THE SURVEY OF INCOME AND
PROGRAM PARTICIPATION

WHO MOONLIGHTS AND WHY? EVIDENCE FROM THE SIPP

No. 224

Jean Kimmel
W. E. Upjohn Institute for
Employment Research
Karen Smith Conway
University of New Hampshire

U. S. Department of Commerce BUREAU OF THE CENSUS

Who Moonlights and Why?

Evidence from the SIPP

Jean Kimmel
Economist
W. E. Upjohn Institute for Employment Research
300 S. Westnedge Avenue
Kalamazoo, MI 49002
(616) 385-0435

and

Karen Smith Conway
Associate Professor of Economics
Department of Economics
University of New Hampshire
Durham, NH 03824
(603) 862-3386

July 1996

The authors, who alternate the order of their names, wish to thank Stanley Sedo and Ben Craig for their helpful comments, and Ken Kline and Rebecca Jacobs for their superb research assistance.

Who Moonlights and Why? Evidence from the SIPP

Moonlighting workers represent a significant percentage of the labor force. The motives for moonlighting include the inability to work sufficient hours or earn enough money on the primary job, as well as a desire for a nonpecuniary aspect of the second job that is unavailable on the primary job. We refer to these motives as primary job constraints and heterogeneous jobs. We present extensive descriptive and econometric evidence using the SIPP, and conclude that both general motives are important.

Multiple-job holding is a significant characteristic of the labor market, with approximately 6 percent of all employed males reporting a second job in 1993 (Mishel and Bernstein, 1995 p.226). Who moonlights and why do they choose to take the second job? Moonlighters are more likely to be relatively young and higher educated than the typical non-moonlighter. And, approximately 40 percent of moonlighters report taking the second job due to economic hardship. However, moonlighting may also result from an increased need for flexibility to combine work and family or the worker's choice to pursue entrepreneurial activities while maintaining the financial stability offered by the primary job.

We can classify the different reasons for moonlighting into two types of motives, which helps us to answer both the "who" and "why" moonlighting questions. First, individuals may hold multiple jobs because an hours constraint on the primary job limits that job's earnings capacity (i.e., causing economic hardship). Second, moonlighting may arise because the labor supplied to the two jobs are not perfect substitutes. We refer to this reason as the heterogeneous jobs or job-packaging motive. For example, working on the primary job may provide the worker with the credentials to acquire a higher paying second job, such as a university psychologist testifying in a jury trial. Or, working on the second job may provide satisfaction not received from the primary job, such as a comedian who has a "regular" job by day and performs at night. In either example, the costs and benefits of both jobs are more complex than the monetary wages paid and the foregone value of leisure. When faced with such nonpecuniary benefits and costs, optimizing behavior may lead a worker to take two jobs.

Previous research on moonlighting, including Shishko and Rostker (1976), O'Connell (1979) Krishnan (1990), and Paxson and Sicherman (1996) , acknowledges that multiple motives may exist but focuses only on the constraint motive. To our knowledge, only three studies focus on the different reasons why people moonlight. Plewes and Stinson (1991) provide survey evidence from the 1989 Current Population Survey of the many distinct reasons for moonlighting reported by workers, while Lilja (1991)

and Conway and Kimmel (1994) explicitly model the joint motives for moonlighting and control for the endogeneity of primary job hours in their estimation of moonlighting hours equations.[1]

This research examines the characteristics of moonlighters and the length of their moonlighting episodes with the goal of understanding who moonlights and why. We begin by studying the personal and job-related characteristics of moonlighters and how the length of the moonlighting episode varies with these characteristics. We also study the linkage between moonlighting and poverty status. Then, we examine the factors contributing to the probability of moonlighting using a probit model. Finally, we estimate a duration model with unobserved heterogeneity to identify formally the determinants of the length of the moonlighting episode when multiple motives may exist. The descriptive analyses reveal that most moonlighters in our sample work fulltime on their primary jobs and 15 to 20 hours a week on lower paying second jobs, and, in spite of those long hours, tend to be somewhat poorer than the average worker. Yet, a significant minority earns a higher wage on their second job, a result consistent with "job-packaging" rather than the constraint motive. The probit model reveals that both pecuniary and nonpecuniary factors are important to the moonlighting choice. And, our duration model results suggest that the structural hazard increases over time and there is significant unobserved heterogeneity. Taken together, these results are consistent with the presence of multiple motives for moonlighting, with the constraint motive being the most common.

II. Moonlighting Trends and Data Description

Despite the importance of moonlighting to workers in today's economy, little comprehensive descriptive information exists in the economics literature. Three notable exceptions are Paxson and Sicherman (1996), Kimmel and Powell (1996), and Levenson (1995). However, Paxson and Sicherman rely on data from the Panel Study of Income Dynamics (PSID), which the authors themselves acknowledge

---

[1]Regets (1992), Lakhani (1994), and Mehay (1991) investigate a closely related issue, whether serving in the military reserves is moonlighting in the usual sense or is instead a case of compensated leisure.

is less than ideal because the specific survey questions make it unclear whether the jobs were held simultaneously or sequentially, and many of the survey responses to critical moonlighting questions are missing or unreported. Also, because the PSID is an annual survey, it may miss short moonlighting episodes. Kimmel and Powell (1996) focus on the comparison between Canada and the United States. And, Levenson (1995) focuses mainly on the links between moonlighting and parttime employment trends and therefore has little to say about the different motivations for moonlighting.

While only six percent of workers report holding two or more jobs at any given time in recent years, a much larger percentage moonlights at some time in their working lives. According to Paxson and Sicherman, over half of men who work continuously moonlight at some point in their working lives. Moonlighting rates peak for workers in their 30's and 40's, perhaps due to some combination of the financial burdens of raising children, purchasing a home, and saving for college. (Unpublished BLS data summarized in Kimmel, 1995) Or, the demands of family may lead workers to increase the flexibility of their work schedules through "job-packaging". Workers in this age group are also more likely to have the work experience and credentials to enjoy some of the nonwage benefits of moonlighting mentioned earlier.

The steadiness in male moonlighting rates in the past 25 years, combined with the increase in labor force participation for wives, and rising overall moonlighting rates suggest that workers are facing growing financial pressures.[2] Real wages for men with a high school education or less have fallen nearly 30 percent in the past two decades, while real home prices and rental rates have risen 20 percent and 13 percent, respectively, during the same period. Additionally, real wages for higher-educated workers have remained stagnant. (*The State of America's Children*, 1994)

Many of these broader moonlighting trends are reflected in the data set used in this paper's empirical analyses. These data are drawn from the 1984 Panel of the Survey of Income and Program Participation, a nationally-representative panel survey data set that contains detailed job information for up

---

[2]Krishnan (1990) investigates the empirical relationship between moonlighting by the primary earner and the labor supply decisions of the spouse.

to two jobs in each four-month survey period or "wave". Included in the job details are specific job start and stop dates, which are used to identify the precise starting and stop dates of each moonlighting episode.[3] We define the primary job as the one for which the individual receives the highest earnings.[4] Our subsample is restricted to prime-aged men who work continuously throughout the 3-year panel, yielding a sample of 203 moonlighting males.[5]

Two problems are encountered in undertaking a study of this kind. The first problem arises because individual and job characteristics may change over the course of the moonlighting episode. Because our data are observed in four-month intervals, a particular moonlighting episode may span several waves and the characteristics may change during the episode. For instance, the wage on the primary job may be different in the last wave of the episode than in the first. Our duration model described in the next section takes account of this problem by permitting time-varying covariates.

In our descriptive analysis, however, the solution is not so straightforward. We must choose a unit of measurement --one observation per moonlighting wave, per moonlighting episode, or per moonlighting individual. Using the wave-level observation describes the typical moonlighter more

---

[3]A potential source of error is that persons with brief periods of nonemployment during a wave may have their weeks employed overstated because of misreported or misleading job start and end dates in the SIPP's monthly records. Extensive data checks confirmed that the possibly overstated hours of work (and therefore understated imputed wages) for persons briefly without jobs is of no empirical importance to our results. We thank Theresa Devine for bringing this subtlety of the SIPP's monthly records to our attention.

[4]Redefining the PJ as the one with the highest hours worked affects only 4.1% percent of the observations.

[5]We restrict our sample to those who moonlight at least five days because shorter episodes are more likely to reflect job changes. However, the definition of job-start and job-end dates in the survey ensures that workers who moonlight repeatedly for short periods (such as on weekends) are included. We also exclude self-employed workers from our sample because their wages reflect the returns to both their labor and capital. Thus, the marginal wage received from working one more hour is nearly impossible to measure. And, we exclude individuals ages 18-25 who are in school, but retain other young people because they moonlight at such high rates and comprise a significant percentage of our sample. Because the self-employed and students tend to moonlight at relatively high rates, these exclusions, in conjunction with the requirement that individuals work all 9 waves of the SIPP, explain the relatively low moonlighting rates in our sample.

accurately at a given point in time, but it weights the analysis more heavily towards longterm moonlighters.[6] If we choose the episode or individual as our unit of measurement (the two will only differ for the people who have more than one episode), we must decide how to assign variable values for the entire episode. This measure more accurately describes the qualities of individuals who moonlight at some point during the panel, and weights all moonlighters equally, regardless of the length of their episodes. Thus, for some of the analyses it makes sense to use the wave as the unit of measurement, while for others the episode is the more logical choice. The unit of measurement is listed at the bottom of each table, and the variable values are from the first wave of the episode when the episode is the unit of measurement. Our total sample consists of 203 individuals, 261 episodes, and 586 observations/waves.

The second problem is that some individuals were moonlighting when the survey began and others were still moonlighting when the survey ended. These individuals are typically referred to as having left-censored and right-censored episodes, respectively. [The survey does not permit reported job start dates to precede the beginning the survey, making it impossible to identify the beginning of a left-censored episode.] These individuals only pose a problem when we examine the length of the moonlighting episode because this variable is measured with error and is biased downward. For instance, an individual who moonlights continuously for the entire panel is both left- and right-censored in our data and has a reported episode length of three years, the entire length of the panel. On the other hand, the individual's personal and job-related characteristics should be valid. Twenty-one percent of our observations involve a left-censored episode (affecting 54 individuals), while the same percent involve a right-censored episode. Only nine individuals are both left- and right-censored.

In section III we focus on the length of the episode by constructing and estimating a duration model. Whereas duration models are designed to deal with the problem of censoring, the presence of

---

[6]For instance, the characteristics of a person who held two jobs for three years would enter into the calculation of the sample means nine times (a three year episode would span nine four-month intervals), whereas the characteristics of a person who held two jobs for three months would enter into the calculation only once.

time-varying covariates requires that we omit left-censored episodes.[7] This selection yields a sample of 149 moonlighters, 207 episodes and 388 observations/waves. Therefore, in order for our descriptive analysis to be comparable to our duration model results, we report the results for the moonlighting sample both with and without left-censored episodes. When the focus of the descriptive analysis is on the length of the moonlighting episode, we report results both with and without left- or right-censored episodes because either type has a mismeasured episode length.

A. Who Moonlights?

Table 1 shows the means of key variables for the full sample of moonlighters, plus the subsample that omits left-censored episodes that is used in estimating the hazard functions. This table also provides means for a male comparison group (also from the 1984 SIPP panel) that does not stratify on moonlighting status. This comparison group permits us to study the differences between male moonlighters and the overall male labor force, thereby shedding additional light on the characteristics of moonlighters and the possible motivations for moonlighting.

Table 1, Part A reports the demographic characteristics of the moonlighters. Here we choose the individual as the unit of measurement because these variables do not vary much over the course of the moonlighting episode or across episodes. The typical moonlighter is 33 years old, which is about two years younger than the average worker. Also, moonlighters tend to more educated than the full sample. See that a greater percentage of the sample of moonlighters comes from the higher education categories, particularly the college-educated category.

The marital status of moonlighters is similar to the comparison group, with over 70 percent of them being married and 7 percent divorced, but they are a little more likely to be single. Moonlighters also tend to have more children (and more *young* children) than the comparison group. About 10 percent of both groups are nonwhite and very few (about 3 percent) report having a physical problem that makes it

---

[7]We discuss this issue further in Section III.

difficult to work. Finally, see that nonlabor income for the four-month period is quite similar across the two groups.

To what extent does our sample reflect the population of moonlighters in the United States? This is somewhat difficult to assess, but a comparison can be drawn between our data and the moonlighting data available from a special supplement to the 1985 Current Population Survey.[8] Compared to the CPS, our SIPP workers moonlight at a lower rate (6.1% versus 3.6%), likely resulting from the criteria used to construct our analysis sample. Specifically, our SIPP sample is comprised of workers with a strong attachment to the labor force who worked each of the nine waves in the SIPP panel. Also, we exclude workers who are self-employed or in school, two subgroups with relatively high moonlighting rates. However, 11.5% of the male workers in our SIPP sample moonlight at some point during the two and a half year panel. Two other noticeable differences between the SIPP and the CPS are the higher rate of moonlighting amongst nonwhites in our SIPP estimating sample, and the lower age. Ten percent more of our SIPP moonlighters are younger than 30 years of age. Other than these two differences, the other basic characteristics are fairly similar, with only minor differences. For example, the SIPP sample of moonlighters is comprised of slightly better educated workers than the CPS sample.

So why the choice of the SIPP over the CPS for the analyses in this paper? The CPS sample of moonlighters is much larger than the sample available in the SIPP, and so might reflect the US population more accurately. Additionally, the CPS has information available concerning the moonlighters self-reported reasons for moonlighting. Unfortunately, the advantages of the CPS stop with these two factors. The link between income levels and moonlighting cannot be addressed well with the CPS because of the extremely high rate of missing income data for moonlighters. Also, the CPS contains no information concerning the duration of moonlighting, a second focal point of this paper. And, because of the panel nature of the SIPP data set, true moonlighters can be distinguished from temporary job changers. Other advantages of the SIPP are the detailed information provided on up to two jobs (including job start and end

---

[8]These CPS numbers are drawn from unpublished analyses conducted by the authors.

dates) and the relatively short length of time (four months) covered by each interview of the survey. Both of these qualities make it possible to identify brief (as well as long) periods of moonlighting, movements into and out of jobs, and the characteristics associated with each job. Because moonlighting may be motivated by short term constraints, being able to observe short moonlighting durations is important.

Table 1, Parts B and C report the primary job and secondary job characteristics. Because these characteristics change during the moonlighting episode, we choose the wave as our unit of measurement. Recall that the primary job is the one for which the individual receives the highest earnings. The wage measures are the reported earnings for the job divided by the reported hours. The average primary job wage is nearly 50 percent higher than the average secondary job wage, at $9.05 an hour versus $6.61 an hour. However, almost 25 percent of the sample received higher hourly wages on their second jobs. This suggests that while the "university professor who consults" model of moonlighting is important, it is not the most common. Average weekly hours worked on the primary job is 40.38, while average weekly hours on the second job are 17.47. The typical moonlighter therefore works fulltime on his primary job and moonlights on a lower paying, part-time job. Also, moonlighters in general earn a lower wage and work fewer hours on the primary job than the average worker.

The most common primary job occupations for moonlighters are service work (19 percent of all moonlighters), professional/technical occupations (19 percent), production and crafts (14 percent), and managerial occupations (14 percent). The most common moonlighting occupations are again service work occupations (27 percent), sales (18 percent), and professional or technical occupations (13 percent). And, approximately 37 percent of the moonlighting episodes are in jobs for which the primary job and secondary job occupations are the same. Compared to the average worker, moonlighters are much more likely to hold primary jobs in a service occupation and are less likely to be in production, craft or repair occupation, or be a machine operator or work in transportation. This suggests that if the service sector continues to grow, as predicted, then so too may the proportion of workers who moonlight.

All of these results point to limited earnings on the primary job as the motive for moonlighting. To further investigate this, Table 2 shows the relationship between working two jobs and poverty status. Again, because household income changes during the moonlighting episode, the wave is the unit of measurement. Based on government standards for poverty status (corresponding to the same time period as the data), column 1 shows the percentage of individuals with household income at four different poverty levels: below the poverty threshold, between one and two times the poverty line, between two and three times the poverty line, and greater than three times the poverty line. Column 2 repeats this percentage using a measure of household income that excludes earnings from the moonlighting job. Column 3 repeats the same percentages for the full comparison group of male workers. However, earnings from both jobs are an endogenous outcome of utility maximizing behavior, so that workers who are "job-packaging" may appear poor if only one job is considered. In other words, job-packagers may have been able to earn more on their primary job had they not chosen to take a secondary job. To address this endogeneity, we also report these poverty figures for the large subgroup of moonlighters who work fulltime (35 hours or more per week) on their primary jobs. Most moonlighters are in this subgroup, and their poverty rates for the different categories are quite similar to the rates for the full moonlighting group.

This table provides a rough picture of the percentages of workers who change their poverty level status by taking a second job. As the table reveals, earnings from the second job have a significant impact on poverty level status. Overall, poverty rates are very low in this sample in part because it is comprised of men who work continuously on at least one job for the entire panel. Still, the percent of the sample in poverty doubles if the income from the second job is taken away. And the percentage that is below two times the poverty threshold rises from approximately 17 percent to almost 25 percent. In total, thirty-five percent change poverty level status when earnings from the second job are excluded. Note also that, with the exception of those below the poverty threshold, moonlighters as a group are somewhat poorer than the comparison group, with or without the earnings from the second job. Thus, even though they are working a second job and working significantly longer total hours, moonlighters are still somewhat poorer than the

average worker. This evidence points to the constraint motive as the most common reason for moonlighting.

An obvious weakness of any descriptive analysis is that it fails to isolate the independent effect of each variable. To remedy this, we estimate a moonlighting participation equation as a function of the characteristics listed in Table 1, Parts A and B. We omit PJ hours and poverty status because both are likely endogenous, and we must exclude all SJ characteristics as they are not observed for those who do not moonlight. We estimate a probit model, which transforms a discretely measured dependent variable into a continuous probability. Then, the resulting estimated coefficients can be transformed into derivatives to describe the effects of the independent variables on the probability of moonlighting.

The results from this exercise, estimated using the comparison group, are reported in Table 3. The results given in the table include the estimated coefficient, followed by its t-statistic in parentheses.[9] Age is inversely related to the probability of moonlighting; that is, a worker becomes less likely to moonlight as he ages. But this negative effect declines mildly as age increases, as is seen with the positive coefficient on age-squared. Having higher education levels (compared to the excluded category of less than 12 years of education) is associated with higher moonlighting rates. That is, ceteris paribus, more educated individuals are more likely to moonlight. This positive effect is strongest for those with 16 or more years of education. Having more children increases the probability of moonlighting, probably reflecting the greater budgetary needs of larger families. But, having a prechool-aged child reduces the incidence of moonlighting, implying that with young children, the value of time exceeds the need for additional income. Being divorced or never-married is not significantly related to the probability of moonlighting, and male workers who report being in fair or poor health are less likely to moonlight.

Having a primary job in a clerical or services occupation is associated with a higher probability of moonlighting. Often times, these moonlighters hold second jobs in the same occupations. Finally, even

---

[9]The second column of Table 3 shows probit results from a subsample that excludes all moonlighting episodes that started prior to the start of the SIPP panel. This subsample is what is used in the duration models in the next section. These results are very similar to the findings using the full sample.

within the regression framework, budgetary factors are important to the moonlighting decision. Having higher amounts of nonlabor income is associated with a decreased probability of moonlighting. And, earning a higher wage on the primary job is also related to a lower probability of moonlighting. But, the PJ wage elasticity of moonlighting is only -0.05.

These probit regression results imply that male workers moonlight for a variety of reasons, which can be summarized as before as primary job constraints or heterogeneous jobs. This is consistent with the descriptive evidence discussed earlier. The PJ wage is important, as is nonlabor income, both suggesting the constraints motive is important. And, after controlling those factors, other personal and job characteristics are still important. For example, the complex role of children in the moonlighting choice reveals that both budgetary as well as family-related time obligations play important roles in the moonlighting decision. This finding is consistent with the self-reported descriptive evidence available in the 1985 CPS moonlighting data.[10]

B. How Long do People Moonlight?

We believe the length of the moonlighting episode may differ in some systematic way depending upon the motive for moonlighting--individuals who moonlight due to constraints might tend to do so for shorter periods of time than those with alternative motives. We also expect the wage on the primary job to be higher than that on the second job if the constraint motive is the primary reason for moonlighting. Table 4 shows figures for the average moonlighting duration (in days), and primary and secondary job wages and weekly hours for a variety of different subgroups of the full moonlighting sample. The unit of measurement is the episode. The final column of the table shows the average moonlighting duration for a sample that excludes both left and right censored episodes, since the duration is mismeasured for these observations. However, by omitting these episodes we are likely excluding those with the longest

---

[10]This is drawn from unpublished analyses conducted by the authors.

durations. By analyzing the relationship between these characteristics and the duration of the episode, we can evaluate our hypothesis.

Divorced men have the longest moonlighting duration with and without the censored episodes (259 and 144 days), but the married group is close (244 and 127 days) and far more significant given its much larger size. Having children in the household is associated with a longer episode. Note that the group with the shortest duration, single males, also have the lowest wages on both jobs. This is probably because they are younger on average. In contrast, nonwhite workers, who also had low wages on both jobs, had moonlighting episodes that lasted longer than average.

The patterns in average duration by occupation are much more affected by excluding the censored episodes. When these episodes are included, clerical and sales workers moonlight the longest, while laborers and crafts workers exhibit the shortest moonlighting durations. However, almost half of the clerical and sales workers have censored episodes so that when those episodes are omitted, these two occupations have a shorter duration than average. None of the other occupations are affected nearly as dramatically, suggesting that the really longterm moonlighters disproportionately are sales and clerical workers. The secondary jobs with the longest moonlighting durations are again sales and clerical workers, although this conclusion is also affected by omitting the censored episodes. The long durations for sales workers are related to the prevalence of parttime jobs in this job sector.

The levels of and difference in wages over the two jobs varies a great deal by occupation as well. Workers with primary or second jobs in a professional or technical occupation tend to have higher wages and smaller PJ-SJ wage differences, as well as above average durations. This suggests that workers in these occupations are more likely to moonlight for alternative reasons. Another interesting pattern emerges from workers in a managerial occupation on either their primary or second job. When the primary job is a managerial one the wages on both jobs tend to be higher, the hours worked lower, and the episode longer than when the second job is a managerial one. This may be evidence of the wide range of managerial jobs available, and reflects that workers with PJ managerial jobs are less likely to be

moonlighting due to constraints than those with SJ managerial jobs. Finally, moonlighters with primary jobs as laborers and/or second jobs in service work tend to earn the lowest wages and work fewer total hours.

Turning to the effects of education, moonlighters at the two extremes appear the most similar in episode duration. Moonlighters with four or more years of college or those with less than a high school education have the longest moonlighting durations on average and the smallest difference in wages. [As expected, the wage *levels* are higher for more educated workers.] This suggests that moderately educated workers (high school education, maybe some college) are most likely to moonlight for short periods of time on jobs that pay much lower wages--a pattern consistent with the constraint motive. Highly educated or poorly educated workers are more likely job-packaging.

Focusing on poverty status, workers with household incomes between two and three times the poverty threshold have the longest moonlighting duration. However, this is also the group with the largest wage difference, so we would expect them to be moonlighting due to the constraint motive and therefore have a shorter duration. The lower income groups (in poverty or 1 to 2 times the poverty threshold) have shorter durations, yet have a much smaller difference in wages. Thus, when we examine the length of the episode it appears that low income workers are more likely to be moonlighting due to constraints, but if we look at the wage differences we arrive at the opposite conclusion. This apparent ambiguity may be due to the very low primary job wages received by the lowest income groups; minimum wage laws and other institutions may prevent the second job wages from being much lower.

As each of these tables have shown, there is significant variety within our sample of moonlighters. Tables 1 and 2 revealed that, on average, moonlighters receive lower wages, work longer hours and are poorer than the average worker. The average moonlighter works fulltime on his primary job and works 15 to 20 hours a week on his second job to receive an hourly wage that pays approximately 25 percent less than his primary job. The vast majority of moonlighters are not in poverty and would not be even if they did not moonlight; however, taking a second job does have a significant impact on their standard level of

living. Yet, these sample averages mask important differences and patterns across variables that are only hinted at in Table 4. For example, the most and least educated workers tend to moonlight the longest and on jobs that pay similar wages. Similarly, workers in professional or technical occupations also tend to moonlight longer and receive similar wages on both jobs. These two findings suggest that these kinds workers are more likely to moonlight for an alternative motive, such as job-packaging.

III. Constructing and Estimating a Duration Model of Moonlighting

The large differences across groups of moonlighters revealed in Tables 1 through 4 emphasize the importance of seeking further insight with a formal econometric model. Estimating hazard functions will explain more accurately the factors underlying moonlighting durations. It also allows us to isolate the effect of each variable, such as education, while dealing with the problems of censored episodes and time-varying variables. And, by estimating a baseline hazard function, we can see whether individuals are more or less likely to continue moonlighting as their episode progresses.

What does economic theory suggest about the duration of moonlighting? How long a worker chooses to moonlight is likely to depend upon his motive for moonlighting. Labor supply constraints typically are believed to be temporary, and the worker will find other avenues for adjusting to any long-lasting constraints, such as finding a new job (e.g. Altonji and Paxson 1988, Paxson and Sicherman forthcoming). Thus, if the worker moonlights in response to a constraint on the primary job then the episode may be fairly short. On the other hand, workers who are moonlighting for other reasons, such as "job-packaging," may tend to moonlight for longer periods.[11]

---

[11]The duration dependence of the hazard function, or how the probability of ending the episode changes over time, also may differ across motives. For instance, the probability that a constrained worker will quit moonlighting may increase as the duration of the moonlighting episode increases, whereas the probability for a heterogeneous jobs moonlighter could be either constant or decreasing over time. Unfortunately, the conventional ways of introducing unobserved heterogeneity and explanatory variables into the duration model do not permit either to affect the duration dependence of the hazard function. Allowing duration dependence to differ by motive is beyond the scope of the present paper, but is a topic worthy of future research efforts. We revisit this issue at the close of our empirical section.

Applying existing theories of moonlighting behavior to the duration of moonlighting suggests that the length of the moonlighting episode should depend upon the same factors that influence the decision to moonlight (as well as the number of hours to supply to the SJ) under both motives, plus any factors that help determine which motive is more important. It also suggests that different individuals will exhibit different probabilities of "leaving" the moonlighting state and that many of these differences will be unobserved. Thus, one way to explore why people moonlight is to identify the factors that are most significant in explaining the length of the moonlighting episode, as well as the duration dependence exhibited by the hazard function and the effect that unobserved heterogeneity has on the estimated structural model.

A. Theoretical Issues

Following Kiefer (1988) and Greene (1993), we define F(t) as the probability that the moonlighting episode will last no longer than t periods. The survival function, S(t), is the probability that the episode will last at least t periods and is therefore equal to 1- F(t). The hazard rate or function, $\lambda(t)$, is the probability that the episode will end at period t, given that it has lasted t periods already, and $\lambda(t) = f(t)/S(t)$, where f(t) is the probability density function associated with F(t). Duration dependence refers to how the hazard rate changes with time, which is the sign of $\partial\lambda(t)/\partial t$. Negative duration dependence means that the probability that the episode will end decreases as t increases, or that $\partial\lambda(t)/\partial t < 0$. Conversely, positive duration dependence suggests that $\partial\lambda(t)/\partial t > 0$. Specification of the hazard function dictates the type of duration dependence permitted by the model.

Unobserved heterogeneity across observations will lead to a downward bias in the estimates of $\partial\lambda(t)/\partial t$ (Kiefer 1988, pp. 671-72). To illustrate, suppose that we have two discrete groups in our sample, (1) "constrained" moonlighters and (2) "job-packaging" moonlighters. Suppose also that Group (1) has a higher, but constant, hazard rate than Group (2)'s constant hazard rate. The estimated hazard function for these two groups combined will exhibit negative duration dependence by the simple fact that over time

more members of Group (1) will quit moonlighting, leaving disproportionately more members of Group (2) in the sample, which are observations with a lower hazard rate. Thus, the hazard rate for the merged sample will decrease over time.

Multiple motives for moonlighting makes controlling for unobserved heterogeneity quite important. However, it is possible that the two main motives are not necessarily mutually exclusive and that the alternative reasons classified under "job-packaging" could have different effects on duration.[12] Thus, even if we could identify the *main* motive for moonlighting (which we cannot with our data), there could still be unobserved heterogeneity as other motives might also be playing a role in the decision. This heterogeneity may thus be better treated as a continuous random variable, rather than the discrete one suggested in the above example. However, we also present an alternative duration model in which the two motives are mutually exclusive.

We estimate our primary duration model of moonlighting behavior using the Weibull distribution for the hazard function, both with and without unobserved heterogeneity. The unobserved heterogeneity is modeled as a continuous random effect that follows a gamma distribution. The technical appendix discusses these distributions and their characteristics in more detail. Duration dependence is revealed through the parameter p and unobserved heterogeneity through the strictly non-negative parameter θ. The Weibull distribution allows for strictly positive, negative or zero duration dependence depending on whether p is greater than, less than or equal to 1.0, respectively. If $\theta = 0$ then there is no unobserved heterogeneity, whereas if it is greater than zero there is. With multiple motives for moonlighting, we expect the Weibull without heterogeneity to suggest negative or constant duration dependence ($p \leq 1$) because of the neglected heterogeneity and the Weibull with heterogeneity to reflect its relative importance ($\theta > 0$).

---

[12]For instance, a worker who is suffering from hours constraints on his PJ may decide to take a second job that he also enjoys more and perhaps shows long run potential, such as a factory worker who moonlights as a handyman, an activity he enjoys more and one that he hopes might lead to a permanent business. Likewise, moonlighting on a job because one enjoys the work may lead to different behavior than the university professor who consults model where the two jobs complement one another.

We expect the probability of ending the moonlighting episode to not only depend on the length of the episode but on other variables as well. The explanatory variables are permitted to shift the hazard function upwards or downwards over all t, but generally are not allowed to affect the duration dependence or slope of the hazard function (Greene 1993, p.721). Variables likely to influence the duration of the moonlighting episode include the wages on both jobs and other characteristics of the two jobs, such as the worker's occupation. For example, the higher the wage on the second job, given the wage on the first job, the more likely one would be to continue moonlighting. Also, as revealed in Tables 1 through 4, several other job-related and personal characteristics are potentially important. Some occupations might lend themselves to greater opportunities for job-packaging, such as the college professor example given earlier, and therefore have a lower hazard rate. Demographic variables also appear to play a role, so variables such as age, education, and family structure are included. Local economic conditions might influence the probability of being constrained on the first job, as well as the severity and permanence of those constraints, so we include the local unemployment rate and seasonal dummy variables. Finally, if one believes that all workers who moonlight are constrained on their primary job (i.e., there is only one reason for moonlighting), then the number of hours on the PJ should not only be important, but exogenous as well. Therefore, we re-estimate all models including PJ hours as an explanatory variable. If the constraint motive is paramount, then this variable may greatly reduce the explanatory power of the other job-related variables.

As mentioned in Section II, many of these variables, such as wages, income and local economic conditions, vary over the length of the moonlighting episode, necessitating the inclusion of time-varying covariates. In essence, including time-varying covariates allows the hazard function to shift upward or downward at each discrete time interval (in our case, the survey wave) as the explanatory variables change. However, permitting time-varying covariates greatly complicates the problem of left-censored episodes. Recall that a left-censored episode is one that begins prior to the start of the survey panel and a right-censored episode is one that is still underway when the survey panel ends. In both cases, the

observed duration is a lower bound on the true duration. If the explanatory variables do not vary over time, then left-censored and right-censored episodes are equivalent--both provide information on durations of at least t, and it is straightforward to deal with this censoring via maximum likelihood estimation (Greene 1993).

With time-varying covariates, however, left-censored episodes have unobserved explanatory variables. The first observation for a left-censored episode is not the first interval of the moonlighting episode, so the explanatory variables do not correspond to the episode's first interval. In fact, because we do not know when a left-censored episode began, we do not know which wave(s) of the moonlighting episode we are observing. If the explanatory variables vary over time then this problem becomes critical because we cannot match the correct variable values with the correct piece of the baseline hazard. Incorporating time-varying covariates dictates that we know the value of each regressor for each interval of the moonlighting episode, particularly the first period. This is not a problem for right-censored episodes because we observe the start of those episodes.

Therefore, in order to permit time-varying covariates, we must eliminate all left-censored episodes. We recognize that by eliminating these observations we may be excluding the people who moonlight the longest. To explore the severity of this problem we also estimate the model in two admittedly inferior ways that include the left-censored episodes: (1) allow time-varying covariates and treat the first wave as the episode's first interval, which we know is incorrect for left-censored observations, and (2) restrict the regressors to be constant over the episode and use the values from the first wave. Although both of these estimates will be biased, they may yield insight into the severity of the left-censoring problem.

In our primary estimating sample, there are 207 total episodes, with an average moonlighting duration of 175.49 days. Sixty-five percent of the moonlighting episodes last one wave or less, and 13 percent cross only two waves. Finally, 67.1 percent of the individuals have one moonlighting episode, 26.8 percent have two episodes, and 6.0 percent have three episodes. The natural log of the number of

days of the moonlighting episode duration is the dependent variable in the survival function. Variables included as explanatory variables are age, age-squared, education categories, number of children in the household, number of children under age 6, dummy variables for whether the individual is married or divorced (never married is the excluded category), the state monthly unemployment rate, the wages on both jobs, nonlabor income, seasonal dummy variables and primary and secondary job occupation dummy variables. An extra set of regressions are estimated that include the log of primary job hours as a regressor.

B. Empirical Results

Table 5 reports the key parameter estimates for the two different hazard functions. For the sake of brevity and because the results are nearly identical, we do not present the estimates for these models when the log of primary job hours is included as a regressor. These results are available upon request from the authors.

These results reveal a pattern across the different hazard specifications consistent with multiple motives for moonlighting. Although the estimated Weibull model without unobserved heterogeneity suggests positive duration dependence ($\hat{p} > 1$), the hypothesis that $p=1$ cannot be rejected. Permitting unobserved heterogeneity, found to be statistically important (θ is statistically significant), more than doubles the estimate of p such that the results now suggest statistically significant positive duration dependence. As discussed earlier, this pattern follows directly from the downward bias in $\hat{p}$ when unobserved heterogeneity is ignored.

Because unobserved heterogeneity is statistically significant, we focus on the results from the Weibull model with heterogeneity. Two variables of interest, the wages on the two jobs, never have a significant effect on the length of the duration episode.[13] Perhaps the lack of statistical significance in the

[13]Our hypotheses regarding the alternative motives for moonlighting suggest that the difference in the two wages might be what is important. Therefore, we re-estimate the model including various measures of this difference, such as the simple difference, the percentage difference and the ratio, and none are

wage variables arises because the wage is a very noisy indicator of the desirability of either job--a result that is consistent with multiple moonlighting motives. In other words, there are benefits and costs to each job that are not reflected in the wage.

Nonlabor income has a positive and statistically significant effect on the probability of ending the moonlighting episode at time t. This result has at least two possible interpretations: (1) individuals with high nonlabor incomes more likely moonlight because of constraints on the primary job, and (2) an individual who moonlights because he is constrained on the PJ may need to do so for a shorter period of time, the higher his nonlabor income. The only demographic variables that are ever statistically significant are the individual's age, being divorced, and the number of children, all of which affect the hazard rate negatively. It is rather surprising that education has no effect, given the results of the descriptive analyses. To explore this further, we also estimate a specification that includes age squared and education squared; however, none of the coefficients are statistically significant and no other results changed.

Similar to the descriptive analyses, we find the SJ occupation is important. Individuals working second jobs in farming or labor, sales, service, or in professional or technical fields moonlight for longer periods. In contrast, the occupation of the primary job appears to be unimportant. Finally, none of the variables reflecting local economic conditions are statistically significant.

The final variable of interest is the hours worked on the primary job. If all workers moonlight because they are constrained on the PJ, then hours worked on the PJ is exogenous and will be very important to the duration of the moonlighting episode. Including this variable has little effect on the other parameter estimates, except for causing some of the coefficients to become less statistically significant and a very slight increase in $\hat{\rho}$. Thus, including hours on the PJ does not substantially diminish the importance of the other variables. The coefficient is negative and statistically significant, suggesting that the more hours an individual works on his first job, the longer his moonlighting episode. This implies that individuals who work more hours on their PJ are more likely to job-package, and moonlight for longer

---

statistically significant. These results are available upon request.

periods. In addition, PJ hours may be an indication of the individual's tastes for work that are not captured by the other demographic variables. (Individuals with high PJ hours like to work more, and therefore moonlight for longer periods.) However, in order for PJ hours to be a valid regressor it must be exogenous, which is inconsistent with the "job-packaging" motive.

How sensitive are these results to the exclusion of left-censored episodes or to permitting time-varying covariates? As discussed in the previous section, permitting time-varying covariates requires eliminating left-censored episodes in order to avoid measurement error in the regressors. However, excluding left-censored episodes also may be omitting those observations with the longest durations. We explore this issue further by estimating the above specifications both with and without left-censored episodes, and both with and without time-varying covariates. (In the latter, we use the value of the variables in the first observed wave of the episode. Note that Table 5 reports the estimates from the specification without left-censored episodes and with time-varying covariates.) This exercise, while yielding biased estimates, may help reveal the separate influences of allowing time-varying covariates and eliminating left-censored episodes. Both experiments lead to very similar results, with the only notable change being a lower estimate of p. Also, by using more data (including left-censored episodes) or asking less of the model (treating the covariates as time-invariant), additional coefficients, such as the local unemployment rate, the summer seasonal dummy and the primary job wage, often become statistically significant and of the correct sign. Thus, our rather bland results in Table 5 may be due to the fact that we have eliminated some of the most interesting data and are asking a lot of the model simultaneously .

In sum, then, our estimates of the determinants of moonlighting duration are fairly robust with respect to the hazard function specified, the treatment of left-censored episodes and the presence of time-varying covariates. However, the estimated duration dependence, $\hat{p}$, is sensitive to these choices. Allowing for observed or unobserved heterogeneity or omitting left-censored episodes increases $\hat{p}$ and suggests positive duration dependence. Although the point estimates of p vary, once unobserved heterogeneity is permitted, they all are statistically significantly greater than 1.0. This suggests that

individuals are more likely to quit moonlighting as their episode lengthens, which hints at the constraint motive as the primary reason for moonlighting.

C. Two Alternative Duration Models

A problem in interpreting our estimates of the determinants of moonlighting behavior is that each variable may have two separate and potentially opposing effects on the hazard function. Each variable may influence the individual's initial reason for moonlighting (constraint versus "job-packaging"), as well as the length of time the individual chooses to moonlight. The estimated effect of nonlabor income is a good example. An individual with high nonlabor income may be less likely to moonlight because he is constrained, and therefore have a lower hazard rate. Conversely, given that the individual has chosen to moonlight, he may need to do so for a shorter period of time. The wages on each job, the local unemployment rate and primary job hours may also have these dual effects on moonlighting behavior. We examine the importance of these two effects with those four variables by estimating an alternative model that includes the value of the variable in the first wave of the episode, plus a second variable that captures the change in the variable across waves over the course of the episode. Additionally, all specifications were re-estimated including both the level and change in primary job hours. Unfortunately, this alternative model is not successful in isolating the separate effects of these key variables (with the exception of PJ hours). The most likely explanation is that most of our moonlighting episodes last less than one wave, eliminating the importance of the change variables. A sample with a more even distribution of episode durations or a shorter time interval would likely yield much more interesting results.

We also estimate a second alternative duration model, one that is more in the spirit of our original view of moonlighting as arising from two distinct motives. This model also clarifies *three* possible influences that any variable may have on moonlighting behavior. Specifically, we specify a different hazard function for each type of moonlighter and a probability that each worker will be of each type. This kind of model is discussed in Kiefer (1988) and can be written as

$$f(t) = \pi\lambda_1 p_1 t^{p_1-1} e^{-\lambda_1 t^{p_1}} + (1-\pi)\lambda_2 p_2 t^{p_2-1} e^{-\lambda_2 t^{p_2}},$$

where the subscripts 1 and 2 denote the hazard functions for constrained and job-packaging moonlighters, respectively, π is the probability of moonlighting due to constraints (and hence having the first hazard function), λ_i is the baseline hazard for each type of worker, p_i is the duration dependence for each type of worker, and t is the length of the episode. Thus, π provides an estimate of the fraction of workers who are moonlighting due to constraints. Our theoretical argument suggests that $p_1 > 1$ (constrained moonlighters exhibit positive duration dependence) and that $p_2 \leq 1$ ("job-packagers" exhibit constant or negative duration dependence).

Ideally, the baseline hazards, the probability of moonlighting due to constraints and even the duration dependencies would be specified as functions of the explanatory variables. Unfortunately, however, due to our small sample estimating that many coefficients is infeasible, so we estimate a stripped-down version of the model that does not include any explanatory variables. While an unfortunate limitation, this simple model provides an estimate of the fraction of workers who are moonlighting due to constraints. It also focuses on the behavioral patterns suggested in our theoretical argument, that the duration dependencies will be different across types of moonlighters. This simple model also provides an excellent starting point for future research in which the effects of explanatory variables can be explored -- i.e., does X affect the probability of being constrained (π), the likelihood of continuing to moonlight given the worker's motive (λ_i) or how that likelihood changes over time (p_i)? Finally, we can find comfort in the fact that few of our explanatory variables are statistically significant in our other hazard models, and that our results from this simple model are quite reasonable and consistent with our theoretical argument.

The specific parameter estimates (and t-statistics) from this model are:

π	λ_1	p_1	λ_2	p_2
0.578	.00954	1.15	0.000658	0.9096
(1.53)	(4.19)	(4.28)	(0.15)	(0.40)

Thus, the probability of ending the moonlighting episode is higher and increases over time for constrained moonlighters as opposed to job-packagers, just as predicted. The probability of moonlighting because of constraints is .578, suggesting that about 60% of our sample is moonlighting for this reason. It is unfortunate that more of our coefficients are not statistically significant, but we believe this is further evidence that a more complicated model is infeasible given our data. Nonetheless, the results from this alternative duration model provide further evidence that while the constraint motive is the most common, other motives for moonlighting exist as well.

V. Concluding Remarks

This research investigates the factors associated with moonlighting using a superior data set and presents the first duration model of moonlighting behavior (to our knowledge) to appear in the literature. Using data from the SIPP, we first examine the characteristics of our moonlighting sample and compare them to our total sample of male workers. The typical moonlighter appears to be somewhat poorer than the average worker, despite working fulltime on his primary job and part-time on a lower paying second job. This finding is due in part to the relative younger age of moonlighters. Yet, there are many exceptions to this depiction, and the length of the moonlighting episode varies with many demographic and job-related variables. The probit model of moonlighting clarifies the importance of education to the moonlighting choice, while the effect of children on the father's moonlighting decision is more complicated.

To isolate the effects of these variables, control for problems in the data, and consider multiple motives for moonlighting, we estimate a duration model of moonlighting. Our results suggest that unobserved heterogeneity is important, and that once it is controlled, the probability of ending the moonlighting episode increases over time. Our interpretation of the factors affecting the length of the moonlighting episode is clouded by the dual effects that any given variable might have--it may affect the reason for moonlighting as well as have a direct effect on the duration of the episode. We propose two

alternative models to address this problem, but our data do not appear to be rich enough to exploit them fully. However, these alternative models clarify the different influences that each variable may have on moonlighting behavior and may prove fruitful for other research endeavors. For example, AFDC recipients have been characterized frequently as belonging to two distinct groups, a group of short-term recipients who use AFDC during a period of transition (such as divorce) and a group of long-term recipients who receive AFDC as a long term source of income. In such a setting, a particular variable such as the local unemployment rate may have two (or three) effects by determining the type of AFDC recipient, as well as having a direct effect on the length of the episode.

In sum, all of our results point to the presence of multiple motives for moonlighting. However, the typically assumed motive of primary job constraints appears the most common, and taking a second job often is not enough to raise the family's income to the level of the average family.

References

Altonji, Joseph G. and Christina H. Paxson, "Labor Supply Preferences, Hours Constraints and Hours-Wage Tradeoffs," *Journal of Labor Economics*, 6, 1988, pp. 254-276.

Borjas, George J., "The Relationship Between Wages and Weekly Hours of Work: the Role of Division Bias," *Journal of Human Resources*, 15(3), 1980, pp. 409-423.

Conway, Karen Smith and Jean Kimmel, "Male Labor Supply Estimates and the Decision to Moonlight," first circulated as Institute Working Paper No. 92-09, revised June 1994.

Greene, William H. *Econometric Analysis*. New York: Macmillan Publishing Co., second edition, 1993.

Kiefer, Nicholas M., "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 26(2), June 1988, pp. 646-679.

Kimmel, Jean. "Moonlighting in the United States," *Employment Research*, (Spring 1995), Volume 2, No. 1, pp. 4-6.

Kimmel, Jean and Lisa M. Powell, "A Comparative Analysis of Moonlighting and Related Policy Issues in Canada and the United States," prepared for presentation at the June 1996 Changes in Working Time in Canada and the United States conference, May 1996.

Krishnan, Pramila, "The Economics of Moonlighting: A Double Self-Selection Model," *Review of Economics and Statistics*, 72(2), May 1990, pp. 361-67.

Lakhani, Hyder, "Civilian Moonlighting and Reenlistment Probabilities of US Army Select Reserve/National Guard Personnel," U.S. Army Research Institute for the Behavioral and Social Sciences, June 1994.

Levenson, Alec R. "Where Have All the Part-Timers Gone? Recent Trends and New Evidence on Dual-Jobs," Milken Institute for Job and Capital Formation Working Paper 95-1, January 1995.

Lilja, Reija, "The Problematic and Unproblematic Second Job," Discussion Paper #107, Labour Institute for Economic Research, Helsinki, 1991.

Mehay, Stephen, "Reserve Participation Versus Moonlighting: Are They the Same?" *Defense Economics*, 2(4), 1991, pp. 325-37.

Mishel, Lawrence and Jared Bernstein. *The State of Working America*, 1994-1995. New York: M.E. Sharp, 1995.

O'Connell, John F., "Multiple Job Holding and Marginal Tax Rates," *National Tax Journal*, 32(1), March 1979, pp. 73-76.

Paxson, Christina H. and Nachum Sicherman, "The Dynamics of Job Mobility and Dual-job Holding," *Journal of Labor Economics*, 14(3), July 1996, pp. 357-393.

Plewes, Thomas J. and John F. Stinson Jr., "The Measurement and Significance of Multiple Jobholding in the United States," *Statistical Journal of the United Nations Economic Commission for Europe*, 8(1), 1991, pp. 57-68.

Regets, Mark, "Military Reserves as Compensated Leisure: A Peculiar Case of Labor Supply," mimeo, 1991.

Shishko, Robert and Bernard Rostker, "The Economics of Multiple Job Holding," *American Economic Review*, 66(3), June 1976, pp. 298-308.

The State of America's Children, 1994. Washington, D.C.: Children's Defense Fund.

Table 1
Variable Means

A. Demographic Characteristics

Variables	Moonlighting Individuals: Full[1]	Moonlighting Individuals: Exclude Left-Censored[2]	Comparison Group[3]
# observations	203	149	1832
Age			
(% of individuals in each range)			
< 30 years of age	0.41	0.44	0.30
30-45 years of age	0.47	0.44	0.51
>45 years of age	0.12	0.12	0.18
Years of Education	13.54	13.31	13.14
Education Categories:			
< 11 years	0.10	0.13	0.15
= 12 years	0.39	0.42	0.39
13-15 years	0.18	0.17	0.20
≥ 16 years	0.33	0.28	0.26
# of kids in household	1.09	0.98	0.91
# of kids < age 6	0.34	0.32	0.26
married	0.73	0.72	0.75
divorced	0.06	0.07	0.07
single	0.21	0.21	0.18
nonwhite	0.11	0.09	0.10
sick	0.03	0.03	0.04
nonlabor income	4042.67	4253.39	3909.18

[1]Includes first wave observation from first moonlighting episode.
[2]Estimating Sample: Includes same as above, but excludes episodes that begin prior to start of panel.
[3]Comparison sample of men from same panel survey, with moonlighters and non-moonlighters; includes first observation per individual.

Table 1
(Continued)

B. Primary Job Characteristics

Variables	Moonlighting Individuals: Full[1]	Moonlighting Individuals: Exclude Left-Censored[2]	Comparison Group[3]
# observations	586	388	16488
PJ Wage	9.05	8.29	10.12
PJ Weekly Hours	40.38	40.49	43.06
PJ Occupations:			
Clerical	0.11	0.08	0.07
Farming, Labor	0.04	0.05	0.07
Managerial	0.14	0.16	0.16
Precision production, crafts or repair	0.14	0.16	0.22
Machine operators; transportation	0.11	0.12	0.18
Professional/Technical	0.19	0.14	0.15
Sales	0.07	0.06	0.09
Service	0.19	0.23	0.07

[1]Unit of measurement is the wave.
[2]Estimating sample: as above but excludes left-censored observations.
[3]Comparison sample: includes each observation for all individuals.

Table 1
(Continued)

C. Secondary Job Characteristics for Moonlighters

Variables	Full Sample[1]	Exclude Left-Censored[2]
# observations	586	388
SJ Wage	6.61	6.42
SJ Weekly Hours	17.47	19.75
Seasons:		
Winter	0.17	0.23
Spring	0.15	0.20
Summer	0.44	0.24
Fall	0.24	0.32
SJ Occupations:		
Clerical	0.06	0.06
Farming, Labor	0.12	0.12
Managerial	0.09	0.09
Precision production, crafts or repair	0.07	0.09
Machine operators; transportation	0.06	0.07
Professional/Technical	0.13	0.15
Sales	0.18	0.15
Service	0.27	0.27

[1]Unit of measurement is the wave.
[2]Unit of measurement is the above, but excludes left-censored observations.

Table 2
Poverty Status and Moonlighting*

Poverty Level	Full HH Income		HH Income less SJ Income		Comparison Group[3]
	Full[1]	FT on PJ[2]	Full[1]	FT on PJ[2]	
Sample:					
< Threshold	1%	0.4%	2.0%	1.3%	2.0%
Between 1 and 2 Times Threshold	16.7	16.2	22.9	22.4	12.2
Between 2 and 3 Times Threshold	25.0	24.5	27.7	27.3	21.5
> 3 Times Threshold	57.3	58.9	47.4	49.0	63.8

* Numbers in table are percentages and reflect percent of relevant sample in each poverty status category.
[1]Includes all moonlighting observations (586 observations).
[2]Sample as above but includes only those individuals working 35 or more hours per week on primary job (531 observations).
[3]Comparison sample of men from same panel survey, with moonlighters and non-moonlighters; includes all observations in this sample (16,488 observations).

Table 3

Moonlighting Probit Equation

	Full Sample[1]	Subsample[2]
Age	-0.059***	-0.088***
	(-2.95)	(-3.88)
Age-squared	0.0007***	0.001***
	(2.61)	(3.45)
Years Educ=12	0.285***	0.021***
	(4.06)	(2.66)
Years Educ 13 to 15	0.263***	0.225***
	(3.29)	(2.60)
Years Educ ≥ 16	0.607***	0.409***
	(7.21)	(4.36)
# children	0.166***	0.111***
	(8.19)	(4.53)
# young children	-0.149***	-0.179***
	(-2.72)	(-2.75)
Divorced	0.050	0.065
	(0.65)	(0.77)
Single	-0.081	-0.202***
	(-1.16)	(-2.52)
Nonwhite	0.146***	0.146**
	(2.44)	(2.15)
Sick	-0.246*	-0.154
	(-1.93)	(-1.15)
Nonlabor Income	-8E-6*	-6E-6
	(-1.70)	(-1.16)

Table 3
(Continued)

	Full Sample[1]	Subsample[2]
PJ Occupations:		
Clerical	0.337***	0.174*
	(4.36)	(1.85)
Farming/Labor	-0.143	-0.143
	(-1.47)	(-1.34)
Managerial	0.015	0.165*
	(0.20)	(1.94)
Machine operators; transportation	-0.113	-0.129
	(-1.59)	(-1.61)
Professional/Technical	0.123	0.070
	(1.63)	(0.78)
Sales	-0.006	-0.084
	(-0.07)	(-0.84)
Service	0.541***	0.556***
	(7.65)	(7.17)
PJ Wage	-0.030***	-0.040***
	(-5.85)	(-6.67)
Elasticity	-0.053	-0.073
Constant	-0.836**	-0.577
	(-2.22)	(-1.36)

[1]Both probit models also include dummy variables for waves 2-9.
[2]This subsample excludes left-censored observations.
*, **, *** indicates significance at the 10 percent, 5 percent, and 1 percent levels.

Table 4
Moonlighting Duration and PJ and SJ Wages and Hours

Group [# obs]	Full Moon Sample[1] DURATION (s. dev.)	PJ Wage	SJ Wage	PJ Weekly Hours	SJ Weekly Hours	Non-censored[2] DURATION (s. dev.)	[# obs]
All [261]	227.8 (268)	8.20	6.02	40.5	21.3	119.2 (135)	[162]
Married [195]	244.4 (281)	8.74	6.46	41.8	21.4	127.2 (141)	[118]
Divorced [14]	259.0 (233)	9.90	6.10	38.1	17.4	143.6 (130)	[9]
Single [52]	156.9 (217)	5.69	4.34	36.1	22.1	86.0 (111)	[35]
Marry w/kids [146]	269.1 (294)	8.58	6.44	42.3	20.8	133.3 (153)	[85]
Marry, no kids [49]	170.9 (224)	9.22	6.52	40.3	23.4	111.6 (106)	[33]
Nonwhite [32]	270.9 (278)	5.91	4.87	39.7	19.9	137.6 (131)	[17]
PJ Occupations:							
Clerical [28]	315.2 (378)	8.20	5.84	37.6	18.9	107.0 (128)	[13]
Farming, Labor [18]	141.9 (151)	5.03	3.72	34.2	19.7	104.3 (104)	[14]
Managerial [34]	215.8 (209)	11.09	8.32	42.0	20.8	129.2 (147)	[21]
Precision production, crafts or repair [40]	148.2 (178)	8.29	5.31	41.3	22.5	85.8 (80)	[30]
Machine operators, transportation [32]	206.2 (206)	6.71	3.90	42.4	21.2	118.8 (108)	[18]

Table 4
(Continued)

Group [# obs]	Full Moon Sample[1] DURATION (s. dev.)	PJ Wage	SJ Wage	PJ Weekly Hours	SJ Weekly Hours	Non-censored[2] DURATION (s. dev.)	[# ob
Professional/ Technical [45]	276.0 (332)	9.81	8.14	43.1	27.1	146.1 (186)	[26
Sales [18]	282.2 (356)	7.87	6.27	40.3	22.1	82.2 (86)	[11
Service [46]	232.7 (244)	6.80	5.23	38.9	17.0	149.5 (163)	[29
SJ Occupations:							
Clerical [17]	288.5 (338)	8.27	5.87	43.5	18.9	92.0 (121)	[8]
Farming, labor [38]	172.4 (246)	6.85	4.94	39.2	23.0	131.0 (182)	[22
Managerial [27]	207.3 (280)	9.13	6.85	44.4	28.3	111.4 (202)	[17
Precision production, crafts or repair [21]	128.7 (150)	9.45	6.74	40.4	23.6	87.2 (88)	[18
Machine operators; transportation [23]	145.6 (180)	7.86	4.92	39.8	31.7	99.4 (97)	[19
Professional/ Technical [39]	252.2 (287)	10.92	10.07	42.1	21.9	127.8 (108)	[24
Sales [35]	304.0 (306)	8.38	4.83	38.0	14.6	157.7 (160)	[21
Service [61]	260.0 (268)	6.45	4.61	39.2	16.8	120.3 (103)	[33

Table 4
(Continued)

Group [# obs]	Full Moon Sample[1] DURATION (s. dev.)	PJ Wage	SJ Wage	PJ Weekly Hours	SJ Weekly Hours	Non-censored[2] DURATION (s. dev.)	[# obs]
Education Categories							
< 12 years [25]	205.2 (193)	6.28	5.20	40.5	21.4	163.3 (179)	[18]
= 12 years [99]	212.4 (257)	7.04	4.43	40.4	21.6	105.8 (104)	[65]
13 to 15 years [50]	196.6 (226)	7.69	4.82	39.2	17.6	99.8 (105)	[30]
> 15 years [87]	269.6 (317)	10.35	8.74	41.2	23.3	132.8 (165)	[49]
Poverty Status:							
< 1 Times Poverty [4]	85.2 (82)	3.23	2.36	31.5	17.0	85.2 (82)	[4]
1-2 Times Poverty [51]	191.1 (233)	6.25	5.47	41.4	17.6	92.4 (85)	[34]
2-3 Times Poverty [72]	261.0 (295)	7.69	4.52	39.0	20.5	149.4 (191)	[39]
> 3 Times Poverty [134]	228.1 (268)	9.36	7.14	41.1	23.4	117.7 (121)	[85]
Moon.No Changes Poverty [157]	246.1 (291)	8.77	6.30	40.5	21.7	124.0 (144)	[94]
Moon.Yes Changes Poverty [104]	200.0 (229)	7.32	5.59	40.4	20.8	112.6 (122)	[68]

Table 4
(Continued)

Group [# obs]	Full Moon Sample[1] DURATION (s. dev.)	PJ Wage	SJ Wage	PJ Weekly Hours	SJ Weekly Hours	Non-censored[2] DURATION (s. dev.)	[# ot
Age Categories:							
< 30 years [103]	168.6 (212)	6.30	4.66	40.4	22.2	104.2 (114)	[75
30 to 45 years [121]	270.7 (303)	9.41	6.78	40.2	20.1	136.0 (161)	[67
> 45 years [37]	251.9 (265)	9.50	7.31	41.5	22.9	119.6 (107)	[20
> 29 years [158]	266.3 (294)	9.43	6.90	40.5	20.8	132.2 (150)	[87

Notes: Number of observations for final column showing figures for just the non-censored observations will differ from numbers shown in left column.

[1]First observation from each moonlighting episode.
[2]Sample as above, but excludes left and right censored episodes.

Table 5
Parameter Estimates for the Duration Model+
(standard errors in parentheses)

	Weibull	Weibull with Heterogeneity
Age	-0.017 (0.01)	-0.026* (0.01)
Years of Education	0.048 (0.05)	0.037 (0.04)
# children	-0.186* (0.10)	-0.089 (0.10)
# young children	0.338 (0.26)	-0.008 (0.24)
Married	-0.188 (0.37)	0.026 (0.36)
Divorced	-0.869** (0.44)	-0.183 (0.43)
Sick	-0.088 (0.60)	-0.272 (0.68)
Nonlabor Income (in 1000s)	0.029 (0.02)	0.037*** (0.02)
Unemployment Rate	0.002 (0.06)	0.019 (0.05)
PJ Wage	-0.015 (0.02)	-0.007 (0.02)
SJ Wage	-0.019 (0.02)	-0.002 (0.02)
Winter	0.003 (0.28)	0.220 (0.25)
Spring	0.055 (0.27)	0.061 (0.25)
Summer	0.072 (0.24)	0.124 (0.24)
Constant	-5.003*** (1.02)	-4.035*** (0.99)
Sigma	0.862*** (0.07)	0.364*** (0.06)
Theta	---	2.541*** (0.75)
P	1.160*** (0.10)	2.746*** (0.44)
Log-Likelihood	-1033.2	-1012.9

+Estimated with time-varying covariates and excludes left censored.
*, **, *** Indicates significance at the 10 percent, 5 percent, and 1 percent levels.

Technical Appendix[1]

The different distributions and corresponding characteristics of the hazard models estimated in Section III are discussed in more detail here. Our primary distribution is the Weibull, which can be written as

$$\lambda(t) = \lambda p (\lambda t)^{p-1}. \tag{1}$$

The Weibull distribution allows for strictly positive, negative or zero duration dependence depending on whether p is greater than, less than or equal to 1.0, respectively. The Log-logistic distribution, with hazard

$$\lambda(t) = \frac{\lambda p (\lambda t)^{p-1}}{[1+(\lambda t)^p]}, \tag{2}$$

allows for first positive and then negative duration dependence as t increases, as long as $p > 1.0$. If $p \leq 1.0$, then the hazard rate always decreases with time.

We choose a common specification for the unobserved heterogeneity; it is modeled as a continuous random effect, v, with a probability density function,

$$f(v) = \frac{k^k}{\Gamma(k)} \quad e^{-kv} v^{k-1}, \tag{3}$$

which is a gamma distribution with a mean of 1.0 and variance of 1/k. The survival function conditional upon v, or S(t |v), is specified as a Weibull distribution,

$$S(t|v) = e^{-(v\lambda t)^p}. \tag{4}$$

The mixed hazard function can then be derived as

$$\lambda(t) = \lambda p (\lambda t)^{p-1} [S(t)]^{1/k}, \tag{5}$$

where

$$S(t) = \int_0^\infty f(v) S(t|v) \partial v = \left[1 + \frac{1}{k}(\lambda t)^p\right]^k. \tag{6}$$

The parameter k, or the more commonly discussed parameter θ, where $\theta = 1/k$, indicates the degree of unobserved heterogeneity in the sample. If $\theta = 0$ then there is no heterogeneity and the hazard written in equation (5) collapses into the simple Weibull hazard written in equation (1). If $\theta > 0$ then unobserved heterogeneity exists and the mixed and structural hazards differ.[2] In particular, if $p \leq 1$, the mixed hazard exhibits negative duration dependence, whereas the structural hazard only exhibits negative duration dependence if $p < 1$. (It is constant if $p=1$.) If $p > 1$, then the structural hazard exhibits positive duration dependence and the mixed hazard is ambiguous.

---

[1]This summary of the hazard function model is drawn from Greene (1993).

[2]Recall that $\theta=1/k$ is the variance of the random variable, v, and therefore must be strictly non-negative.

Unemployment Insurance and Precautionary Saving

Eric M. Engen
Board of Governors of the Federal Reserve System
20th and Constitution Avenues
Mail Stop 83
Washington, D.C. 20551
(202)452-2980
eengen@frb.gov

Jonathan Gruber
Department of Economics
MIT
50 Memorial Drive
Cambridge, MA 02139
(617)253-8892
gruberj@mit.edu

April 1996

Models of precautionary saving imply that households will tend to hold more assets when faced with greater income uncertainty. We test this hypothesis by considering both theoretically and empirically the effect of unemployment insurance (UI) on precautionary savings behavior. Simulations of a stochastic life cycle model suggest that increasing the UI benefit replacement rate for wages lost during unemployment will significantly lower the asset holdings of the median worker, and that this effect increases with unemployment risk and falls with a worker's age. We test these implications by matching data on expected UI benefit replacement rates to asset holdings of households in the Survey of Income and Program Participation (SIPP). Our empirical results are consistent with the predictions of the model. We find that raising the replacement rate for UI by 10 percent lowers financial asset holdings by 0.7% to 2.8%, which amounts to $50-$147 for the average worker. We also find evidence that the crowdout effect of UI on household saving is stronger for those facing higher unemployment risk and weaker for older workers.

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 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.

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 suggests that precautionary saving is a significant, and perhaps the most important, determinant of individual wealth accumulation. And in the 1992 Survey of Consumer Finances more households report precautionary saving as an important motive for their saving than any other reason.[1] Empirical studies of precautionary saving, however, have often produced more 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 job loss. 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' wealth accumulation. Furthermore, the extent of the income insurance available to workers who are unemployed varies exogenously with the benefit schedule of the UI system in their state of residence. Thus, differences in workers' expected UI benefits provides a natural source of variation for testing the presence of the precautionary savings motive. The purpose of this paper is to carry out such a test.

We begin by developing a realistic model of household savings decisions which allows us to derive 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 time limited nature of unemployment insurance, should we expect UI to have significant effects on household

[1]Kennickell and McCluer (1994). Similar responses were reported in the 1983 and 1989 Surveys of Consumer Finances.

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 financed (in a balanced-budget manner) by a mandatory payroll tax. Thus, the model captures the insurance effects of UI but their are no direct wealth effects of the UI program.

Simulations of the model show that government provision of insurance for idiosyncratic unemployment risk that is not insurable in the private market can cause a small but significant decline in savings: for the median household, raising the UI benefit replacement rate by 10 percent lowers median wealth-income ratios by about 1.1 percent. This negative effect of UI benefits on asset accumulation is shown to increase with unemployment risk and decrease with age. However, UI does increase the welfare of risk-averse, prudent savers even as it reduces asset accumulation.

In the second part of our analysis, we test the implications of the precautionary saving model by estimating the crowdout effects of UI on wealth using household-level data from the Survey of Income and Program Participation (SIPP).[2] We employ a reduced-form empirical specification for households' wealth 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 financial asset holdings: increasing

[2]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, 1994), 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.

the benefit replacement rate by 10 percent lowers financial wealth-to-income ratios by 0.7% to 2.8%. This effect tends to be strongest 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 precautionary motive is an important determinant of individual savings behavior, and that the unemployment insurance system, while potentially raising welfare, is also lowering wealth accumulation.

The paper proceeds as follows. In Part I, we provide a brief overview of the related literature on precautionary savings behavior. Part II discusses the structure of the stochastic life-cycle simulation model. Part III presents simulation results analyzing the saving effects of UI in the model. Part IV outlines our empirical procedures for testing some of the implications of the precautionary saving model. Part V presents the basic empirical results, while Part VI presents some extensions. Part VII concludes.

I. Previous Literature on Precautionary Savings and Social Insurance

A. Simulation Studies

There is a sizeable literature which has considered the role of precautionary savings; see Deaton (1992) and Browning and Lusardi (1995) for reviews of recent developments. 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.[3]

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, Shoven, and Spivak (1987), and Hubbard and Judd (1987) show that, in partial equilibrium models with a

[3]See Skinner (1988), Zeldes (1989), Caballero (1991), Deaton (1991), Carroll (1993), Carroll and Samwick (1993), Engen (1993b), Hubbard, Skinner, and Zeldes (1994) and Aiyagari (1994).

functioning capital market but no annuity market, Social Security can significantly decrease precautionary saving when lifespan is uncertain, while increasing individual welfare.[4] Hubbard, Skinner, and Zeldes (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 taxing away individual savings through means-testing to qualify for government assistance. Kotlikoff's (1988) simulation analysis of public health insurance provides similar conclusions.

In work more closely related to this study, Hansen and Imrohoroglu (1992) simulate the 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.[5]

While providing an interesting analysis of optimal UI benefits determination, Hansen and Imrohoroglu's model probably 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. Asset-income ratios generated by their

---

[4]This reduction in precautionary savings and increase in welfare is somewhat mitigated after accounting for general equilibrium effects. With borrowing constraints, the savings reduction is mitigated further and the welfare increase can be reversed (Hubbard and Judd, 1987). Samwick (1994) demonstrated that the negative offset to saving from Social Security can be reduced if precautionary savings for income uncertainty comprises most of the saving undertaken by workers subject to the Social Security payroll tax.

[5]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, Kiefer, and Possen (1985), Atkeson and Lucas (1995), and Davidson and Woodbury (1995).

model are 0.5 or less, without UI, and as low as zero with a UI program; in fact, the aggregate wealth-income ratio in the U.S. is above 3, and the median household wealth-income ratio tabulated from the 1992 Survey of Consumer Finances is 1.8. 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.

B. Empirical Studies

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.

Early attempts to find supporting evidence for the precautionary motive are 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. However, occupation is an extremely crude measure of consumer uncertainty, and, as Skinner notes, if less risk-averse individuals self-select into occupations with riskier incomes then the precautionary saving theory does not necessarily imply that people in these occupations will save more.

One set of recent studies used households' subjective assessments of their income uncertainty. Guiso, Jappelli, and Terlizzese (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 variance for their next year's income. But this subjective variance may not be an accurate estimate of the longer-horizon uncertainty required by the theory. Their estimates of the risk facing the individual are only one-tenth to one-fifth as large as the estimates of income uncertainty from U.S. panel data. Furthermore, individual assessments of a given level of uncertainty may also be correlated with tastes for savings, so that there may be little predictive power for the effect of an exogenous change in the uncertainty facing all persons. This also is a potential problem with the study by Lusardi (1993), which also uses Italian data and finds a somewhat bigger wealth effect than Guiso *et al.* using a subjective earnings variance measure.

Another set of studies has examined the effect of uncertainty differentials across job types and education levels on consumption patterns. Dardanoni (1991), using data on British households, found average consumption across occupation and industry groups to be significantly lower when income variance was greater. But Dynan (1993) finds that the quarterly variance of consumption expenditures is not a significant predictor of the quarterly growth rate of consumption, when this variance is instrumented by education, occupation, or industry. And 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 has no statistically significant effect on current consumption, although a rise in Kimball's (1990) "equivalent precautionary premium," which can only be calculated under certain restrictive assumptions, is found to significantly decrease current consumption.

There are important problems with these approaches, however. For example, the cross-sectional variance of income used by Dardanoni is not necessarily related to the time-series variation

of income required by the precautionary saving model. And it is likely that Dynan's innovative approach asks too much of the limited, available data. In her data, the Consumer Expenditure Survey, consumption growth rates are only available for 3 quarters, and a three-period measure of the variance is quite likely to be an extremely noisy measure of actual quarterly reevaluations of expected lifetime income. Finally, as Carroll and Samwick (1994) note, the theoretical relationship between current consumption, or consumption growth, and income uncertainty can be quite subtle in the precautionary saving model. On the other hand, there is a clearer theoretical implication for the relationship between uncertainty and wealth: greater income uncertainty increases precautionary stocks of saving.[6]

Several recent studies have considered the effect of uncertainty on wealth holdings. 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. But for this sample income uncertainty over the next 15 years may well be a function of initial wealth holdings; for example, those with higher wealth initially may be more likely to retire, yielding the impression that their income is more uncertain due to a large reduction in income in some future year. Mercurio (1995) estimates a strong precautionary saving effect using data from the SIPP, but it is not clear whether this study is merely finding that households with faster-growing incomes save more, or whether there is evidence of precautionary saving. Moreover, he finds that broader measures of wealth, which includes housing equity, are more sensitive to income uncertainty than narrower, more liquid measures of household assets; this result is conterintuitive to general models of saving which include precautionary saving. Carroll and Samwick (1995) 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.

---

[6]In a standard precautionary savings model, for example, individuals facing a higher variance of lifetime income will have lower consumption at young ages but higher consumption at older ages; they will have a higher ratio of wealth-to-income at all ages, however.

They find that unpredictable income uncertainty is a potentially important predictor of household wealth-income ratios. However, their results are sensitive to whether older workers and the self-employed and farmers are included in the sample.[7]

The mixed results of these studies may be at least partially attributable to the difficult of 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. 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]

II. Theoretical Framework

In this section, we examine unemployment insurance in the context of a stochastic life-cycle simulation model.[9] Our goal in doing so is to answer two questions which help guide the empirical analysis. First, under realistic assumptions about the risk of unemployment and the generosity of

---

[7]In related work, Guiso, Jappelli, and Terlizzese (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. Starr-McCluer (1994) finds that having health insurance reduces households' financial asset holdings, once the endogeneity of health insurance coverage is taken into account. 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. Levin (1995) studies how uncertainty about medical expenses and government health programs (i.e., Medicare and Medicaid) affects the demand for health insurance.

[8] There are results in related literatures, using potentially more plausible identification of risk, which lend support to the precautionary motive. 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. And Fishback and Kantor (1994) find that increases in the generosity of insurance for on-the-job injuries in the early 20th century lowered household savings.

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

unemployment insurance, is UI predicted to have a significant effect on household asset accumulation? Second, how does this effect vary with observable characteristics of the population? Simulation analysis is necessary since it is only under implausibly simple conditions that analytical solutions for household wealth can be derived.

Our multiperiod 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.

A. 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, $\{C_{t+1},...,C_D\}$,[11] to maximize expected lifetime utility (ELU):

$$ELU_t = E_t[\sum_{j=t}^{D} (1+\delta)^{t-j} U(C_j)], \tag{1}$$

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.

---

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

Consumption, actual and planned, is restricted to be non-negative: $C_{t+j} \geq 0 \ \forall j=0,...,D-t$.

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 substitution), which implies that risky labor income leads to precautionary saving, and is of the iso-elastic form:

$$\begin{aligned} U(C_t) &= C_t^{1-\gamma}/(1-\gamma) \qquad \text{if } \gamma \neq 1, \\ &= \ln(C_t) \qquad \text{if } \gamma = 1. \end{aligned} \tag{2}$$

where γ denotes the relative risk aversion coefficient ($\gamma>0$) and the intertemporal elasticity of substitution for consumption is equal to $1/\gamma$.[12]

B. 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:

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

[12]The structural preference parameters for this stochastic life-cycle model are estimated using household panel data in Engen (1993a). Estimates of the intertemporal elasticity of substitution, $1/\gamma$, are in the range of 0.30 and 0.38 which implies a risk aversion coefficient, γ, between 2.6 and 3.3. Estimates of the time preference rate, δ, are in the range of 0.04 to 0.08. For the simulations of the model, these estimates provide a range of estimates for γ and δ that are consistent with values used in many other studies.

with general taxes on consumption, interest income, and labor income represented by τ^c, τ^i, 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 $W_{D+1} \geq 0$.[13]

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. 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.[14] The magnitude of the probability of being unemployed and receiving zero labor income will significantly affect the amount of precautionary saving undertaken. The baseline probability of being employed given employment in the previous period is set to $p^1=.97$, and the probability of remaining unemployed the next quarter if unemployed in the current quarter is $p^0=.35$. These probabilities were calculated by the authors using monthly data on employment and unemployment for 1984-1987 from

[13]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.

[14]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.

the Panel Study of Income Dynamics.

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 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: $B_t=\{\min[\overline{B},\max(\underline{B},\theta Y_{t-1})]\}$. The gross replacement rate for UI benefits is denoted by θ, with the maximum UI benefit level equal to $\overline{B}$ 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.

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.[15] Unemployment benefits are available for 2 quarters, consistent with the 26 week limit facing most workers in the U.S.[16]

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.[17] The payroll tax is levied only on labor income up to a limit of $\overline{Y}$, 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

---

[15]This average is quite similar to the average replacement rate, and is almost exactly equal to the median replacement rate, in our SIPP sample.

[16]In practice, benefit duration can vary as a function of several factors: earnings over the previous 5 quarters; state extensions of benefits when unemployment exceeds a "trigger" level; and federal extension of benefits in national recessions. We show the effects of varying UI benefits durations below.

[17]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).

determined endogenously so that tax revenues cover UI payouts. As a result, on average, there are no income effects from the UI system: any effects on savings behavior occur solely through reduced income risk conditional on a mean level of income. 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.[18] This may have a second order effect on savings behavior, which should be quite small relative to the effects of risk reduction.

Tax revenue from the income tax and a sales tax on consumption 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.[19] The individual receives a certain stream of pension income, P_t, when retired until death. 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.[20]

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 human capital adjustment

---

[18]In fact, we 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. There is further redistribution under the UI system in reality; we discuss this in the conclusion, where we compare our empirical and theoretical results.

[19]Part of the tax on wages consists of the SS payroll tax.

[20]The sum of pension and SS 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.

factor 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 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. Appendices A and B provide some details concerning the specification of the stochastic labor earnings process. Appendix C provides a brief discussion of the technique for solving the model.

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 (ie. 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 (1994) 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.

III. UI and Household Saving Simulations

Table 1 shows the effect of a UI program on individual median asset-to-income ratios, for individuals age 25 to 64, and expected consumer welfare in the stochastic life-cycle model. The model generates a median asset-income ratio of 3.52 for working individuals in the absence of UI. The mean asset-income ratio, not shown in the table, is somewhat higher (4.35), reflecting the skewed nature of wealth accumulation. The effects of UI on median and mean asset ratios are quite similar,

so the results for means are not reported in Table 1. Precautionary saving represents approximately 68 percent of total saving in this initial scenario.

Introducing a UI program dramatically reduces saving accumulation; the median asset-income ratio falls by 24% 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.[21] Thus, UI can have a sizeable impact on household asset accumulation decisions in a plausible savings model.

The next four rows of Table 1 alter the risk-aversion coefficient and time preference rate used in our baseline simulations. The absolute effects of UI are larger for more risk averse workers, while the rate of time preference has little absolute effect. Overall, the results are fairly robust to variations in these key parameters.

While it is useful to demonstrate this robustness, variations in these parameters do not provide useful testable implications since we cannot observe households' underlying preferences in standard micro-data sets. Therefore, we next consider two observable dimensions along which the effect of UI might be expected to vary. The first is the underlying level of unemployment risk. In rows six and seven of Table 1, the probability of being unemployed ($1-p^1$), conditional on currently being employed, is changed from the initial simulation value of 3 percent. For the median 25-64 year old, the reduction in the asset/income ratio from introducing UI is twice as large at a probability of unemployment of 4 percent than at a probability of unemployment of 2 percent. The effect continues to be roughly linear as the unemployment risk rises further. Thus, in our empirical analysis, we should expect to find that higher unemployment risk should be positively correlated with a higher UI

[21]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 <u>only</u> the gains in terms of smoothing consumption. See Gruber (1994) for a further discussion of the limitations of this type of model for doing welfare analysis of UI.

crowd-out effect. However, note that at the same time, the model suggests that UI offers a larger welfare gain when unemployment risk is higher.

A second potential source of variation in household saving responses to UI is differences 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 saving, thus reducing the savings effects of UI on 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 of UI effects is shown in the final set of rows of Table 1. In absolute terms, the effects of introducing UI rise somewhat with age. But in percentage terms, they fall dramatically. For those below 45, the effect of introducing UI is about 2.6 times as large as for those above age 45.

Of course, in our empirical work we will not observe the introduction of UI but rather variations in benefits generosity. In Table 2, we therefore present the effects of raising or lowering the replacement rate by 10% on asset accumulation and welfare. The first row shows the median asset-income ratio for workers age 25-64; we also report the average UI benefit replacement rate and payroll tax rate necessary to finance the UI program. The following rows also show the welfare change relative to the baseline parameterization. As noted above, the average replacement rate in this baseline specification corresponds closely to the median replacement rate in our SIPP sample. As the slope coefficient θ rises, increasing the generosity of the UI system, there is a decrease in asset accumulation and a rise in welfare. Overall, raising the replacement rate by 10% lowers the wealth/income ratio by 1.1%. Lowering it by 10% has a slightly larger effect, raising the wealth/income ratio by 1.5%. Although not shown on the table, continually lowering the benefit replacement rate gives larger proportional increases in the asset-income ratio. This is because the marginal utility of the additional income, and thus the reduced need for precautionary savings, is lower as UI replaces a larger share of

pre-unemployment resources.

The remainder of the table shows how the effects of changing the replacement rate vary with the underlying unemployment risk and with age. The basic results from Table 1 are reiterated in this analysis: the (percentage) crowdout effects of UI on wealth accumulation are larger for younger workers (slightly more than twice as large for those under 45 as for those over 45); and the crowdout of UI rises with the underlying level of unemployment risk (with the crowdout effects rising roughly linearly with the level of unemployment risk). These findings therefore provide two implications that are testable in our cross-sectional data analysis.[22]

Part IV: Empirical Strategy

A. Data

In order to test some of the implications of the precautionary saving model and further investigate the effect of UI on wealth accumulation, we use the largest annual data set available with information on both wealth and 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,

[22]A natural addition to the analysis is to incorporate the stochastic life-cycle model into a general equilibrium setting. A general equilibrium model allows for feedback effects of UI policy changes: the lower capital stock raises the rate of return, lowers wages, and increases taxes on this smaller wage base, tempering the rise in welfare and the decline in the capital stock. It is difficult, however, to realistically incorporate the aggregate savings implications of the UI program into our modelling framework; which does not capture the cross-sectional richness of savings behavior in the U.S.; as we demonstrate empirically, the effects of variations in UI generosity on the aggregate capital stock are likely to be small, even as the program does have a sizeable effect on the median worker. Adding standard GE feedbacks to the model in earlier drafts reduced somewhat both the savings and welfare effects of the program, but the basic pattern of results was similar.

individuals are asked about their wealth holdings in selected waves. For the panels in the earlier years of the SIPP (1984 through 1986), wealth inventories were collected twice, in the fourth and seventh wave. For 1987 and 1990, there is wealth information only in the fourth wave; there is no data for 1988 and 1989.[23] As a result, we use wealth information for a total of 8 waves spanning the 1984 to 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;[24] 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 24,903 observations.[25]

We use two measures of wealth: gross financial and net financial assets. Gross financial assets consist 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; net assets subtracts off unsecured

---

[23]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.

[24]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.

[25]Since there are two wealth interviews in 1984-86, 5562 households have two observations for wealth. Deleting either the first or second wealth observation for these households has little effect on the results.

debt from this total. The former has the advantage that gross asset holdings most closely approximate the steady state stock of accumulated precautionary savings, whereas net assets also reflect short run smoothing through consumption loans. Furthermore, debt holding may not proxy for financial constraint but rather access to debt markets.[26] On the other hand, changes in unemployment insurance generosity should be reflected in borrowing as well as savings behavior. Thus, it is of interest to examine separately both gross and net asset holdings.

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.[27] 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, there are two costs to this approach: first, we ignore any potential response of non-financial assets to UI generosity; and second, we limit the ability to extend our empirical results to aggregate behavior. We therefore discuss below the effect of expanding our dependent variable to include less liquid measures of wealth as well.

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.,

---

[26]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.

[27]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.

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.[28]

B. Empirical Framework

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

$$WEALTH_i = \alpha + \beta_1 X_i + \beta_2 UI_i + \beta_3 \delta_j + \beta_4 \tau_t + \epsilon_{ijt} \quad (4)$$

where i indexes households, j indexes states, and t indexes years. WEALTH is 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 labor 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; 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

[28]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 on 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.

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).[29]

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,

[29]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.

higher wage workers have higher wealth to income ratios. Thus, if wages are not controlled for in the estimation, there will 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 should be sufficient to pick up any nonlinearities in UI benefits schedules. Our results are not sensitive to the inclusion of additional higher order terms in the average weekly wage.

A related issue 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 have therefore rerun our basic specifications using log wealth as the dependent variable, which overcomes this problem: there is no common element in the dependent and independent variables in this case. The results are similar to those reported below, although the estimated savings response is somewhat larger at the mean level of savings. A problem with this specification, however, is that we lose any information on those with zero or negative wealth holdings; this is especially troublesome in the case of net financial assets (where we lose almost one-half of our observations). We therefore use the wealth/income ratio as the dependent variable for the results reported below. 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.[30]

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.[31]

The means and distributions of the SIPP data are presented in Tables 3A and 3B. Table 3A presents the means and medians of the covariates included in our model. Twenty-four percent of the sample is female, and thirteen percent is non-white. Roughly three-quarters of the spouses are working. The average real weekly wage is $418. The average replacement rate, after SS taxes only, is 43.6 percent of previous wages; the median is 45.6 percent.

For our wealth measures, we present both the mean and the distribution in Table 3B. Two facts about wealth are immediately apparent. First, the distribution of wealth is highly skewed. In order to account for this, we will use robust regression techniques which first exclude influential outlying observations, then iterate towards a solution by downweighting those observations with larger residuals (Berk, 1990). 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 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 the means and medians somewhat, as the second panel of the Table shows. But asset holdings remain quite low even in this case.

---

[30]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.

[31] 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).

Part V: Empirical Results

A. Basic Results

Our basic regression results are presented in Table 4, for both gross and net financial assets. In both cases, there is a significant negative correlation between UI generosity and wealth levels. In the first column, the coefficient indicates that raising the replacement rate by 10 percent (0.0436 percentage points) lowers the gross asset to income ratio by 0.16 percentage points, or 0.7% at the mean gross financial asset to income ratio. For net financial assets, the replacement rate coefficient is also highly significant, and it is much more sizeable; raising the replacement rate by 10 percent would lower mean net savings by 2.8%.[32] This suggests that there is a response of borrowing behavior, as well as savings behavior, to UI generosity levels. 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.[33]

While large in percentage terms, these effects are fairly small when expressed in terms of dollars. We find that raising the replacement rate by 10 percent reduces savings by 0.16-0.47% of income, or $50 to $147 at the mean income in our sample. This small finding in dollars terms is not surprising given the low wealth 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 mean real weekly wage in our sample, raising the replacement rate by 10% would raise weekly income during unemployment spells by $18. The average person who becomes unemployed spends 13.1 weeks unemployed, so that the 10% replacement rate rise would raise his

---

[32]The sample for net wealth is somewhat smaller due to missing data on uncollateralized borrowing. The change in coefficient across the two columns is not due to the smaller sample when we use net wealth; the coefficient for gross wealth is quite similar in this sample to the full sample result.

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

income over the spell by $236.[34] Thus, the reduction in wealth from this 10% replacement rate rise is 21% to 63% as large as the increase in UI benefits during an unemployment spell.[35]

The control variables generally have their expected signs. Savings is predicted to fall at young ages (below age 29 for both gross and net assets), and rise thereafter (evaluated at the mean 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 and work variables. At the mean values, being married with a non-working spouse is associated with an increase in the asset/income ratio of 0.009 percentage points in column (1); if there is a spouse who works, the increment is only 0.004 percentage points. The savings propensity rises more steeply with age for higher wage earners.

B. Sensitivity to Specification

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

---

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

[35]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. 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 (see the conclusion for further discussion of this redistribution). This is impossible to do with our short panel of data on employment.

Note also that these findings are not inconsistent with those of Gruber (1994), whose significant consumption smoothing results implied that UI generosity increases were not simply crowding out other forms of consumption insurance (such as savings). Gruber finds that raising the replacement rate by 10 percent increased consumption while unemployed by 1.3%. For an individual who becomes unemployed, the same 10% replacement rate rise increases net (of the savings decline) resources by 4 to 8%, so that this is a plausible consumption response.

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 provide 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 different savings propensities for high and low wage earners 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 first row of Table 5. 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.

C. Sensitivity to Wealth Measurement

As noted above, we have restricted our analysis of the SIPP data to financial assets, and we have excluded imputed values. The former restriction seems to us an important one given the arguments presented above about liquidity and consumption responses. Nevertheless, it is also useful to examine how our results are affected by moving to a broader measure of assets. Thus, in the second row of Table 5, we use as our wealth measure net financial assets plus retirement assets (in IRA or Keogh accounts) plus housing wealth. This represents the bulk of collateralizable wealth for

the households in our sample.[36]

In absolute value, the effect of unemployment insurance on this broader measure of wealth lies in between the response of gross and net assets shown in Table 4. 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 0.28%. 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 4. We return to this point in the conclusion.

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 next row of Table 5, therefore, we include the imputed observations in our regressions. The estimated coefficients are larger than those reported in Table 4. The means and medians of wealth 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.

Part VI: Extensions

A. 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 the individual. In this section we test by proposition by

[36]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).

running regressions of the form:

$$WEALTH_i = \alpha + \beta_1 X_i + \beta_2 UI_i + \beta_3 UNEM_i + \beta_4 UNEM_i{*}UI_i + \beta_5 \delta_j + \beta_6 \tau_t + \beta_7 w_i{*}\delta_j + \beta_8 w_i{*}\tau_t + \epsilon_{ijt} \quad (4)$$

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, β_2 will be small, since it now measures the effect on those facing low unemployment risk, and β_4 will be large (in absolute value), since it 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 9 major industry categories and 9 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 wavely, 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.[37]

Our findings are reported in the first panel of Table 6. In fact, we do see the hypothesized pattern. The main effect (β_2, the coefficient on the replacement rate) is insignificant in both specifications, while the interaction term (β_4, the coefficient on the replacement rate times risk) is sizeable; it is only significant for gross assets, however.[38] That is, the crowdout effect on savings is largest for those individuals who face the greatest unemployment risk. The gross asset results imply that increase the level of unemployment risk from 2% to 4% would increase the crowdout effect of a 10% replacement rate rise from 0.01% of income to 0.09% of income; the net asset results imply that this unemployment risk rise would increase the crowdout effect of a 10% replacement rate rise from 0.32% of income to 0.41% of income. Thus, the underlying level of unemployment risk is an important predictor of the crowd-out effects of UI.

B. 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:

$$\mathrm{WEALTH}_{ijt} = \alpha + \beta_1 X_{ijt} + \sum_{k=0}^{K}(\underline{\gamma}\mathrm{AGECAT}^k_{ijt} + \underline{\pi}\mathrm{AGECAT}^k_{ijt}*\mathrm{UI}_{ijt}) + \beta_2\delta_j + \beta_3\tau_t + \beta_4\tau_t + \beta_5 w_{ijt}*\delta_j + \beta_6 w_{ijt}*\tau_t + \epsilon_{ijt} \quad (5)$$

where AGECAT^k_i is a dummy for individual i being in age category k, up to K different

[37]A higher likelihood of unemployment causes a second-order positive correlation between UI and wealth through consumption smoothing: for those who actually lose their job, their wealth may fall less because UI is more generous. A potential solution to this problem is to focus only on those workers who do not lose their jobs. Using this group introduces an additional bias, however, since their probability of unemployment is most likely lower than average, so that the industry/region unemployment risk will systematically overstate their unemployment risk. In fact, our results are fairly similar if we exclude those who actually become unemployed.

[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.

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 6. For gross assets, the coefficients do decline across the age distribution; for net assets, they are roughly constant. At the same time, however, wealth levels are rising with age, so that the *relative* effect of UI is falling sharply as individuals age for both measures of assets. In brackets under each of the estimates, we present the mean and median wealth to income ratio for each age group. Relative to either the mean or the median, the proportional effect of UI on savings is falling 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.

C. 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.[39] 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

[39]See Lundberg (1985) for direct evidence on the response of spousal labor supply to the unemployment of the household head. Note that "head" vs. "spouse" in our context is defined by the SIPP interviewing convention. In almost all cases, the man is labelled as the head for married couples, regardless of who is the primary income earner. This provides further motivation for using a family replacement rate, as described below.

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 equation (5).

The results of this exercise are in the next panel of Table 6. 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 fairly sizeable for gross assets, but it is small for net assets.

D. Excluding Single Mothers

One potential problem with our approach is that we have attributed our finding to differences in UI generosity across individuals. In fact, UI is only one element of the safety net provided by states, and savings should respond to all components of the state social insurance package.[40] The two most important potential omitted programs are Workers' Compensation, which provides income support for on the job accidents, and Aid to Families with Dependent Children (AFDC), which provides cash welfare. By including state effects and state*wage interactions, fixed state differences in other social insurance program generosity will be controlled for. But correlated changes within states over time in these programs, may lead to omitted variables bias.

We are unable to fully differentiate the effects of UI and Workers' Compensation in our results, although this distinction is addressed to some extent with our unemployment risk interaction model above. But we can differentiate the effects of AFDC, since eligibility for this program is

[40]There are also a large number of social insurance programs at the federal level which may affect precautionary savings, such as Social Security or Medicare. These programs do not vary at the state level, however, so that they should not have a differential effect on workers in our sample once we condition on demographic differences.

largely restricted to single mothers.[41] In the final panel of Table 6, we exclude single mothers from our sample. In fact, the coefficients rise slightly relative to our findings in Table 4. This is consistent with the view that single mothers are less responsive to UI generosity, since there is another social insurance program on which they can rely if they lose their jobs.

VII. 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; the results are large in percentage terms, but small in absolute amounts, reflecting the small average asset holdings of non-elderly families in the U.S.

In terms of comparative statics, the predictions of our model and the results of our empirical analysis are quite consistent: UI crowds out individual savings, and this crowdout rises with the extent of unemployment risk and falls with age. In terms of magnitudes, however, the results are quite different. In the model, increasing the average replacement rate by 10% lowered wealth by 3 percent of income. In the empirical results, this same change is predicted to lower asset/income ratios by only 0.2 to 0.5 percent of income. There are several sources of the difference in these results. Understanding these sources, and incorporating them into the modelling of

[41]Some married couples can get AFDC under the AFDC-UP option, but this represented only 5% of the caseload in 1990 (U.S. Committee on Ways and Means, 1992).

savings for unemployment, should be a high priority for future work.

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 modelled the effect of UI on financial assets, which represent a small share of total asset holdings. In a richer model which incorporated different types of assets with different degrees of liquidity, the response of financial assets only to UI generosity might be smaller than that predicted by our model.

The second is that we have ignored all other forms of insurance for unemployment in our modelling. While we noted earlier that there is evidence of less than full insurance for unemployment spells, 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. This is evident in our empirical finding that the effects of UI are stronger for single workers and weaker for female heads (who can rely on AFDC).

Third, the model does not incorporate the redistribution that is a feature of the UI system in the U.S. The UI system redistributes in two ways: from high wage to low wage workers, since benefits are not directly proportional to wages but taxes are proportional to payroll;[42] and from low unemployment risk to high unemployment risk workers, since employer experience rating is imperfect. Thus, while there is no income effect on average from the UI system, under the assumption of full shifting to wages on average of the payroll tax, there may be relative income effects across groups. These relative income effects may in turn affect savings behavior. The redistribution from high to low wage workers raises the crowdout effect of UI, since savings propensities rise with income. The effect of the redistribution from low to high risk workers, however, is ambiguous: high unemployment risk workers should save more in a precautionary framework but, as we have highlighted, if risk-loving

[42]Computing the extent of redistribution here is difficult because the payroll taxes are themselves capped, limiting the redistribution from high wage workers.

individuals select into these types of jobs they may be saving less on average.

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. While having fairly sizeable effects on the median worker in the model, the UI program may have much smaller effects on the aggregate capital stock, for two reasons. First, financial assets represent only a small part of the aggregate capital stock, and we are unable to demonstrate any effect of UI on other types of capital accumulation. Second, the median worker that we model in this paper only does a small share of aggregate savings. Most savings is done by higher income individuals whose income does not depend in an important way on the generosity of the UI system. This point can be addressed by extending models of this type to model the underlying differences in savings behavior across families of different income levels. 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, Skinner, and Zeldes (1994, 1995) provides one rationale for the wider distribution of accumulated wealth.

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.

References

Abel, Andrew B., "Precautionary Saving and Accidental Bequests," American Economic Review 75 (September 1985): 77-791.

Abowd, John M., and David Card, "Intertemporal Labor Supply and Long-Term Employment Contracts," American Economic Review 77 (March 1987): 50-68.

_____., "On the Covariance Structure of Earnings and Hours Change," Econometrica 57 (March 1989): 411-446.

Adams, J, "Permanent Differences in Unemployment and Permanent Wage Differentials," Quarterly Journal of Economics, 100 (1985), 29-50.

Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," Quarterly Journal of Economics 109 (August 1994a): 659-684.

Altonji, Joseph and Aloysius Siow, "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data," Quarterly Journal of Economics 102 (May 1987): 293-328.

Anderson, Patricia M. and Bruce D. Meyer (1994). "The Effect of Unemployment Insurance Taxes and Benefits on Layoffs Using Firm and Individual Data". NBER Working Paper #4960.

Anderson, Patricia M. and Bruce D. Meyer (1995). "The Incidence of a Firm-Varying Payroll Tax: The Case of Unemployment Insurance". NBER Working Paper #5201.

Atkeson, Andrew and Robert E. Lucas, Jr., "Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance," Journal of Economic Theory 66 (June 1995): pp.64-88.

Baily, Martin Neil, "Some Aspects of Optimal Unemployment Insurance" Journal of Public Economics, 10 (1978): 379-402.

Berk, R. A., "A Primer on Robust Regression," in J. Fox and J. S. Long, eds, Modern Methods of Data Analysis. Newbury Park, CA: Sage Publications, 1990.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," mimeo, April 1995.

Caballero, Ricardo J., "Earnings Uncertainty and Aggregate Wealth Accumulation," American Economic Review, 81 (1991): 859-871.

Carroll, Christopher, "Buffer Stock Saving and the Permanent Income Hypothesis," mimeo, Board of Governors of the Federal Reserve System, 1993.

_____., "How Does Future Income Affect Current Consumption?" Quarterly Journal of Economics 109 (February 1994): 111-148.

_____., and Andrew Samwick, "How Important is Precautionary Saving?" mimeo, Board of

Governors of the Federal Reserve System, 1995a.

_____., "The Nature and Magnitude of Precautionary Wealth," mimeo, Board of Governors of the Federal Reserve System, 1995b.

Cochrane, John H. "A Simple Test of Consumption Insurance". Journal of Political Economy, 99 (1991): 957-976.

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

Dardanoni, Valentino, "Precautionary Savings Under Income Uncertainty: A Cross-Sectional Analysis," Applied Economics, 23 (1991): 153-160.

Davidson, Carl and Stephen Woodbury, "Optimal Unemployment Insurance," mimeo, 1995.

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

_____., Understanding Consumption, Oxford: Oxford University Press, 1992.

Diamond, Peter, "A Framework for Social Security Analysis," Journal of Public Economics, 8 (1977), 275-298.

Duncan, Gary J., and Daniel H. Hill, "An Investigation of the Extent and Consequences of Measurement Error in Labor-Economic Survey Data," Journal of Labor Economics 3 (1985): 508-532.

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

Easley, David, Nicholas Kiefer, and Uri Possen, "An Equilibrium Analysis of Optimal Unemployment Insurance and Taxation," Quarterly Journal of Economics 100 (suppl., 1985): 989-1010.

Employment and Training Administration (various years). "Significant Provisions of State Unemployment Insurance Laws". Washington, D.C.: U.S. Department of Labor.

Engen, Eric, "A Stochastic Life-Cycle Model with Mortality Risk: Estimation with Panel Data." Mimeo, UCLA, 1993a.

_____., "Consumption and Savings in a Life-Cycle Model with Stochastic Earnings and Mortality Risk." Mimeo, UCLA, 1993b.

Fishback, Price V., and Shawn Everett Kantor, "Insurance Rationing and the Origins of Workers' Compensation". NBER Working Paper #4943, December 1994.

Fisher, Malcolm, "Exploration in Savings Behavior," Oxford University Institute Statistics Bulletin 18

(August 1956): 201-277.

Fleming, J.S., "Aspects of Optimal Unemployment Insurance: Search, Leisure, Savings, and Capital Market Imperfections," Journal of Public Economics 10 (December 1978): 403-425.

Friedman, Milton, A Theory of the Consumption Function, Princeton: Princeton University Press, 1957

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," NBER Working Paper #4750, May 1994.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, "Earnings Uncertainty and Precautionary Savings," Journal of Monetary Economics, 30 (1992): 307-337.

_____., "Income Risk, Borrowing Constraints and Portfolio Risk," CEPR Discussion Paper No. 888, January 1994.

Guiso, Luigi and Tullio Jappelli, "Does Uninsurable Uncertainty Increase the Demand for Insurance against Insurable Risks?" mimeo, 1994.

Hamermesh, Daniel S., "Social Insurance and Consumption: An Empirical Inquiry." American Economic Review, 72 (1982): 101-113.

Hamermesh, Daniel S., and Daniel T. Slesnick, "Unemployment Insurance and Household Welfare: Microeconomic Evidence 1980-93," NBER Working Paper #5315.

Hansen, Gary D., and Ayse Imrohoroglu, "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard," Journal of Political Economy 100 (February 1992): 118-142.

Hoynes, Hilary, Michael Hurd and Harish Chand, "Imputed Wealth, Subjective Survival Probabilities, and Social Security Wealth," Mimeo, NBER, May 1995.

Hubbard, R. Glenn and Kenneth L. Judd, "Social Security and Individual Welfare: Precautionary Saving, Liquidity Constraints, and the Payroll Tax," American Economic Review, 77 (September 1987), 630-646.

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

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," Journal of Political Economy 103 (April 1995): 360-399.

Hurd, Michael, "Mortality Risk and Bequests," Econometrica 57 (1989): 779-814.

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan (1993). "Earnings Losses of Displaced Workers," American Economic Review, 83: 685-709.

Kazarosian, Mark "Precautionary Savings--A Panel Study," mimeo. 1994.

Kennickell, Arthur, and Martha Starr-McCluer, "Changes in Family Finances from 1989 to 1992: Evidence from the Survey of Consumer Finances," Federal Reserve Bulletin (October 1994).

Kimball, Miles, "Precautionary Saving in the Small and in the Large," Econometrica 58 (1990): 53-73.

Kotlikoff, Laurence J., "Health Expenditures and Precautionary Saving," in L. Kotlikoff, ed., What Determines Saving? Cambridge: MIT Press, 1988.

_____., John Shoven, and Avia Spivak, "Annuity Markets, Savings, and the Capital Stock," in Z. Bodie, John Shoven, and David Wise, eds., Issues in Pension Economics, Chicago: University of Chicago Press, 1987.

Levin, Laurence, "Demand for Health Insurance and Precautionary Motives for Savings among the Elderly," Journal of Public Economics 57 (July 1995): pp. 337-368.

Levine, Phillip B. (1990). "Unemployment Insurance Benefit and Tax Schedules, 1978-1987". Mimeo, Princeton University.

Lillard, Lee A., and Robert J. Willis, "Dynamic Aspects of Earnings Mobility," Econometrica 46 (September 1978): 985-1012.

Lundberg, Shelley, "The Added Worker Effect," Journal of Labor Economics, 1 (1985): 11-37.

Lusardi, Annamaria, "Precautionary Saving and Subjective Earnings Variance," Center for Economic Research, 1993.

MaCurdy, Thomas E., "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," Journal of Econometrics 18 (1982): 83-114.

Mercurio, Matthew, "Saving with Liquidity Constraints Using SIPP Panel Data," mimeo, Princeton University, 1995.

Meyer, Bruce D., "Unemployment Insurance and Unemployment Spells," Econometrica 58 (July 1990): 757-782.

Mortensen, Dale, "A Welfare Analysis of Unemployment Insurance: Variations on Second-Best Themes," Carnegie-Rochester Conference Series on Public Policy 19 (Autumn 1983): 67-97.

Murphy, Kevin M., and Robert H. Topel, "Unemployment, Risk, and Earnings: Testing for Equalizing Wage Differences in the Labor Market," in Kevin Lang and Jonathan S. Leonard, eds. Unemployment and the Structure of Labor Markets. New York: Basil Blackwell, 1987.

Nelson, Julie, "On Testing for Full Insurance Using Consumer Expenditure Survey Data," Journal of Political Economy 102 (April 1994): 384-394.

Palumbo, Michael, "The Effects of Income Uncertainty on Saving and Insurance Purchases: Evidence

from a Survey of Maine's Workers, 1890," mimeo. 1992.

Runkle, David E., "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data," mimeo. Federal Reserve Bank of Minneapolis. November 1988.

Samwick, Andrew, "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving," mimeo. 1994.

Sandmo, Agnar, "The Effect of Uncertainty on Saving Decisions," Review of Economic Studies 37 (1970): 353-360.

Shavell, Steven and Laurence Weiss, "The Optimal Payment of Unemployment Insurance Benefits over Time," Journal of Political Economy 87 (December 1979): 1347-1362.

Sheshinski, E. and Yoram Weiss, "Uncertainty and Optimal Social Security Systems," Quarterly Journal of Economics 96 (1981): 189-206.

Skinner, Jonathan, "Variable Lifespan and the Intertemporal Elasticity of Consumption," Review of Economics and Statistics 67 (November 1985): 616-623.

_____., "Risky Income, Life Cycle Consumption, and Precautionary Saving," Journal of Monetary Economics 22 (September 1988): 237-255.

Starr-McCluer, Martha, "Health Insurance and Precautionary Saving," mimeo, Federal Reserve Board, 1994.

Topel, Robert, "On Layoffs and Unemployment Insurance", American Economic Review, 73 (1983): 541-559.

U.S. Committee on Ways and Means, Green Book 1992: Background Material on Programs Under the Jurisdiction of the Committee on Ways and Means, Washington D.C.: GPO, 1992

Walker, James R. 1993. "Migration Among Low-Income Households: Helping the Witch Doctors Reach Consensus," mimeo, University of Wisconsin-Madison, 1993.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations form Certainty Equivalence," Quarterly Journal of Economics 104 (May 1989): 275-298.

_____., "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy 97 (April 1989): 305-346.

Table 1: Effects of Unemployment Insurance on Asset Accumulation and Consumer Welfare

	Without UI	With UI		
Parameters	Median Asset/Income Ratio	Median Asset/Income Ratio	Δ Asset/Income Ratio	%Δ consumer welfare
Baseline	3.52	2.68	-0.84	8.9
γ=2	3.07	2.29	-0.78	5.0
γ=3	3.92	3.00	-0.92	13.5
δ=0.06	3.20	2.34	-0.86	9.1
δ=0.04	3.89	3.04	-0.85	9.0
(1-p1)=.07	6.10	4.24	-1.86	19.7
(1-p1)=.04	4.24	3.12	-1.12	11.8
(1-p1)=.02	2.66	2.19	-0.47	4.8
Age (Baseline)				
25-34	1.08	0.67	-0.41	---
35-44	2.32	1.60	-0.72	---
45-54	4.07	3.22	-0.85	---
55-65	5.88	5.03	-0.85	---

Note: Table 1 compares steady-state asset-income ratios in a partial equilibrium framework. The % change in consumer welfare is calculated as an equivalent variation measure of the mean change in consumer welfare as a percentage of mean expected lifetime income.

Baseline parameters for model simulations:

risk aversion coefficient (γ)=2.5
time preference rate (δ)=0.05 (annual rate)
maximum lifespan (D)=75 years old
variance of log earnings innovations (σ_ϵ^2)=0.05 (annual)
AR(1) coefficient (ρ)=0.85
employment/unemployment probabilities: p0=.35, p1=.97
initial UI payroll tax=0.082
real net return=0.035 (annual)
UI benefit formula rate (θ)=0.5
minimum UI benefit=$325/qtr.
maximum UI benefit=$3,900/qtr.
initial average benefit replacement rate = 0.47
initial average UI payroll tax rate = 8.2%
UI payroll tax earnings limit ($\bar{Y}$) = $7,000
UI benefits received for a maximum of 2 consecutive periods (1/2 year).

Table 2: Effects of Alternative Unemployment Insurance Benefit Replacement Rates on Household Assets and Consumer Welfare

Parameters	Median Asset/Inc Ratio	Δ Asset/Inc. Ratio	%Δ consumer welfare	Average UI Replace Rate	UI Tax Rate
(1-pl)=0.03. θ=.5	2.68	---	---	0.47	8.2
θ=.55	2.65	-0.03	0.2	0.52	8.8
Age: 25-34	0.66	-0.01	---	---	---
35-44	1.58	-0.02	---	---	---
45-54	3.19	-0.03	---	---	---
55-65	5.01	-0.02	---	---	---
θ=.45	2.72	0.04	-0.3	0.42	7.9
Age: 25-34	0.68	0.01	---	---	---
35-44	1.62	0.02	---	---	---
45-54	3.26	0.04	---	---	---
55-65	5.06	0.03	---	---	---
(1-pl)=0.02, θ=.5	2.19	---	---	0.47	5.5
θ=.55	2.17	-0.02	0.1	0.52	6.0
θ=.45	2.21	0.02	-0.1	0.42	5.3
(1-pl)=0.04, θ=.5	3.12	---	---	0.47	10.8
θ=.55	3.08	-0.04	0.3	0.52	11.5
θ=.45	3.17	0.05	-0.4	0.42	10.4
(1-pl)=0.07, θ=.5	4.24	---		0.47	17.2
θ=.55	4.15	-0.09	0.7	0.52	18.6
θ=.45	4.34	0.10	-0.8	0.42	16.3

Note: The UI payroll tax rate (τ^p) is determined endogenously to satisfy the government budget constraint. Initial parameters described in note to Table 1.

Table 3A: Means and Medians of SIPP Covariates

	Mean	Median
Age	40.55 (10.32)	39
Female	0.240	
Years of Education	13.37 (2.96)	13
White	0.873	
Married	0.684	
Spouse's Education (if married)	12.78 (2.86)	12
Number of Kids	0.99 (1.16)	1
Real Weekly Wage	417.8 (232.4)	376.3
Real Family Income	29,176 (16,643)	26,030
After SS Tax Replacement Rate	0.436 (0.135)	0.456
Number of Obs	24903	24903

Note: Based on authors' tabulations of '84-'90 SIPP data. Standard deviations in parentheses.

Table 3B: Distribution of SIPP Wealth Data

	Without Imputed Values			
	Gross	Net	Gross/Inc	Net/Inc
Mean	10229 (49429)	7625 (47992)	0.251 (0.673)	0.166 (0.761)
10th %ile	0	-4229	0	-0.175
25th %ile	232	-1177	0.010	-0.050
50th %ile	1614	325	0.056	0.014
75th %ile	7418	5786	0.218	0.178
90th %ile	24297	22413	0.657	0.622
Number of Obs	24903	23162	24093	23162
	With Imputed Values			
	Gross	Net	Gross/Inc	Net/Inc
Mean	12974 (52089)	7625 (47992)	0.308 (0.810)	0.215 (0.859)
10th %ile	0	-4301	0	-0.172
25th %ile	303	-1175	0.013	-0.048
50th %ile	2166	654	0.071	0.024
75th %ile	9667	8017	0.279	0.232
90th %ile	31767	29665	0.809	0.757
Number of Obs	33434	33434	33434	33434

Note: 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.

Table 4: Basic Results

	(1) Gross/Income	(2) Net/Income
Replacement Rate	**-0.0378**	**-0.1072**
	(0.0119)	**(0.0255)**
Age	-0.0014	-0.0034
	(0.0004)	(0.0009)
Age Squared (*100)	0.0021	0.0056
	(0.0005)	(0.0011)
Years of Education	0.0029	0.0013
	(0.0002)	(0.0005)
Female	0.0031	-0.0129
	(0.0016)	(0.0035)
Married	-0.0138	-0.0192
	(0.0033)	(0.0072)
White	0.0115	0.0160
	(0.0016)	(0.0034)
Spouse's Ed.	0.0018	0.0014
	(0.0002)	(0.0005)
Number of Kids	-0.0059	-0.0102
	(0.0005)	(0.0011)
Spouse Works	-0.0044	-0.0180
	(0.0015)	(0.0034)
Income ($*10^4$)	0.0119	0.0152
	(0.0016)	(0.0036)
Income^2 ($*10^{10}$)	-0.0758	-0.0996
	(0.0115)	(0.0256)
Income * Age ($*10^7$)	0.0569	0.0188
	(0.0298)	(0.0065)
Number of Obs	24903	23162

Note: Standard errors in parentheses. Regression also includes a quartic in wages and a full set of panel dummies and state dummies.

Table 5: Sensitivity Checks

	(1) Gross/Income	(2) Net/Income	(3) (Net + Ret. Assets + Housing)/Income
Include Wage Interactions	-0.041 (0.013)	-0.104 (0.027)	
Broader Measure of Assets			-0.064 (0.136)
Number of Obs			21164
Including Imputed Values	-0.068 (0.014)	-0.139 (0.026)	
Number of Obs	33434	33434	

Note: Standard errors in parentheses. Each row reports the coefficient on the replacement rate from a separate regression which also includes the control variables listed in the Table 4 and the note to that Table.

Table 6: Extensions

	(1) Gross/Income	(2) Net/Income
	Unemployment Risk Interactions	
Repl Rate	0.017 (0.024)	-0.055 (0.052)
RR*Unemployment Risk	-0.921 (0.355)	-0.976 (0.765)
	Age Interactions	
RR* 25-34 Years Old	-0.055 (0.014) [0.134] {0.035}	-0.143 (0.031) [0.026] [0]
RR* 35-44 Years Old	-0.040 (0.014) [0.209] {0.054}	-0.075 (0.030) [0.121] {0.011}
RR* 45-54 Years Old	-0.022 (0.015) [0.326] {0.083}	-0.078 (0.033) [0.247] {0.043}
RR* 55-64 Years Old	-0.006 (0.018) [0.545] {0.172}	-0.068 (0.039) [0.501] {0.140}

Table 6: Extensions, Continued

	(1) Gross/Income	(2) Net/Income
	Marital Status Interactions	
RR* Single	-0.053 (0.014) [0.255] {0.041}	-0.097 (0.031) [0.161] {0.006}
RR* Married	-0.028 (0.014) [0.249] {0.064}	-0.082 (0.029) [0.162] {0.019}
	Excluding Female Heads	
Replacement Rate	-0.044 (0.014)	-0.116 (0.030)

Note: Standard errors in parentheses; group-specific means in [] and medians in { }. Each panel reports the coefficients on the replacement rate and its interactions from separate regressions which also includes the control variables listed in the Table 4 and the note to that Table.

Appendices

A. Mean, Expected Age-Earnings Profile

The specification of the human capital adjustment factor determines an individuals expected age-earnings profile, and 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.

B. Stochastic Process for Wages

The stochastic process for log labor earnings is expressed as:

$$y_t = \bar{y}_t + u_t$$

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; $e_t \sim N(0,\sigma_e^2)$.

Empirical studies by Lillard and Willis (1978), MaCurdy (1982), Abowd and Card (1989, 1991), Hubbard, Skinner, and Zeldes (1994), and Carroll and Samwick (1995) 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 (1986) 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 .003 on a quarterly basis). These are conservative values for the variance and persistence of shocks to income compared to these other studies.

C. 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,...,D, by:

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

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,..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:

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

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,...,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=7,000. Calculating the sample medians, or means, for each age gives a representative age profile for consumption, wealth, earnings, and saving.