Income uncertainty

---

The concept that some household saving may be undertaken as a precaution for a "rainy day" has long been recognized in the savings literature. Recent simulation studies suggest that precautionary saving is a significant, and perhaps the most important, determinant of individual wealth accumulation. Moreover, in the 1995 Survey of Consumer Finances more households report precautionary saving as their most important motive for saving than any other reason.[1] Empirical studies of precautionary saving, however, have produced somewhat mixed conclusions. This empirical ambiguity may stem, at least in part, from the difficulty in identifying and measuring exogenous indicators of the income uncertainty facing an individual.

A key element of the uncertainty in future income for working households, and thus a potential determinant of precautionary saving, is the risk of lost wages stemming from unemployment. This uncertainty is mitigated in the U.S. by the presence of unemployment insurance (UI), which on average replaces 45 percent of a covered worker's lost earnings for up to 26 weeks after a qualifying loss of a job. A testable prediction of a precautionary saving model is that this type of income insurance should reduce households' asset accumulation. Furthermore, the extent of the income insurance available to unemployed workers varies exogenously with the benefit schedule of the UI system in their state of residence. This paper therefore uses differences in workers' expected UI benefits to provide a source of variation for testing the presence of a precautionary savings motive.

We begin by developing a model of household savings decisions which allows us to present testable implications to help guide the subsequent empirical work. In particular, we address two questions. First, given the low risk of unemployment faced by many households and the limited time of eligibility for unemployment insurance, should we expect UI to have a

[1] Kennickell et al. (1997). Similar responses were reported in the 1983, 1986, 1989, and 1992 Surveys of Consumer Finances.

noticeable effect on household wealth accumulation? Second, how will these effects vary according to observable characteristics of workers, providing us with testable implications in a cross-section of wealth data?

Our framework for addressing these questions is a stochastic life-cycle model in which risk-averse, prudent individuals face uncertainty over both their employment status and their wage, conditional on employment. Income uncertainty stemming from temporary unemployment risk is mitigated by an unemployment insurance system which is completely financed, in the aggregate, by a mandatory payroll tax. Thus, the model captures the insurance effects of a UI program but there are no direct aggregate wealth effects of the UI program; i.e., because the program is self-financing, on average, the only aggregare wealth effects are caused by the program's affect on precautionary saving.[2]

It is necessary to simulate solutions to the model in order to provide empirically testable implications because analytical solutions for asset holdings in the model are not tractable. Model simulations suggests that government provision of insurance for idiosyncratic unemployment risk that is not insurable in the private market can cause a significant decline in savings, although UI does increase the welfare of risk-averse, prudent savers. The negative effect of UI on asset accumulation is shown to increase, in percentage terms, with unemployment risk and decrease with age.

In the second part of our analysis, we test these implications of the precautionary saving model by estimating the "crowdout" effect of UI on financial assets using household-level data from the Survey of Income and Program Participation (SIPP).[3] We employ a reduced-form empirical specification for households' financial assets which controls for many potential determinants of savings behavior that might be spuriously correlated with UI benefit replacement rates, such as wages and fixed location effects. We find that there is a significant crowd-out effect of UI generosity on households' financial asset holdings: our estimates imply that reducing the benefit replacement rate of UI by 50 percent would raise the average household's financial asset-to-income ratio by 14 percent, or $241. This effect tends to be strongest, in percentage terms, for those facing the greatest risk of unemployment and for younger workers, which is consistent with the predictions of our precautionary-saving model. Overall, our empirical results support the premise that the

[2] In principle, there also could be aggregate wealth effects stemming from redistribution by the program across individuals with different saving propensities.

[3] While previous work has examined the effect of UI on unemployment durations (e.g. Meyer, 1990), layoff propensities (e.g. Topel, 1983), and consumption (Hamermesh, 1982; Hamermesh and Slesnick, 1995; Gruber, 1997), there has been no previous empirical examination, that we are aware of, of the savings effects of UI. Related historical analysis is presented by Palumbo (1992); he found that membership in trade unions and benefit societies, which provided some measure of private unemployment insurance, lowered the odds of saving for working households in Maine in 1890.

precautionary motive is an important determinant of individual savings behavior. Moreover, our results imply that the unemployment insurance system, while potentially raising welfare, diminishes household asset accumulation.

The paper proceeds as follows. In Section 1, we provide a brief overview of the related literature on precautionary savings behavior. Section 2 discusses the structure of our stochastic life-cycle simulation model and presents simulation results analyzing the effects of UI on asset accumulation in the model. Section 3 outlines our empirical procedures for testing some of the implications of the precautionary saving model. Section 4 presents the basic empirical results, while Section 5 presents some extensions. Section 6 concludes.

1. Previous literature on precautionary savings and social insurance

1.1. Simulation studies

There is a sizeable literature which has considered the role of precautionary savings.[4] A number of recent studies have simulated the role of income uncertainty in individual savings decisions, and find that precautionary saving can explain a sizeable portion, and perhaps the majority, of U.S. savings.[5]

The interaction of social insurance programs, other than UI, with precautionary savings motives has been investigated as well. Sheshinski and Weiss (1981), Abel (1985), Kotlikoff et al. (1987), and Hubbard and Judd (1987) show that, in partial equilibrium models with a functioning capital market but no annuity market, Social Security can significantly decrease precautionary saving when lifespan is uncertain, while increasing individual welfare. Hubbard et al. (1994, 1995) develop a simulation model of precautionary savings with uncertainty concerning earnings, uninsured necessary medical expenditures, and lifespan, and include a minimum level of consumption provided by a means-tested social insurance program. They show that social insurance can significantly reduce savings through two mechanisms: by mitigating the need for precautionary savings through the provision of a welfare safety net for consumption, and by effectively taxing away individual savings through means-testing to qualify for government assistance. Kotlikoff's (1988) simulation analysis of public health insurance provides similar conclusions.

There has been one previous simulation analysis of the effect of the UI program on asset accumulation. Hansen and Imrohoroglu (1992) calculate the

[4] Deaton (1992) and Browning and Lusardi (1996) provide reviews of recent developments.

[5] See Skinner (1988), Zeldes (1989a, b), Caballero (1991), Deaton (1991), Carroll (1993), Carroll and Samwick (1995a, b), Engen (1993b), Hubbard et al. (1994) and Aiyagari (1994).

potential welfare benefits of an unemployment insurance program in a dynamic general equilibrium model with infinitely lived agents whose only motivation for saving in a non-productive asset is to buffer unemployment spells. They find that welfare can improve with the introduction of UI, and saving drops to zero with the optimal benefit replacement rate and no moral hazard.[6]

While providing an interesting analysis of optimal UI benefits determination, Hansen and Imrohoroglu's model almost certainly overstates the degree to which UI crowds out savings in the U.S., for two reasons. First, since the only motivation for savings in the model is to buffer unemployment shocks, the level of asset accumulation is unrealistically low. Second, the UI program in their model misrepresents the structure of the program in the U.S. in ways which will overstate its impact on precautionary savings. In their model, a worker can receive UI benefits indefinitely, while in practice benefits are generally limited to 6 months; and the UI benefit replacement schedule in their model is a constant proportion of pre-unemployment wages, which ignores actual minimum and maximum UI benefit levels. As a result of these shortcomings, the (percentage) effect of an unemployment insurance system on savings in their model is likely to be unrealistically large. It remains an open question whether, in the context of a more realistic model, UI has important effects on savings.

1.2. Empirical studies

Despite the strong predictions of simulation models, econometric investigations to empirically assess the role of precautionary savings have reached mixed conclusions. Much of the empirical difficulties facing previous studies are related to identifying and measuring exogenous indicators of income uncertainty facing a household. Testing the precautionary saving model requires constructing a measure that both captures income uncertainty and is uncorrelated with other characteristics that may influence saving. These requirements have proved difficult to meet for most previous empirical approaches. Perhaps as a result, the previous findings of the literature are distinctly mixed.

There have been two general approaches to measuring uncertainty for the purposes of testing the precautionary motive. The first is to use direct measures of the uncertainty of an individual's income. Guiso et al. (1992), using data on Italian households, find that consumption is only slightly lower, and asset accumulation only slightly higher, for consumers reporting a greater subjective

[6] The optimal benefit replacement rate and welfare are found to vary substantially depending upon the degree of moral hazard. Other theoretical studies of optimal UI benefits include Bailey (1978), Fleming (1978), Shavell and Weiss (1979), Mortensen (1983), Easley et al. (1985), Atkeson and Lucas (1995), and Davidson and Woodbury (1995).

variance for their next year's income; but Lusardi (1993), using a similar approach, finds somewhat larger effects on assets.[7] Kazarosian (1994) finds that the variance of a household's income over the next 15 years is a positive predictor of wealth holdings for a sample of households headed by older (45–59) year old men in the National Longitudinal Survey, and Mercurio (1995) estimates that households with more variable incomes save more.

The second approach is to use a proxy for individual uncertainty, based on job characteristics or education. This was the approach followed by early attempts to find supporting evidence for the precautionary motive in Fisher (1956) and Friedman (1957), as well as Skinner (1988), who tabulated saving rates by occupation. While Fisher and Friedman find some evidence that individuals save more when in occupations assumed to have riskier income—consistent with the precautionary saving hypothesis—Skinner found that the highest risk occupations, the self-employed and sales workers, had lower rates of savings.

More recently, Carroll (1994), using U.S. data from both the Panel Study of Income Dynamics (PSID) and the CES, finds that measures of the coefficient of variation or the variance of future income uncertainty by education and job type has no statistically significant effect on current consumption, although a rise in Kimball's (1990) "equivalent precautionary premium", is found to significantly decrease current consumption. Dynan (1993) finds that the quarterly variance of households' consumption expenditures in the Consumer Expenditure Survey (CES) is not a significant predictor of the quarterly growth rate of consumption, when this variance is instrumented by education, occupation, or industry. On the other hand, Dardanoni (1991), using data on British households, found average consumption across occupation and industry groups to be significantly lower when income variance was greater. Carroll and Samwick (1995a,b) estimate a wealth model that separates the predictable and unpredictable components of income uncertainty, and they instrument the latter using the education and occupation of the household head. They find that unpredictable income uncertainty is a potentially important predictor of household wealth–income ratios.

The mixed results of these studies may be at least partially attributable to the difficulty in calculating an exogenous measure of income uncertainty. Determinants of income uncertainty such as education, occupation, and industry are all, to some extent, choice variables that reflect the same underlying tastes that drive wealth accumulation. Moreover, the most obvious

[7] In related work using this approach, Guiso et al. (1994) find that income risk significantly affects household asset portfolio choice by reducing the demand for risky assets, which is consistent with a precautionary saving model. Guiso and Jappelli (1994) provide empirical evidence in support of the proposition that people facing higher uninsurable risk buy more insurance against other risks that are insurable.

correlations of these observable characteristics with unobservable preferences (time preference and prudence) would tend to bias down empirical estimates of the magnitude of a precautionary saving effect. Similarly, actual income uncertainty or subjective assessments of risk are likely to be correlated with underlying tastes for savings. Thus, it is difficult to assess whether these mixed findings are merely a result of the potential correlation between the measure of income uncertainty and an omitted taste parameter.[8]

Our empirical approach attempts to avoid these problems with previous studies. Variations in unemployment insurance are directly correlated with the income risk facing families, and the model below demonstrates that this insurance has a quantitatively important effect on savings incentives in a precautionary framework. And, due to legislative differences across unemployment insurance regimes, unemployment insurance benefits vary across individuals for reasons exogenous to tastes for savings: two plausibly identical workers may have very different net income uncertainty because of the UI regimes in which they reside. Thus, variation in UI generosity provides a very useful source for testing the precautionary motive.

The most closely related empirical study to our own is that of Fishback and Kantor (1994), who find that the introduction of insurance for on-the-job injuries in the early 20th century lowered household savings. Our test is in the same spirit as that paper, but we use much newer data and a different social insurance program, allowing us to assess the relevance of the precautionary motive for current debates over the determination of savings.

2. Theoretical framework

In this section, we examine an unemployment insurance program in the context of a stochastic life-cycle simulation model.[9] Our goal is to investigate theoretically two questions which help guide the empirical analysis. First, under realistic assumptions about the risk of unemployment and the generosity of unemployment insurance, is UI predicted to have a significant negative effect on household asset accumulation? Second, how does this effect vary with observable characteristics of the population? Simulation analysis is necessary to help guide the subsequent empirical analysis since it is only under

[8] There are some related papers which consider the effects of different types of uncertainty (besides income variation) on savings. Starr-McCluer (1996) finds that having health insurance does not lower households' financial asset holdings, even when the endogeneity of health insurance coverage is taken into account. Skinner (1985), Hurd (1989), and Engen (1993a) have shown that household consumption growth is significantly sensitive to mortality rates, which is consistent with precautionary saving because of lifespan uncertainty and the absence of perfect annuity markets.

[9] See Engen (1993a,b) for further details on the estimation and specification of a related model and its solution and saving properties.

implausibly simple conditions that tractable analytical solutions for household wealth and the effect of UI can be derived.

Our multi-period model extends the standard life-cycle framework to include income risk arising from two sources: a positive probability of unemployment, and uncertainty about wages, conditional on employment. Individuals are forward-looking, risk-averse and prudent,[10] and therefore save as a precaution against downturns in future earnings as well as saving for retirement.[11]

2.1. Consumer preferences

An individual enters the model at age 21 and lives with certainty until age 75. In period t, an individual chooses consumption, C_t, and makes a contingency plan for future consumption, $\{C_{t+1}, \ldots, C_D\}$,[12] to maximize expected lifetime utility (ELU):

$$\mathrm{ELU}_t = E_t\left[\sum_{j=t}^{D}(1+\delta)^{t-j}U(C_j)\right], \tag{1}$$

E_t is the expectations operator conditional on information available at period t, D is the period corresponding to the known time of death, δ is the pure rate of time preference, and $U(\cdot)$ is the instantaneous utility function. The model is calibrated at a quarterly frequency to allow for unemployment spells and UI eligibility periods to be less than a year.

Several assumptions make the model tractable while still capturing important elements of household saving decisions for evaluating the effects of UI benefits. First, individuals do not have an explicit bequest motive. Second, labor supply and retirement decisions are assumed to be exogenous. Third, utility is assumed to be separable over time, and separable within a time period for consumption and leisure; in addition, private consumption and leisure are assumed to be completely separable from any government-provided good. Finally, the utility function for private consumption is assumed to exhibit constant relative risk aversion (constant intertemporal elasticity of

[10] Prudence refers to the propensity to save as a precaution against uncertainty (Kimball, 1990).

[11] A key assumption in this model is that there are no private markets for insuring riskiness in wages or unemployment prospects, inducing the individual to save as a precaution for these risk. This assumption follows the literature which motivates the existence of social insurance programs through market failures such as adverse selection or moral hazard (i.e., Diamond, 1977). There may in fact be market and non-market insurance available to individuals, such as the labor supply of spouses, support from friends, and other social insurance programs. But empirical work by Hamermesh (1982), Cochrane (1991), Nelson (1994), and Gruber (1997) suggests that these alternatives provide much less than full consumption insurance against adverse income or unemployment shocks, so that precautionary savings remains an important consumption smoothing device.

[12] Consumption, actual and planned, is restricted to be non-negative: $C_{t+j} \geqslant 0$, $\forall_j = 0, \ldots, D-t$.

substitution), which implies that risky labor income leads to precautionary saving:

$$U(C_t) = C_t^{1-\gamma}/(1-\gamma) \quad \text{if } \gamma \neq 1,$$
$$= \ln(C_t) \quad \text{if } \gamma = 1. \tag{2}$$

where γ denotes the relative risk aversion coefficient ($\gamma > 0$) and the intertemporal elasticity of substitution for consumption is equal to $1/\gamma$.[13]

2.2. Budget constraint

The choice of consumption depends on the consumer's financial wealth at age t, W_t. Wealth after receiving net labor income, Y_t^* (if employed), net UI benefits, B_t^* (if unemployed and eligible for benefits), pension income, P_t^* (if retired), and net (certain) capital income, $(W_{t-1} - C_{t-1}^*)r_t^*$, but before consuming C_t, is determined by the wealth transition equation

$$W_t = [W_{t-1} - C_{t-1}^*](1 + r_t^*) + (1 - R_t)[S_t Y_t^* + (1 - S_t)B_t^*] + R_t P_t, \tag{3}$$

where after-tax values of consumption and income are determined as:

$$C_t^* = C_t(1 + \tau^c),\ r_t^* = r_t(1 - \tau^r),\ Y_t^* = Y_t(1 - \tau^y)$$
$$- \min\{\tau^P Y_t,\ \tau^p \bar{Y}\},\ \text{and}\ B_t^* = B_t(1 - \tau^y).$$

with general taxes on consumption, interest income, and labor income represented by τ^c, τ^r, and τ^y, respectively, and a payroll tax to finance UI, τ^p. Since the consumer's lifespan is certain and a consumer is not allowed to die in debt, a constraint on the consumer's wealth is $W_{D+1} \geqslant 0$.[14]

The exogenous and certain labor force status is denoted by R_t, with $R_t = 1$ if the consumer is retired and out of the labor force and $R_t = 0$ if the individual is in the labor force. Exogenous but uncertain employment status (conditional on being in the labor force) at time t, S_t, is equal to one ($S_t = 1$) if the consumer is employed, and is equal to zero ($S_t = 0$) if unemployed. (Appendix A discusses the parameterization of the unemployment risk in more detail.)

Net UI benefits are denoted by B_t^* where B_t is gross UI benefits and τ^y is the personal income tax rate. UI benefits are considered taxable income by the

[13] The structural preference parameters for this stochastic life-cycle model are estimated using household panel data in Engen (1993a). Following this work, we assume a time preference rate $\delta = 0.05$ (annualized), and a coefficient of relative risk aversion, $\gamma = 2.5$.

[14] Nonsatiation and a certain time of death will make this constraint an equality. Note that if there is a positive probability of being unemployed in every period and being unemployed results in zero labor income, then a consumer will never want to borrow if marginal utility goes to infinity as consumption goes to zero (which is the case for constant relative risk aversion). With UI, a consumer may want to borrow against the small certain component of expected UI benefits, but borrowing against UI benefits is not allowed in the model in order to mirror actual lending practices.

personal income tax, which is true for most workers after 1982 and all workers after 1986. If unemployed ($S_t = 0$) then the worker will receive gross UI benefits B_t which are a function of their previous labor earnings: $B_t = \{\min[\bar{B}, \max(\underline{B}, \theta Y_{t-1})]\}$. The gross replacement rate for UI benefits is denoted by θ, with the maximum UI benefit level equal to $\bar{B}$ and the minimum benefit level equal to $\underline{B}$. However, if a worker remains unemployed for two periods, but the benefit duration is only one period, then $B_t = 0$.[15]

Net labor earnings, if employed, are denoted by Y_t^* with the payroll tax rate used to finance UI benefits, τ^p, and the personal income tax rate, τ^y, both assessed on labor income.[16] The payroll tax is levied only on labor income up to a limit of $\bar{Y}$, which is set to \$7000, the current maximum for the Federal Unemployment Tax. The UI budget is in balance at each point in time in the model; τ^p is determined endogenously so that tax revenues cover UI benefits. As a result, on average, there are no wealth effects from the UI system; any effects on savings behavior occur solely through reduced income risk conditional on a mean level of income.[17]

Other tax revenue is used by the government to finance both a government-provided consumption good (assumed to be separable in the consumer's utility function) and social security benefits. The individual receives a certain stream of pension income, P_t, when retired until death.[18]

[15] The UI formula replacement rate, θ, is set initially set equal to 0.50 of the previous period level of earnings so long as the benefit is above a minimum level of \$325/quarter and below a maximum level of \$3,900/quarter. The average replacement rate is then slightly below θ as more UI beneficiaries are constrained by the maximum benefit level than bolstered by the minimum. The average replacement rate in our simulations is quite similar to the average replacement rate, and is almost exactly equal to the median replacement rate, in our SIPP sample. Unemployment benefits are available for 2 quarters, consistent with the 26 week limit facing most workers in the U.S.

[16] In practice, the UI payroll tax is levied on firms, not on workers. We assume that there is full shifting in a lump sum fashion of the payroll tax amount (which is a lump sum for most workers due to the low taxable maximum). Full shifting of UI taxes on average is consistent with the evidence in Anderson and Meyer (1995).

[17] There is, however, some redistribution that occurs as a result of the fact that workers who experience unemployment see their wage growth delayed until they become reemployed, so that they end up lifetime poorer than other workers in their cohort, but this has second order effects on savings. We do understate substantially the effect of job loss on future earnings prospects. Jacobson et al. (1993) find that job loss among high tenure manufacturing workers induces significant long run earnings losses. Adding this to our model would increase the baseline level of precautionary savings, but would not change the marginal effect of UI; as a result, the percentage UI effects would be lower.

[18] This income represents both net payments from social security and a private pension; these are lump sum payments which are exogenous to the individual's own earnings history. The sum of pension and social security benefits replace 41 percent of the average of the three highest expected mean years of wages. This replacement rate is consistent with estimates in Boskin and Shoven (1987). Real pension income is assumed to decline by 1 percent, reflecting that private pensions only partially cover (on average) cost of living increases.

Conditional on being employed, gross labor earnings are $Y_t = w_t z_t h$, where w_t is the wage rate per effective labor unit in period t, z_t is an exogenous age-specific human capital adjustment factor, and the exogenous number of labor hours supplied is h. Individuals age 21–65 work h hours if employed ($R_t = 0, S_t = 1$), and they are retired ($R_t = 1$) during ages 66 to D. The log of labor earnings, $y_t = \ln(Y_t)$, is assumed to be normally distributed around a mean level of earnings. To capture the persistence of observed earnings shocks over time, the error structure for the log of labor earnings is specified as an AR(1) process. (Appendix B provides some details concerning the specification of the stochastic labor earnings process. Appendix C provides a brief discussion of the technique for solving the model.)

2.3. Simulation results

Fig. 1 shows the effect of a UI program on an (median) individual's asset-to-income ratio over the working portion of their life cycle—from age 25 to 64. The solid line represents the age profile of the asset–income ratio in the absence of a UI program.[19] The dashed line represents the (median) steady-state asset–income profile in the presence of a UI program.

It is immediately clear that, although the income insurance provided by UI does not completely replace lost wages and there is a time limit for eligibility, UI has an important negative effect on asset accumulation. The asset–income ratio is lower at all ages when UI is present. Over all ages 25–64, the median asset–income ratio falls by 24%, from 3.52 to 2.68. At the same time, by insuring the otherwise uninsurable risk of job loss, UI raises the expected utility of a risk-averse, prudent consumer by 8.9 percent.[20] Engen and Gruber (1995) provide more detail on these simulation results, and in particular illustrate that the conclusion of a sizeable effect of UI on asset holdings is robust to reasonable variation in the preference parameters.

We next consider two observable dimensions along which the effect of UI might be expected to vary. The first is variation in household saving responses to UI along the life cycle. As workers age, their savings will respond less (in percentage terms) to the presence of UI for two reasons. First, older workers will have generally been able to accumulate a large stock of precautionary

[19] The mean asset–income profile, not shown in the figure, is above the median profile, reflecting the skewed nature of wealth accumulation. However, the qualitative effects of UI on median and mean asset ratios are similar.

[20] It is important to note that this welfare gain ignores several additional effects of unemployment insurance coverage, such as additional welfare costs through increased unemployment durations and increased temporary layoffs, and additional (potential) welfare gains through improved job matches arising from subsidized search. The welfare gains reported in Table 1 reflect *only* the gains in terms of smoothing consumption. See Gruber (1997) for a further discussion of the limitations of this type of model for doing welfare analysis of UI.

Fig. 1. Unemployment insurance and asset–income ratios.

saving, thus reducing the savings effects of UI at the margin. Second, older workers will have resolved much of the inherent uncertainty in lifetime wages and employment and thus face fewer years in which they can be unemployed. The age pattern, in absolute terms, can be seen in Fig. 1. The absolute saving decline caused by a UI program rises somewhat with age. But in percentage terms, the saving offset falls dramatically.

We illustrate this more directly in Fig. 2, which shows the *percentage decline* in assets–income caused by a UI program in the baseline specification of the precautionary saving model; i.e., the percentage difference between the two asset profiles in Fig. 1.[21] Generally, the negative effect of a UI program on household asset accumulation is more than two times as large for younger workers than for older workers. This is a result of the feature that precautionary saving is a substantially larger component of the total assets of younger workers than older workers. Younger workers face more lifetime uncertainty in employment and wages and have a less urgent demand for retirement saving. Alternatively, older workers have a higher relative demand for retirement saving and less demand for precautionary saving against employment and wage uncertainty.

[21] Note that measuring the asset decline in percentage terms corresponds more closely with our empirical work discussed below because our empirical analysis uses a log specification for households' asset–income ratios.

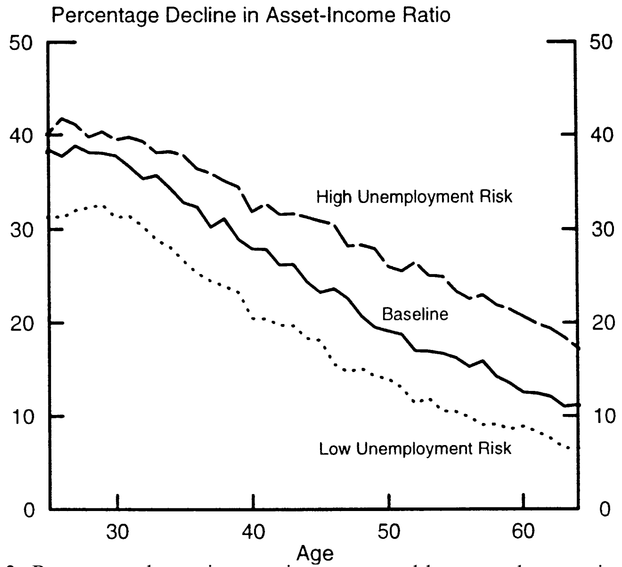

Fig. 2. Percentage change in asset–income caused by unemployment insurance.

A second potential source of variation in household saving responses to UI is the underlying level of unemployment risk. Fig. 2 also shows the impact of different levels of unemployment risk on the crowd-out effect of UI on savings. The solid line represents the percentage decline in the wealth–income profile caused by UI in our baseline scenario where the probability of unemployment (conditional on being employed) is 3 percent. If the (conditional) probability of unemployment is higher—7 percent—the percentage decline in the asset–income ratio caused by UI is larger, as shown by the long dashed line (above the solid line) in Fig. 2. Alternatively, if the (conditional) unemployment risk is lowered to 2 percent (low unemployment risk), the negative effect of UI on wealth is smaller, represented by the short-dashed line below the solid line in Fig. 2. Thus, in our empirical analysis, we should expect to find that higher unemployment risk should be positively correlated with a higher UI crowd-out effect.[22]

Of course, in our empirical work we do not observe the introduction of UI but rather variations in benefits generosity. However, simulations of the model show qualitatively similar effects of raising or lowering the replacement rate on asset accumulation, as demonstrated by Fig. 3. A 10 percent decrease in the benefit replacement rate increases the median worker's asset–income ratio by

[22]Note, however, that the model suggests that UI offers a larger welfare gain when unemployment risk is higher. See Engen and Gruber (1995) for a further discussion.

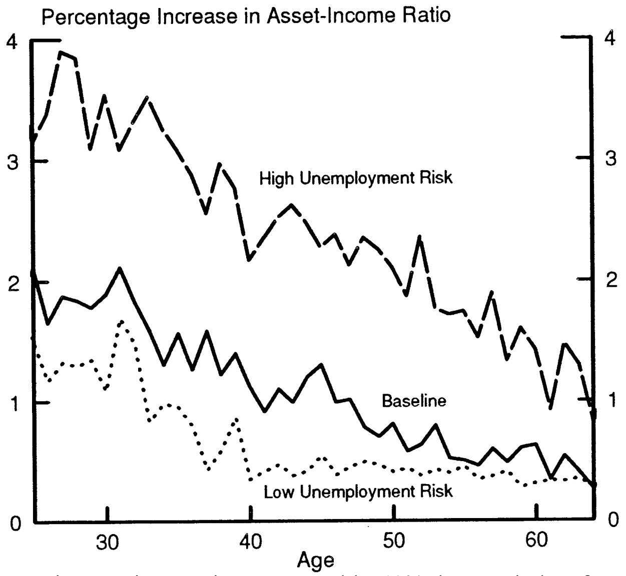

Fig. 3. Percentage increase in asset–income caused by 10% decrease in benefit replacement rate.

1.5 percent, or 4% of income.[23] Fig. 3 also shows that the magnitude of the (percentage) effect of a change in the UI benefit replacement rate on assets is greater for younger workers and those that face higher unemployment risk.

Thus, simulations of this model of precautionary saving yield two important conclusions that can guide the empirical work. First, even the limited income insurance provided by the unemployment insurance system is predicted to have an important crowd out effect on individual savings. Second, there are two clear testable cross-sectional predictions of the model: this crowd out should fall with age and rise with unemployment risk.

3. Empirical strategy

3.1. Data

In order to test these implications of the precautionary saving model, we use the largest annual data set available with information on both wealth and

[23] The effects are somewhat non-linear, in that variation in the replacement rate closer to zero has larger effects than does variation in the replacement rate at higher levels due to concavity of the utility function and the Inada conditions. But the range of observed UI benefit rates in our data, the effects are roughly linear for changes in the replacement rate.

socio-economic characteristics, the Survey of Income and Program Participation (SIPP). The SIPP is a longitudinal survey that follows a large cross section of individuals over a period of roughly 2 and 1/2 years. Each year a new panel of individuals is introduced, and these persons are interviewed for a series of waves (four month periods) over the coming months. In addition to a core set of questions about the preceding four months in each wave, individuals are asked about their wealth holdings in selected waves. For the panels in the earlier years of the SIPP (1984–1986), wealth inventories were collected twice, in the fourth and seventh waves. For 1987 and 1990, there is wealth information only in the fourth wave; there are no data for 1988 and 1989.[24] As a result, we use wealth information for a total of eight waves spanning the 1984–1990 period.

Our unit of observation in the SIPP sample is the household; since the wealth summary measures are collected only at the household level, we excluded households with more than one family in residence. We select the sample according to the following criterion: the respondent household head must be between age 25 and age 64, inclusive (in order to avoid the effects of either schooling or retirement on wealth accumulation); they must have some wage and salary earnings from a non-self employment job in each wave preceding and including the wave of the wealth survey;[25] they cannot have changed marital status in the waves preceding and including the wealth observation; and they must report values for some wealth data. The resulting sample has 24,904 observations.[26]

Our measure of wealth is gross financial assets, which consists of interest earning assets in banking and other institutions, household equity in stocks and mutual funds, and other assets such as bonds and checking accounts. We consider gross—not net—financial assets because gross assets most closely approximate the steady-state stock of accumulated precautionary savings, whereas net assets also reflect short run smoothing through consumption loans.

[24] In 1985, the first interview is actually in the third wave. For 1988 and the seventh wave of 1987 and 1990, there is some wealth data collected, but it is incomplete and is not comparable to the other waves.

[25] This selection rule allows us to measure the average replacement rate using information on each wave up to and including the wealth wave, as described below. It also avoids using individuals who may have already decumulated their wealth (in a manner correlated with UI generosity) in response to an unemployment spell. On the other hand, it does skew somewhat the representativeness of our sample. We have experimented with the sensitivity of our results by also excluding from our sample individuals who have unemployment spells within a wave (but not lasting an entire wave); this had no effect on our results, suggesting that our sample selection is not seriously biasing our inferences.

[26] Since there are two wealth interviews in 1984–1986, 5562 households have two observations for wealth. Deleting either the first or second wealth observation for these households has little effect on the results.

Furthermore, debt holding may not proxy for financial constraint but rather access to debt markets.[27] We focus only on financial wealth, and not on total household net worth, for two reasons. First, precautionary savings for unemployment should be concentrated in assets which can be liquidated quickly in the event of job loss, which may exclude categories such as housing.[28] Second, non-financial assets have some component of consumption value as well as savings, and this consumption value may respond positively, rather than negatively, to UI generosity. For example, as UI benefits rise, and individuals feel less need for precautionary savings, they may purchase more housing. Thus, we focus only on financial assets, which have no direct consumption value.

Of course, the limitation to gross financial assets has two potential costs. First, if there is some reaction of non-financial assets, or of borrowing, to UI, we will misstate the impact. Second, we limit the ability to extend our results to aggregate behavior. Thus, in our sensitivity analysis below, we will consider the effects of broadening our definition of wealth on our estimates.

Roughly one-quarter of the households in our data set have imputed wealth information, and the SIPP imputation methodology has been criticized by a number of commentators (Curtin et al., 1989; Hoynes et al., 1995). We therefore exclude imputed values for the majority of our analysis. On the other hand, the probability of imputation is not random; for example, it rises with income (and therefore with wealth). Thus, in analyzing the sensitivity of our results below, we show the results with imputed values included as well.[29]

3.2. Empirical framework

In order to model the effect of UI on wealth accumulation, we estimate regressions of the form

$$\mathrm{WEALTH}_i = \alpha + \beta_1 X_i + \beta_2 \mathrm{UI}_i + \beta_3 \delta_j + \beta_4 \tau_t + \varepsilon_{ijt}, \tag{4}$$

[27] That is, an individual with slightly negative net assets may actually have more precautionary savings than a person with slightly positive net assets, if they have better access to unsecured borrowing and therefore a lower boundary condition on net assets.

[28] We also do not include assets in IRA or Keogh accounts, nor pension assets. The former are potentially liquid, in that individuals can take loans against their assets and they only face a small penalty if they withdraw the assets early, but the ultimate perceived liquidity of these assets remains a point of debate.

[29] Curtin et al. (1989) provide an extensive review of the quality of the SIPP wealth data, relative to other household surveys and national aggregates. They find that the SIPP understates the wealth holdings of very high income households and of entrepreneurial assets, but that it is otherwise of equal quality to other household surveys. Since we focus primarily financial assets, we avoid the problem of problematic measurement of entrepreneurial holdings. And the fact that wealth is under-reported for the very wealthy should influence our estimation, since we use a robust regression technique which minimizes the influence of these observations.

where i indexes households, j indexes states, and t indexes years. WEALTH is the household financial assets normalized by average household income. The underlying role of both precautionary (and life-cycle) savings is to maintain a given living standard, so that wealth holdings should be considered relative to some measure of permanent income. This also corresponds directly to the wealth-to-income ratios presented in the previous theoretical section. We proxy for permanent income by an average of the income of the household over the waves up to and including the wealth wave.

X is a vector of demographic and economic characteristics of the household or household head/spouse: head age and age squared, sex, marital status, race, and education; a quartic in the head's weekly wage; spousal education (which is set to zero for single heads); the number of children; and an interactions of the weekly wage with head's age (to capture differential effects of earnings on savings behavior over the life cycle). UI is the replacement rate for this individual under his or her state/year UI system; once again, we use the replacement rate, rather than the benefit level, to capture the generosity of UI relative to the baseline living standard. δ_j and τ_t are vectors of state and year specific dummies.

Our key regressor is the unemployment insurance benefits for which the individual is eligible. To create this variable, we have built a simulation program which models each state's UI system for the period 1983–1991. The basis for this program is Employment and Training Administration (various years), which reports semi-annual information on state benefit schedules; in addition, it was augmented by information from a number of states and from Levine (1990). The inputs to the program are the individual's weekly wage in each wave and their number of dependents (in some states). UI benefits are then divided by this weekly earnings variable to calculate the replacement rate. Finally, the replacement rate was averaged over the waves up to and including the wealth interview. In this way, we hope to proxy for permanent differences in UI regimes which affect individuals' long run wealth accumulation decisions.

UI benefits were subsidized by the U.S. tax system in two ways over our time period: for workers in low income families (income below \$12,000 if single or \$18,000 if married), UI benefits were exempt from income taxation before 1987; and for all workers UI benefits were exempt from payroll taxes. We adjust our replacement rates to account for the second subsidy, but not the first. This is because an individual's income, and thus their income tax rate, is a function of their wealth holdings (through the returns on those assets).[30]

[30] The correlation between our post-SS tax replacement rate and the true post-income and SS tax replacement rate is 0.99, and our results are all very similar if the latter measure is used, since the tax exemption only affected low income households in the first three years of our sample. In order to qualify for UI, individuals must have a (state-specific) minimum level of earnings. To approximate this eligibility criterion, we drop any observation where weekly wages are below the minimum weekly UI benefit for that state and year.

An important potential problem with this specification is that the replacement rate for an individual may be correlated with omitted tastes for savings. In particular, due to the progressive nature of benefit schedules, replacement rates are a negative function of earnings. At the same time, higher wage workers have higher wealth-to-income ratios. Thus, if wages are not controlled for in the estimation, there would be a spurious negative correlation between the replacement rate and wealth. Even controlling for wages linearly, however, is not sufficient, since the replacement rate is a non-linear function of wages; in some states this function is quite complicated. Thus, we control for a quartic in wages, which appears to be sufficient to pick up any non-linearities in UI benefits schedules. Our results are not sensitive to the inclusion of additional higher order terms in the average weekly wage, or to alternative specifications of the wage controls (such as splines).

A final potential problem is that, even with detailed wage and individual demographic controls, there are other omitted factors correlated with both UI benefits and wealth accumulation. In particular, we are concerned that (a) there may correlated changes through time in UI benefits and savings decisions, and that (b) there may by systematic differences between the states in their UI systems which are correlated with wealth levels in the state. To control for these factors, we include fixed effects for time (SIPP panel dummies) and state.[31] While these controls do reduce the effective variation in our UI variable, their inclusion is important for ruling out alternative hypotheses for our findings based on time trends or locational taste differences.[32]

The means and distributions of the SIPP data are presented in Tables 1A and B. Table 1A presents the means and medians of the covariates included in our model. Twenty-four percent of the sample is female, and 13 percent is non-white. The median real weekly wage is \$376. The average replacement rate, after SS taxes only, is 43.6 percent of previous wages; the median is 45.6 percent.

Table 1B shows the distribution of wealth holdings in our sample, both with and without imputed values. Two facts about wealth are immediately apparent. First, the distribution of wealth is highly skewed. Second, wealth holdings are quite low for the majority of the working population. The median person in our sample only has sufficient assets to cover less than one month of family income; even the worker at the 75th quantile has financial wealth covering

[31] The SIPP does not identify all states individually, grouping Iowa, North Dakota, and South Dakota together and Alaska, Idaho, Montana, and Wyoming together; we exclude individuals in these states from the analysis.

[32] As with other studies of the effect of UI on behavior, we are unable to completely rule out the possibility that our findings are due to individual migration decisions which are correlated with UI generosity. There is no empirical evidence on this issue in the context of UI. But past research suggests that the generosity of cash welfare benefits has only a small effect on the migration decisions of potentially eligible females (Walker, 1993), for whom welfare provides a much larger total subsidy than does UI for the unemployed (since UI is time limited and welfare is not).

Table 1

(A) Means and medians of SIPP covariates[a]		
	Mean	Median
Age	40.55	39
	(10.32)	
Female	0.239	
Years of education	13.37	13
	(2.96)	
White	0.873	
Married	0.684	
Spouse's education	12.78	12
(if married)	(2.86)	
Number of kids	0.99	1
	(1.16)	
Real weekly wage	417.8	376.3
	(232.4)	
Real family earnings	29,176	26,028
	(16,643)	
After SS tax	0.436	0.456
Replacement rate	(0.135)	
Number of obs	24,904	24,904

(B) Distribution of SIPP wealth data[b]

	Without imputed values		With imputed values	
	Assets	Assets/Inc	Assets	Assets/Earn
10th %ile	0	0	0	0.001
25th %ile	276	0.012	392	0.017
50th %ile	1668	0.058	2281	0.075
75th %ile	7456	0.222	9898	0.285
90th %ile	24,406	0.661	32,174	0.816
Number of obs	24,904	24,094	33,434	33,434

[a] Based on authors' tabulations of '84–'90 SIPP data. Standard deviations in parentheses.

[b] Based on authors' tabulations of '84–'90 SIPP data. Standard deviations in parentheses. First panel shows results excluding observations where some element of wealth has been imputed by the SIPP; the second panel includes these observations.

less than 3 months of income. Part of the reason for these low wealth holdings is that we have excluded imputed values; including these values raises wealth holdings somewhat, as the second set of columns shows. But assets remain quite low even in this case, both in absolute value and as a share of earnings.

The extreme skewness of our dependent variable implies that standard ordinary least squares estimation may be inappropriate; several outlying observations could have a significant impact on our estimates even with our very large sample. We therefore use for our base specification a robust regression technique designed to minimize the influence of outliers (Berk, 1990).[33] We will also assess the sensitivity of our findings to an alternative specification that deals with this skewness, using the log of wealth.

4. Empirical results

4.1. Basic results

Our basic regression results are presented in Table 2. There is a highly significant negative relationship between UI generosity and wealth holdings. The estimated coefficient on the UI benefit replacement rate indicates that raising it by 10 percent (0.0456 percentage points at the median) lowers the asset-to-income ratio by 0.16 percentage points. At the median value of gross financial asset holdings in our sample, this is a 2.8 percent reduction. In other words, reducing the generosity of the UI benefit replacement rate by half, which represents a move from roughly the median to the 10th percentile replacement rate in our sample, would raise financial assets by 14 percent. Thus, our basic empirical results show a strong effect on wealth accumulation from the presence of the UI system, as in the previous theoretical section.[34] There is clear evidence from this finding for a precautionary saving motive.

While large in percentage terms, these effects are fairly small when expressed in terms of dollars. We find that halving the generosity of the UI system would raise savings by only 0.8 percent of income, or \$241. This small finding in dollars terms is not surprising given the low financial asset holdings of the median family.

A natural comparison is to the dollar income risk faced by a family where the head becomes unemployed. At the median real weekly wage in our sample, halving the replacement rate lowers weekly income during unemployment spells by \$85.80. The average person who becomes unemployed spends 13.1 weeks unemployed, so that this halving of benefits would lower his income over the spell by \$1124.[35] Thus, the reduction in wealth from halving benefits generosity is only 21 percent as large as the increase in UI benefits during an unemployment spell.[36]

[33] This technique first excludes influential outlying observations, then iterate towards a solution by downweighting those observations with larger residuals.

[34] We return in the conclusion to discuss the set of issues involved in comparing our empirical results to the magnitudes predicted by the theory.

[35] The unemployment duration figure was provided to us by Patty Anderson, based on tabulations of administrative wage history data.

[36] Note also that these findings are not inconsistent with those of Gruber (1997), whose significant consumption smoothing results implied that UI generosity increases were not simply crowding out other forms of consumption insurance (such as savings). Gruber's results imply that halving the replacement rate would lower consumption during unemployment spells by 13.5%, a small magnitude which is consistent with the crowdout documented here.

Table 2
Basic results[a]

Replacement rate	− 0.0362
	(0.0114)
Age	− 0.0008
	(0.0004)
Age squared	0.0017
(∗100)	(0.0005)
Age∗wage	− 0.0001
(∗1000)	(0.0002)
Wage	− 0.0089
(∗100)	(0.0029)
Wage^2	0.0048
(∗10^4)	(0.0006)
Wage^3	− 0.0048
(∗10^7)	(0.0006)
Wage^4	0.0015
(∗10^{10})	(0.0002)
Years of education	0.0031
	(0.0002)
Female	0.0053
	(0.0015)
Married	− 0.0135
	(0.0032)
White	0.0123
	(0.0015)
Spouse's education	0.0022
	(0.0002)
Number of kids	− 0.0066
	(0.0005)
Number of obs	24,904

[a] Standard errors in parentheses. Regression also includes a full set of panel dummies and state dummies.

Of course, this is only a very rough approximation to the ‘‘crowdout’’ of wealth holdings from UI. An appropriate crowdout calculation would involve comparing the reduction in wealth holdings to the expected future stream of UI payments, factoring in not only payments per unemployment spell but also the likelihood and number of future spells. This is impossible to do with our short panel of data on employment. Moreover, such a calculation would incorporate the effective redistribution of the UI system from high income workers, who are constrained by the benefits maximum and face a lower risk of layoff, to low income workers, who are bolstered by the minimum and face a higher risk.

The control variables generally have their expected signs. Savings are predicted to fall at young ages (below age 30), and rise thereafter (evaluated at the median wage level). Higher education is associated with more savings, while being a female head is associated with more savings, but even more borrowing, for a negative net asset effect. The marital dummy has a negative coefficient, but it cannot be interpreted without consideration of the spousal education variable. At the median spousal education level, there is essentially a zero effect of being married. The savings propensity rises less steeply with age for higher wage earners.

4.2. Sensitivity to specification

In this section, we assess the robustness of our conclusions to our choice of specification and wealth measure. In the first panel of Table 3, we show the results of three variations to our specification. In the second, we show the effects of varying the wealth measure that we use.

As highlighted above, the extreme skewness in the wealth data suggests the use of a regression specification that is not as sensitive to outliers such as OLS. An alternative to robust regression is simply to use the log of the asset–income ratio as the dependent variable. The potential concern with using a log specification, however, is that this requires us to exclude the 11 percent of our sample with zero wealth holdings. This may bias our conclusions through sample selection if having zero (versus non-zero) assets is a function of the generosity of the UI system. But a prerequisite for this sample selection bias to

Table 3
Sensitivity checks[a]

Vary specification	
Log of financial asset/Income ratio	− 1.496
	(0.270)
Log of financial assets	− 1.618
	(0.277)
Include age interactions	− 0.0420
	(0.0123)
Vary measure of wealth	
Net financial assets	− 0.1027
	(0.0271)
Net financial + retirement + housing equity	− 0.0620
	(0.1359)
Gross financial − include imputed values	− 0.0676
	(0.0136)

[a] Standard errors in parentheses. All regressions include the control variables listed in Table 2 and the note to that table. First panel varies the specification of Eq. (4): first two rows use a different specification of dependent variable, and third row includes interactions of wage with state of residence and year. Second panel varies the measure of wealth used, in regressions that include interactions of wage and state/time.

be important for our conclusions is that there be a relationship between being a non-zero saver and UI generosity. In fact, there is no such relationship, as revealed by a substantively and statistically insignificant coefficient on the replacement rate in a regression of a dummy for being a zero saver. This is unsurprising, since there is little reason for a partial income replacement program to affect the desire to hold no assets at all.

In the first row of Table 3, we show our basic model estimated using log wealth as the dependent variable. The coefficient on the replacement rate is once again negative and highly significant. Our estimates from this model indicate that raising the replacement rate by 10 percent (0.0456 percentage points at the median) lowers the gross asset to income ratio by 6.8 percent. This effect is somewhat larger than our estimate from the linear model, but it is once again consistent with the presence of a precautionary motive.

One potential problem with our specification is that the normalization of wealth by average family income introduces a potential dependence between the dependent variable and the UI replacement rate. Measurement error which caused reported earnings to be too high would affect both the wealth ratio (since income would be too high) and the replacement rate (since wages would be too high). We can once again address this potential problem using the log of absolute wealth, rather than the wealth/income ratio, as our dependent variable. Doing so in the next row of Table 3 yields results which are somewhat larger, but once again highly supportive of a precautionary motive.

Our specification of the wealth/income regression attempted to control for a number of omitted covariates that were correlated with both wealth accumulation and with UI generosity. There are two further potential source of omitted correlation that we have ignored, however. First, while we have controlled for fixed state effects and overall differences between high and low wage earners, part of the identification of our model is coming from *relative* differences in the savings behavior of high and low wage earners across states. That is, the *ratio* of replacement rates for high and low wage earners will differ across states according to their relative maxima/minima and the slopes of their benefits schedules. But it is possible that states with a more skewed distribution of wealth may choose more progressive benefits schedules. As a rough control for this, we can add to our model interactions of each of the fixed state effects with weekly wages; this will allow for (linear) savings differences along the wage distribution by state. Second, there may also be correlated changes through time in the progressivity of the UI system and the distribution of wealth holdings in the U.S. Thus, we also add to the model interactions of weekly wages with year dummies to capture any change in savings propensities across different types of wage earners over time.

The results of adding these controls to our specification are presented in the third row of Table 3 (once again returning to our linear regression framework). In fact, this change has relatively little effect on our basic result or standard errors. We therefore include this additional set interactions in our remaining estimated models, to allow for these alternative hypotheses for the wealth-UI correlation.

We employ three variants on our gross financial asset measure in the second panel of Table 3. First, we use net financial assets, allowing borrowing to respond to UI as well as savings. Once again, we find a highly significant negative effect of UI, consistent with a precautionary motive. The implied effect of UI is higher than with gross assets, with a 10% replacement rate increase estimated to decrease the net asset/income ratio by 0.5 percent.

Second, we broaden our asset measure further, to include retirement savings and home equity. This represents the bulk of collateralizable wealth for the households in our sample.[37] In absolute value, the effect of unemployment insurance on this broader measure of wealth lies in between the response of gross and net financial assets. In percentage terms, the effect is much smaller; this coefficient implies that a 10 percent rise in the replacement rate would lower broad asset holdings by only 0.39%. The estimate, however, is very imprecise due to the noise which is introduced into the estimation by including other assets. While this imprecision precludes strong conclusions here, the results do suggest that the effect of UI operates through financial assets only, so that the effect of the program on total wealth holdings is much smaller than is implied by Table 2. We return to this point in the conclusion.

Finally, we have also excluded from our analysis any observations with imputed information on wealth holdings, on the grounds that the relatively poor SIPP imputation procedure is unlikely to yield reasonable inferences for this population. Nevertheless, since this is not a random subsample of the population, by excluding them we may limit the applicability of our results to the entire population. In the final row of Table 3, therefore, we include the imputed observations in our regressions. The estimated coefficients are larger than those reported in Table 2. Medians wealth holdings are larger for this sample as well, however, so that in percentage terms these findings do not differ meaningfully from those in our basic regression models.

5. Extensions

5.1. Unemployment risk

An important implication of our model is that the savings effect of UI should rise with the amount of unemployment risk faced by

[37] The remaining major source of asset holdings is entrepreneurial assets. We are reticent to include these in our measure of personal wealth due to the criticisms of these SIPP data by Curtin et al. (1989).

the individual. In this section we test by proposition by running regressions of the form

$$\begin{aligned}\mathrm{WEALTH}_i = {} & \alpha + \beta_1 X_i + \beta_2 \mathrm{UI}_i + \beta_3 \mathrm{UNEM}_i + \beta_4 \mathrm{UNEM}_i^* \mathrm{UI}_i \\ & + \beta_5 \delta_j + \beta_6 \tau_t + \beta_7 w_i^* \delta_j + \beta_8 w_i^* \tau_t + \varepsilon_{ijt},\end{aligned} \tag{5}$$

where UNEM is a measure of the unemployment risk faced by worker i, and w is the weekly wage; following the discussion above, we are now including interactions of the weekly wage with the state and year dummies. If the effects of UI are stronger as unemployment risk rises, then β_2 will be smaller than the coefficient estimate in Table 2, while there will be a significant interaction term β_4, which measures the differential effect on those facing high unemployment risk.

Our measure of unemployment risk is the probability of being unemployed at a point in time in the worker's industry-by-region-by-time cell, for nine major industry categories and nine regions of the country. We have computed these using the Current Population Survey's Merged Outgoing Rotation Group file, which provides monthly information on the employment status of approximately 150,000 workers. We compute the measures for each set of four months corresponding to the SIPP interview waves for individuals in our data set, and then average them over the waves preceding and including the wealth interview.

We include in the model a full set of industry dummies, as well as our set of state dummies. We also include a full set of industry*region interactions, so that we completely control for differences across individuals that are related to their differential unemployment risk, and identify the model *solely* from the interaction of that risk with UI benefits. That is, unlike the earlier empirical precautionary savings literature, we are not identifying our model based on average differences in savings across job types, but rather by the differential effects of UI across differentially risky job categories. This interaction should be exogenous to underlying tastes for savings. We also include an interaction of unemployment risk with the weekly wage, to ensure that we are not just identifying the model from differences in savings behavior among high and low wage workers in different industries or regions.

Our findings are reported in the first panel of Table 4. In fact, we do see the hypothesized pattern. The main effect (β_2, the coefficient on the replacement rate) is insignificant, while the interaction term (β_4, the coefficient on the replacement rate times risk) is sizeable and negative.[38] That is, the crowdout effect on savings is largest for those individuals who face the

[38] Recall that the main effect on unemployment risk, β_3, is not estimable since the model includes a full set of industry*region interactions to control for underlying differences in behavior that are correlated with unemployment risk levels.

Table 4
Extensions[a]

Unemployment risk interaction	
Replacement rate	0.0123
	(0.0223)
Replacement rate * Unemployment risk	− 0.9278
	(0.3252)
Number of obs	24,651
Age interactions	
RR * 25–34 years old	− 0.0537
	(0.0137)
	[0.043]
RR * 35–44 years old	− 0.0409
	(0.0133)
	[0.069]
RR * 45–54 years old	− 0.0250
	(0.0147)
	[0.111]
RR * 55–64 years old	− 0.0141
	(0.0173)
	[0.220]
Marital status interactions	
RR * married	− 0.0316
	(0.0131)
	[0.083]
RR * single	− 0.0483
	(0.0139)
	[0.055]

[a] Standard errors in parentheses. All regressions include the control variables listed in Table 2 and the note to that table, as well as full set of interactions of wage with state and year. First panel shows coefficients of interest from estimation of Eq. (5). Second panel shows coefficients of interest from estimation of Eq. (6). Final panel estimates equation of the form of (6), but with marital status interactions rather than wage interactions. Figures in square brackets are median wealth holdings for the population subgroup to which the interaction refers.

greatest unemployment risk. The results imply that increasing the level of unemployment risk from 3% to 7% would raise the savings effect of halving the replacement rate from 0.35% of income to 1.20% of income; that is, this rise in the unemployment rate would triple the savings increase from cutting UI generosity. Thus, the underlying level of unemployment risk is an important predictor of the crowd-out effects of UI. This is consistent with the notion that our estimates reflect true precautionary savings effects.

5.2. Life cycle effects

As noted in the theory section, the life-cycle precautionary savings model yields the strong prediction that UI generosity will explain a higher fraction of savings for younger workers. We can test this supposition quite readily with our SIPP data. To do so, we run regressions of the form

$$\begin{aligned}\text{WEALTH}_{ijt} = \alpha + \beta_1 X_{ijt} + \sum_{k=0}^{K} (\underline{\gamma}\text{AGECAT}_{ijt}^{k} + \underline{\pi}\text{AGECAT}_{ijt}^{k} * \text{UI}_{ijt}) \\ + \beta_2\delta_j + \beta_3\tau_t + \beta_4\tau_t \\ + \beta_5 w_{ijt} * \delta_j + \beta_6 w_{ijt} * \tau_t + \varepsilon_{ijt},\end{aligned} \tag{6}$$

where AGECAT_i^k is a dummy for individual i being in age category k, up to K different monotonically increasing categories, and $\underline{\gamma}$ and $\underline{\pi}$ are vectors of coefficients on these dummies and their interactions with the replacement rate. We use four categories, following the simulations presented earlier: 25–34 years old; 35–44 years old; 45–54 years old; and 55–64 years old.

The results from this interacted regression are reported in the second panel of Table 4. Indeed, these coefficients decline monotonically across these age groups. The result is large and significant for the youngest age groups, but the result falls as workers age, so that by age 55–64 the effect is insignificant. At the same time, moreover, wealth levels are rising with age. In parentheses under each of the estimates, we present the median-wealth-to-income ratio for each age group. Relative to the median, the proportional effect of UI on savings is falling very sharply as individuals age, as predicted by the theory. Thus, our findings confirm the general prediction that precautionary savings for unemployment risk represents a smaller share of asset accumulation at older ages.

5.3. Spousal labor supply

Thus far, our analysis has focused solely on heads of household, and ignored working spouses. In fact, in our sample almost three-quarters of spouses work. This raises two important issues for the analysis. First, for those families with working spouses, we may mismeasure the net effect of the UI system on total family income risk by ignoring the insurance provided to secondary earners. Second, a natural source of insurance for unemployment risk is the labor supply of spouses, lowering the need for savings as insurance against this event, and reducing the crowdout effect. Indeed, Cullen and Gruber (1996) find that the labor supply of wives of unemployed husbands is strongly negatively correlated with UI benefits. As a result, we would expect increases in UI generosity to have a smaller effect on married couples than on singles.

To extend our analysis to consider spousal labor supply, we first create a family replacement rate which is a weighted average of each spouse's replacement rates, where the weight is that spouse's average earnings. We then control for a quartic in a weighted average of spousal wages, and an interaction of this weighted average wage with time and state. Finally, we allow the effect of the replacement rate to differ for married and single heads, using the same type of interactive framework as in Eq. (6).

The results of this exercise are in the final panel of Table 4. The effect of UI on wealth holdings is significant for both married and single heads. We do find, however, that the effect of unemployment insurance is larger for single heads, which is consistent with the presence of insurance through spousal labor supply for married heads. The difference is relatively small in absolute value, but it is large relative to median wealth holdings (median wealth/income ratios are shown in square brackets); the percentage effect is over twice as large for singles as for marrieds. Thus, once again, this extension is supportive of the notion of precautionary savings that is used as insurance against unemployment spells.

6. Conclusions

Low rates of savings among individuals in the U.S. has stimulated both concern among policy-makers over means of raising the savings rate and interest among academics in modelling how individuals make savings decisions. One view which has gained significant theoretical attention in recent years is that much of individual wealth accumulation is driven by precautionary motives. But empirical tests of the precautionary motive have produced quite mixed results. Our approach in this paper has been to use the exogenous variation in income risk across individuals that arises from the unemployment insurance system to test for the precautionary motive. Our results suggest that individuals do save less when UI is more generous, providing empirical confirmation that the precautionary motive is an important one in practice. The results are large in percentage terms, suggesting that a halving of the generosity of the UI system (a policy change within the range of our data) would raise individual savings by one-third. However, the findings are quite small in dollar terms, reflecting the small average asset holdings of non-elderly families in the U.S.

While our finding is consistent in direction with the predictions of our model, it is quite different in its magnitude. In percentage terms, our empirical estimates are somewhat larger than those of the model; we find that a 10 percent rise in replacement rate would lower savings by 2.8 percent, while the model predicts only a 1.5 percent effect. But this larger effect in percentage terms reflects a much smaller base of savings in the data than in the model; in absolute terms, we estimate a change in savings of 0.16 percent of income from a 10 percent replacement rate rise, as opposed to 4 percent of income in the model.

The difference in these findings arise from several sources. The first is that we only have one type of asset in the model, and it is fully liquid. In the empirical work, however, we have only modeled the effect of UI on financial assets, which represent a small share of total asset holdings; as noted above, we are unable to draw firm conclusions as to the effect of UI on broader measures of assets. Thus, our small estimate dollar effects may misrepresent the aggregate savings impact of this program. Second, we do not allow for any other forms of insurance (other than savings) in the model. In reality, there still may be partial insurance which mitigates the need for precautionary savings for the event of unemployment, through the labor supply of family members, transfers from extended family and friends, or other social insurance programs. Finally, and perhaps most importantly, we have not incorporated in the model the important underlying heterogeneity in savings behavior that we see in the data. The median worker that we model in this paper only does a small share of aggregate savings. Most savings are done by higher income individuals whose income does not depend in an important way on the generosity of the UI system. The challenge is to find a theoretical basis for these differences, rather than to build them in ad-hoc; the means-tested welfare programs in Hubbard et al. (1994, 1995) provide one rationale for the wider distribution of accumulated wealth. Nevertheless, despite these limitations in comparing the results, the basic prediction of our model is borne out by the data: more generous unemployment insurance induces a significant reduction in asset holdings.

At the same time that it lowers savings, unemployment insurance can raise welfare through completing the missing market for state-contingent insurance. Furthermore, this program has a host of other costs and benefits which must be considered in optimal program design, as discussed in Gruber (1994). The important point that this paper raises is that a critical input to such optimal benefit calculations is the potential welfare costs from reduced savings. Future research on this topic could usefully extend our findings to consider measurement of these welfare costs.

Appendix A

A.1. Unemployment risk

Employment status is assumed to follow a first order Markov process:

		Previous period $(t-1)$ Employed	Unemployed
Current period (t)		$S_{t-1}=1$	$S_{t-1}=0$
Employed	$S_t=1$	p_t^1	$(1-p_t^0)$
Unemployed	$S_t=0$	$(1-p_t^1)$	p_t^0

It is assumed that the worker knows the actual probabilities in this stationary, homogenous Markov process.[39] The baseline probability of being employed given employment in the previous period is set to $p^1 = 0.97$, and the probability of remaining unemployed the next quarter if unemployed in the current quarter is $p^0 = 0.35$. These probabilities were calculated by the authors using monthly data on employment and unemployment for 1984–1987 from the Panel Study of Income Dynamics.

A.2. Stochastic process for wages

The human capital adjustment factor in labor income, z_t, allows individuals to earn more or less (on average) than other workers at time t because of different productivity/skill levels among individuals of different ages. However, z_t, and therefore the mean level of labor earnings, is the same for all individuals within an age cohort. The specification of this human capital adjustment is based on the empirical estimation of a log earnings regression equation using panel data on labor earnings for employed heads of households and spouses (when present) that are high school graduates, using the Panel Study of Income Dynamics from 1982 to 1986 (Engen, 1993b). We find that z_t is increasing until peaking at age 46 and then decreasing until retirement at age 65. The mean age-earnings profile for consumers exhibits approximately the same profile as z_t since labor supply, h, is a fixed constant and the wage rate, w_t, grows at a constant rate, g, (reflecting productivity growth) in a steady state. The quarterly pattern of wages is set to be a piece-wise linear function of the annual wage regression.

The stochastic process for log labor earnings is expressed as:

$$y_t = \overline{y_t} + u_t,$$

where the realized log of labor earnings, y_t, is the sum of the mean of log net earnings, $\bar{y}_t$, and an error term, u_t. To capture the persistence of observed earnings shocks over time, the error structure for the log of labor earnings is specified as an AR(1) process:

$$u_t = \rho u_{t-1} + e_t,$$

where e_t is normally distributed; $e_t \sim \mathrm{N}(0, \sigma_e^2)$.

Empirical studies by Lillard and Willis (1978), MaCurdy (1982), Abowd and Card (1989, 1987), Hubbard et al. (1994), and Carroll and Samwick (1995a,b)

[39] In particular, we do not allow these parameters, or labor supply more generally, to be a function of the generosity of the UI system. There is an uncertain relationship between UI benefit generosity and the probability of becoming unemployed (Feldstein, 1978; Topel, 1983; Anderson and Meyer, 1994), but there is strong evidence that more generous benefits lengthen unemployment durations (Meyer, 1990). However, making p^1 or p^0 or labor supply choices endogenous would greatly complicate the model.

provide estimates of the stochastic process for earnings (or the covariance structure of earnings from which a stochastic process can be derived) using survey data from the Panel Study of Income Dynamics (PSID). But Duncan and Hill (1985) and Altonji and Siow (1987) find evidence that measurement error for reported earnings in the PSID is a problem for inferring this process. Reporting error should be less of a problem with administrative earnings data collected from $W-2$ tax returns. Engen (1993b) analyzes the covariance structure of earnings reported on $W-2$ statements included with tax returns, and finds that an AR(1) process best describes the stochastic structure of income shocks. Estimates for the parameters ρ and σ_e^2 are $\hat{\rho}=0.82$ and $\hat{\sigma}_e^2 = 0.05$ for joint tax returns, and $\hat{\rho}=0.92$ and $\hat{\sigma}_e^2=0.10$ for individual tax returns. For the UI policy simulations ρ is set equal to 0.85 and σ_e^2 is set equal to 0.05 (on an annual basis, or 0.003 on a quarterly basis). These are conservative values for the variance and persistence of shocks to income compared to these other studies.

A.3. Model solution

During the retirement phase of the life-cycle, Social Security benefits and asset income are nonstochastic so optimal consumption can be calculated analytically for any wealth level, W_t, at any age $t = R, \ldots, D$, by

$$C_t = \left[W_t + \sum_{i=t+1}^{D} P_i \right] \Big/ H_t.$$

where

$$H_t = (1+\tau^c) \sum_{j=t}^{D} \left[\frac{\left[(1+\delta)^{t-j} \prod_{k=t}^{j} (1+r_k^*)\right]^{1/\gamma}}{\prod_{k=t}^{j} (1+r_k^*)^{j-t}} \right].$$

During the working phase of the life-cycle, $t = 1, \ldots, R-1$, with stochastic income and constant relative risk aversion utility, the optimal consumption/saving problem cannot be solved analytically, but only described by the Euler conditions:

$$U_c(C_t) = \frac{(1+r_t^*)}{(1+\delta)} E_t U_c(C_{t+1}) \quad \forall t = 1, \ldots, (R-1),$$

where U_c represents the marginal utility of consumption. A numerical algorithm is used to solve for optimal consumption and saving. The general idea is to discretize the wealth and income state space—age and employment status are already discrete—and then use the wealth transition equation and the Euler equation to solve the consumption/saving problem recursively backwards. Values of consumption for wealth and income levels between the discrete steps are linearly interpolated. For instance, the Euler equation is solved at age $t\,(\forall t = 1, \ldots, R-1)$ by numerically searching for a value of C_t that solves the Euler equation conditional on the state variables W_t, Y_t, and S_t. The expected marginal utility of optimal next period consumption, C_{t+1}, is integrated numerically as a function of expected wealth, W_{t+1}, income, Y_{t+1}, and employment status, S_{t+1}, in period $t+1$, using the wealth transition equation and the stochastic processes for employment and earnings. Optimal values of C_t are found for different discrete levels of W_t, S_t, and Y_1, over plausible ranges, and optimal consumption is approximated for levels of wealth and income between these levels. The procedure is repeated for period $t-1$ and the algorithm continues back inductively to the initial period. The result is a numerically approximated consumption function that is a function of a four-dimensional grid of state variables; wealth, W_t, current labor income, Y_t, employment status, S_t, and age, t.

A random number generator is used to simulate (log) labor income shocks from a normal distribution, so that a realized labor income path is projected for an individual. For a given realization of earnings, an individual's corresponding lifetime consumption and saving decisions are calculated. This procedure is repeated for a large number of individuals, N, in a given cohort, who each receive different shocks from the same distribution, and thus end up with different realized lifetime earnings profiles. Reported calculations are for $N = 7000$. Calculating the sample medians, or means, for each age gives a representative age profile for consumption, wealth, earnings, and saving.

References

Abel, A.B., 1985. Precautionary saving and accidental bequests. American Economic Review 75, 77–791.

Abowd, J.M., Card, D., 1987. Intertemporal labor supply and long-term employment contracts. American Economic Review 77, 50–68.

Abowd, J.M., Card, D., 1989. On the covariance structure of earnings and hours change. Econometrica 57, 411–446.

Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. Quarterly Journal of Economics 109, 659–684.

Altonji, J., Siow, A., 1987. Testing the response of consumption to income changes with (noisy) panel data. Quarterly Journal of Economics 102, 293–328.

Anderson, P.M., Meyer, B.D., 1994. The effect of unemployment insurance taxes and benefits on layoffs using firm and individual data. NBER Working Paper No. 4960.

Anderson, P.M., Meyer, B.D., 1995. The incidence of a firm-varying payroll tax: the case of unemployment insurance. NBER Working Paper No. 5201.

Atkeson, A., Lucas Jr., R.E., 1995. Efficiency and equality in a simple model of efficient unemployment insurance. Journal of Economic Theory 66, 64–88.

Bailey, M.N., 1978. Some aspects of optimal unemployment insurance. Journal of Public Economics 10, 379–402.

Berk, R.A., 1990. A primer on robust regression. In: Fox, J., Long, J.S. (Eds.), Modern Methods of Data Analysis. Sage, Newbury Park, CA.

Boskin, M., Shoven, J., 1987. Concepts and measures of earnings replacement during retirement. Issues in Pension Economics, NBER Conference volume, University of Chicago Press.

Browning, M., Lusardi, A., 1996. Household saving: micro theories and micro facts. Mimeo, April.

Caballero, R.J., 1991. Earnings uncertainty and aggregate wealth accumulation. American Economic Review 81, 859–871.

Carroll, C., 1993. Buffer stock saving and the permanent income hypothesis. Mimeo, Board of Governors of the Federal Reserve System.

Carroll, C., 1994. How does future income affect current consumption? Quarterly Journal of Economics 109, 111–148.

Carroll, C., Samwick, A., 1995a. How important is precautionary saving? Mimeo, Board of Governors of the Federal Reserve System.

Carroll, C., Samwick, A., 1995b. The nature and magnitude of precautionary wealth. Mimeo, Board of Governors of the Federal Reserve System.

Cochrane, J.H., 1991. A simple test of consumption insurance. Journal of Political Economy 99, 957–976.

Cullen, J.B., Gruber, J., 1996. Spousal labor supply as insurance: does unemployment insurance crowd out the added worker effect? NBER Working Paper No. 5608, May.

Curtin, R.T., Juster, F.T., Morgan, J.N., 1989. Survey estimates of wealth: an assessment of quality. In: Robert, E.L., Tice, H.S. (Eds.), The Measurement of Saving, Investment, and Wealth. University of Chicago Press, Chicago.

Dardanoni, V., 1991. Precautionary savings under income uncertainty: a cross-sectional analysis. Applied Economics 23, 153–160.

Davidson, C., Woodbury, S., 1995. Optimal unemployment insurance. Mimeo.

Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59, 1221–1248.

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

Diamond, P., 1977. A framework for social security analysis. Journal of Public Economics 8, 275–298.

Duncan, G.J., Hill, D.H., 1985. An investigation of the extent and consequences of measurement error in labor-economic survey data. Journal of Labor Economics 3, 508–532.

Dynan, K.E., 1993. How prudent are consumers? Journal of Political Economy 101, 1104–1113.

Easley, D., Kiefer, N., Possen, U., 1985. An equilibrium analysis of optimal unemployment insurance and taxation. Quarterly Journal of Economics 100, 989–1010.

Employment and Training Administration, various years. Significant Provisions of State Unemployment Insurance Laws. US Department of Labor, Washington, DC.

Engen, E., 1993a. A stochastic life-cycle model with mortality risk: estimation with panel data. Mimeo, UCLA.

Engen, E., 1993b. Consumption and savings in a life-cycle model with stochastic earnings and mortality risk. Mimeo, UCLA.

Engen, E., Gruber, J., 1995. Unemployment insurance and precautionary savings. NBER working paper # 5252.

Fishback, P.V., Kantor, S.E., 1994. Insurance rationing and the origins of workers' compensation. NBER Working Paper No. 4943, December.

Fisher, M., 1956. Exploration in savings behavior. Oxford University Institute Statistics Bulletin 18, 201–277.

Fleming, J.S., 1978. Aspects of optimal unemployment insurance: search, leisure, savings, and capital market imperfections. Journal of Public Economics 10, 403–425.

Friedman, M., 1957. A Theory of the Consumption Function. Princeton University Press, Princeton.

Gruber, J., 1994. The incidence of mandated maternity benefits. American Economic Review 84 (3), 622–641.

Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. American Economic Review, to appear.
Guiso, L., Jappelli, T., Terlizzese, D., 1992. Earnings uncertainty and precautionary savings. Journal of Monetary Economics 30, 307–337.
Guiso, L., Jappelli, T., Terlizzese, D., 1994. Income risk, borrowing constraints and portfolio risk. CEPR Discussion Paper No. 888, January.
Guiso, L., Jappelli, T., 1994. Does uninsurable uncertainty increase the demand for insurance against insurable risks? Mimeo.
Hamermesh, D.S., 1982. Social insurance and consumption: an empirical inquiry. American Economic Review 72, 101–113.
Hamermesh, D.S., Slesnick, D.T., 1995. Unemployment Insurance and household welfare: microeconomic evidence 1980–93. NBER Working Paper No. 5315.
Hansen, G.D., Imrohoroglu, A., 1992. The role of unemployment insurance in an economy with liquidity constraints and moral hazard. Journal of Political Economy 100, 118–142.
Hoynes, H., Hurd, M., Chand, H., 1995. Imputed wealth, subjective survival probabilities, and social security wealth. Mimeo, NBER, May.
Hubbard, R.G., Judd, K.L., 1987. Social security and individual welfare: precautionary saving, liquidity constraints, and the payroll tax. American Economic Review 77, 630–646.
Hubbard, R.G., Skinner, J., Zeldes, S.P., 1994. The importance of precautionary motives in explaining individual and aggregate saving. Carnegie-Rochester Conference Series on Public Policy 40, 59–126.
Hubbard, R.G., Skinner, J., Zeldes, S.P., 1995. Precautionary saving and social insurance. Journal of Political Economy 103, 360–399.
Hurd, M., 1989. Mortality risk and bequests. Econometrica 57, 779–814.
Jacobson, L., LaLonde, R., Sullivan, D., 1993. Earnings losses of displaced workers. American Economic Review 83, 685–709.
Kazarosian, M., 1994. Precautionary savings—a panel study. Mimeo.
Kennickell, A.B., Starr-McCluer, M., 1997. Household savings and portfolio change: Evidence from the 1983–1989 SCF Panel. Finance and Economics Discussion Series 96-18.
Kimball, M., 1990. Precautionary saving in the small and in the large. Econometrica 58, 53–73.
Kotlikoff, L.J., 1988. Health expenditures and precautionary saving. In: Kotlikoff, L. (Ed.), What Determines Saving? MIT Press, Cambridge.
Kotlikoff, L.J., Shoven, J., Spivak, A., 1987. Annuity markets, savings, and the capital stock. In: Bodie, Z., Shoven, J., Wise, D. (Eds.), Issues in Pension Economics. University of Chicago Press, Chicago.
Levine, P.B., 1990. Unemployment insurance benefit and tax schedules, 1978–1987. Mimeo, Princeton University.
Lillard, L.A., Willis, R.J., 1978. Dynamic aspects of earnings mobility. Econometrica 46, 985–1012.
Lusardi, A., 1993. Precautionary saving and subjective earnings variance. Center for Economic Research.
MaCurdy, T.E., 1982. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. Journal of Econometrics 18, 83–114.
Mercurio, M., 1995. Saving with liquidity constraints using SIPP panel data. Mimeo, Princeton University.
Meyer, B.D., 1990. Unemployment insurance and unemployment spells. Econometrica 58, 757–782.
Mortensen, D., 1983. A welfare analysis of unemployment insurance: variations on second-best themes. Carnegie-Rochester Conference Series on Public Policy 19, 67–97.
Nelson, J., 1994. On testing for full insurance using consumer expenditure survey data. Journal of Political Economy 102, 384–394.

Palumbo, M., 1992. The effects of income uncertainty on saving and insurance purchases: evidence from a survey of Maine's workers, 1890. Mimeo.
Shavell, S., Weiss, L., 1979. The optimal payment of unemployment insurance benefits over time. Journal of Political Economy 87, 1347–1362.
Sheshinski, E., Weiss, Y., 1981. Uncertainty and optimal social security systems. Quarterly Journal of Economics 96, 189–206.
Skinner, J., 1985. Variable lifespan and the intertemporal elasticity of consumption. Review of Economics and Statistics 67, 616–623.
Skinner, J., 1988. Risky income, life cycle consumption, and precautionary saving. Journal of Monetary Economics 22, 237–255.
Starr-McCluer, M., 1996. Health insurance and precautionary saving. American Economic Review 86, 285–295.
Topel, R., 1983. On layoffs and unemployment insurance. American Economic Review 73, 541–559.
Walker, J.R., 1993. Migration among low-income households: helping the witch doctors reach consensus. Mimeo, University of Wisconsin-Madison.
Zeldes, S.P., 1989a. Optimal consumption with stochastic income: deviations form certainty equivalence. Quarterly Journal of Economics 104, 275–298.
Zeldes, S.P., 1989b. Consumption and liquidity constraints: an empirical investigation. Journal of Political Economy 97, 305–346.