PENSIONS, SAVING, AND RETIREMENT UNDER UNCERTAINTY

by

Andrew Alan Samwick

A.B., Economics
Harvard College, 1989

Submitted to the Department of
Economics in Partial Fulfillment of
the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

at the
Massachusetts Institute of Technology
June 1993

Signature of Author ———————

May 4, 1993
Department of Economics

Certified by ———————

James M. Poterba
Professor of Economics
Thesis Supervisor

Certified by ———————

Jerry A. Hausman
John and Jennie S. MacDonald Professor of Economics
Thesis Supervisor

Accepted by ———————

Richard S. Eckaus
Ford International Professor of Economics
Departmental Committee on Graduate Studies

Abstract

The broad issue under study in the sequel is the consequences of the simultaneous provision of social insurance by private and public institutions. The first two chapters (co-authored with Christopher Carroll) demonstrate the nature and magnitude of precautionary saving by households. Because prudent households are necessarily risk averse, these chapters show that social insurance against labor income risk has the potential to improve welfare. With this proposition as background, the remaining chapters use a unique dataset of pension plan formulas to examine the effects of Social Security and employer-provided pensions on both retirement behavior and risk-sharing between firms and workers. The effects of pensions and Social Security on retirement are largely independent, whereas the interaction of the two in the context of risk sharing is quite important.

In Chapter One, we use the *Panel Study of Income Dynamics* to provide some of the first direct evidence that wealth is systematically higher for consumers with greater income uncertainty. However, the apparent pattern of precautionary saving is not consistent with a simple, unconstrained life-cycle optimization problem with patient consumers because wealth is not sufficiently sensitive to uncertainty in permanent income. Instead, the results are found to be more consistent with a buffer stock model of saving in which consumers hold wealth to buffer consumption against near term fluctuations in income but are far less responsive to uncertainty in lifetime (or permanent) income than in the standard model.

In Chapter Two, we extend the precautionary saving model to estimate the fraction of aggregate wealth that is held because some groups of households face greater labor income uncertainty than others. We numerically solve the buffer stock model under a variety of income distributions in order to derive a simple equation for the relationship between wealth and income uncertainty. The adoption of the buffer stock model allows us to use a theoretical, rather than merely statistical, measure of income uncertainty and more rigorously test for a proper specification. We then estimate this equation using data from the PSID and use the empirical estimates to simulate the distribution of wealth that would result if all households had income uncertainty equal to that which is observed for the lowest-uncertainty group. Our results indicate that between 25 and 40 percent of aggregate wealth is held as a result of the "extra" labor income uncertainty faced by most households relative to the uncertainty faced by households in the lowest-uncertainty group.

In Chapter Three, I assess the applicability of the *Pension Provider Survey* of the households in the *Survey of Consumer Finances* 1983 for subsequent economic analysis. A unique feature of these data is that they allow the details of households' pension formulas to be linked to the economic and demographic information typically found in survey data. I illustrate the considerable variation in pension entitlements and accrual rates using representative workers from the household data with an emphasis on the

consequences of plan provisions that may affect retirement decisions but cannot be appropriately modeled in the absence of such detailed data. The retirement incentives of different pension plans are most sensitive to tenure, the wage level, the form of the pension (defined benefit versus defined contribution), and integration with the Social Security benefit formula.

In Chapter Four, I apply this uncommon dataset to the estimation of the combined effect of Social Security and pension benefits on the probability of retirement in a cross-section of the population near retirement age. The improved quality of pension data does not change the basic result of past studies which did not account for pension incentives: Social Security has statistically significant but economically modest effects on retirement behavior. Pensions also have aggregate effects on the probability of retirement comparable to those of Social Security. Moreover, accounting for pension incentives for retirement shifts much of the predicted impact of changes in Social Security to ages before Social Security benefits can be collected, suggesting that pensions may influence retirement in part through relaxing liquidity constraints.

In Chapter Five, I synthesize the main features of the preceding papers to explore the interaction of nonlinear pension formulas and within-job wage uncertainty. In particular, two types of pension formulas become more valuable to the worker in the presence of wage uncertainty. The first has the property that increases in wage uncertainty raise the expected value of pension benefits, thereby increasing total compensation to offset some of the utility lost due to the wage uncertainty. The second has the property that pension benefits are independent of the actual wages paid, thereby reducing the effect of pre-retirement wage uncertainty on lifetime utility. I first document the considerable heterogeneity in the degree to which actual pension formulas possess these properties and variation in the amount of wage uncertainty faced by workers. I then show that workers who face more wage uncertainty also tend to have pensions with these risk-compensating or risk-absorbing properties, suggesting that pensions are serving as a means of partially insuring workers against labor income uncertainty.

Thesis Supervisor: James M. Poterba
Title: Professor of Economics

Acknowledgements

Four years ago I made the decision to pursue a Ph.D. in economics because I was shaken by the encroachment of the federal government into an ever widening circle of economic activity. Only through careful study of the public finance activities of government institutions would I be equipped to propose worthwhile alternatives. Before presenting the results of my analyses to date, I would like to thank some of the many persons who have fostered my interest in economics or supported me in my academic endeavors.

Beyond all others, my parents, Marilyn and Gary, instilled in me a belief that I could accomplish anything with intellect and perseverance. With the hope that I will eventually fulfill their high expectations, I dedicate this thesis to them. I also thank my brother Matthew and my grandparents for their love and support over the past twenty-three years.

While attending college at Harvard, I found in my first economics professor, Martin Feldstein, a kindred spirit in the belief that one important element of a more prosperous economy would be more limited yet thoughtful social insurance programs. Beginning with his advising of my undergraduate thesis and continuing through research assistantships and other collaborations, I have learned a great deal from him about the political economy of Social Security. I thank him for many years of inspiration and teaching.

During the summer after my second year in graduate school, I was extremely fortunate to have Christopher Carroll invite me to the Federal Reserve Board to be his research assistant. The first two chapters of this thesis represent the results of our continued work together. It has been a pleasure to work with one of the rising stars in the field, and I have benefitted greatly from his friendship and mentoring.

While at MIT, my research was advised by an extremely capable trio. In addition to being the finest instructor I have ever had, James Poterba was a prolific source of research ideas and suggestions. Most importantly, he kept my research on track by helping me to recognize the most important elements of work in progress. Jerry Hausman was an entertaining and thoughtful adviser on the econometric aspects of the thesis. His early work in the field of retirement was also a guide to my own research. Finally, David Wise provided considerable support for the research on retirement in the final three chapters. His uncompromising and continued interest in my work, his recent papers in the field of pensions and retirement, and the opportunity to be a part of the Aging group at the NBER were all invaluable in facilitating my research.

I was also fortunate to have superb classmates and colleagues at MIT and NBER. Special thanks go to Mark McClellan, Jonathan Skinner, Robin Lumsdaine, and Jonathan Gruber for innumerable research suggestions on each of the papers in this thesis. Kevin Berner, Brigitte Madrian, Bill Miracky, Neal Rappaport, and Kim Rueben were fellow commiserators through the peaks and valleys of graduate school, and I drew vitality from them. I especially thank Chad Jones for sharing my interests in economics at the graduate and undergraduate levels for the better part of the past seven years.

Finally, I would be sorely remiss if I did not recognize Anil Shrivastava, Joshua Young, and Kathryn Mostoller for their bountiful friendship and patience with an "imperfect perfectionist" during my years of thesis research. I would also like to express my appreciation to the National Science Foundation, the Lynde and Harry Bradley Foundation, and the National Institute on Aging for providing me with fellowships during various parts of the last four years.

To Marilyn and Gary

Table of Contents

Preface 7

1 The Nature of Precautionary Wealth (joint with Christopher Carroll) 12

2 How Important Is Precautionary Saving? (joint with Christopher Carroll) 66

3 Retirement Incentives in the 1983 *Pension Provider Survey* 110

4 The Joint Effect of Social Security and Pensions on the Timing of Retirement: Some New Evidence 149

5 Wage Risk Compensation Through Employer-Provided Pensions 191

Preface

Over the last thirty years the federal government's role in the provision of social insurance has grown remarkably. Whereas nominal gross domestic product and personal income grew at annual rates of 8.2 and 8.4 percent between 1961 and 1991, government transfers to persons grew at an annual rate exceeding 11 percent over this period. Government transfers now stand at 13.5 percent of GDP. Most of both the level and the trend are due to growth in the Social Security program, which is comprised of Old-Age and Survivors Insurance, Disability Insurance, and Health Insurance policies.[1] The ongoing debates about universal health care programs do not suggest a change in this trend.

Unlike a pure public good such as national defense, however, both public and private institutions typically exist to provide social insurance. Designing and implementing efficient government programs requires a detailed knowledge of not only the effects of those programs in isolation but also as they interact with existing private insurance mechanisms. Similarly, evaluating the welfare consequences of the rapid growth of government provided social insurance requires an understanding of individual behavior in the presence of uncertainty. In this dissertation, I consider several questions which bear directly on either or both of these two issues.

The first two chapters investigate the nature and magnitude of precautionary saving against labor income uncertainty. Several components of government transfer

---

[1]Other government transfers include unemployment insurance, government retirement benefits, veteran benefits, and Aid to Families with Dependent Children. All figures are from Tables B-1 and B-23 in *Economic Report of the President 1993*. Washington, DC: United States Government Printing Office.

programs are based on insuring consumption against uncertainty in labor income. The most obvious include unemployment benefits and AFDC, but progressive income tax rates or schedules for Social Security benefits are also advocated for this reason.

In the first paper, we demonstrate that wealth holdings are significantly increased in the presence of both permanent and transitory shocks to income. This provides some of the first empirical evidence that households do engage in precautionary saving. We also show that the relative elasticities of wealth with respect to the variance of permanent and transitory shocks to income are consistent with a "buffer-stock" model of saving in which households accumulate small but positive stocks of wealth to insure consumption against fluctuations in income. Households in such a model are both impatient, in the sense that in the absence of uncertainty they would wish to borrow against future income to finance current consumption, and prudent, in the sense that uncertainty about future income paths causes them to engage in precautionary saving. In contrast to certainty equivalence models, which have been the standard bearer in macroeconomics, a buffer stock model predicts large effects on saving rates of insuring even transitory shocks to labor income. Similarly, the standard models of consumption in public finance typically ignore the consequences of both income uncertainty and buffer-stock behavior on the effects of policy changes.

The second paper adopts the buffer stock model as a representation of consumer behavior in order to estimate the fraction of aggregate wealth that results from precautionary saving against labor income uncertainty. This paper incorporates several improvements to the measures of uncertainty and the specification of the relationship between wealth and uncertainty made possible by the adoption of the buffer stock model. The finding of a significant effect of income uncertainty on the level of wealth is reconfirmed, and simulations reveal that between 25 and 40 percent of aggregate wealth is due to the presence of income uncertainty beyond the minimum observed in the sample of workers. We also find that uncertainty is responsible for relatively more of the wealth holdings in the tails of the permanent income distribution than in the middle. This suggests that government programs which seek to provide insurance against income uncertainty should be tailored differently to those with different earnings capacities.

The emphasis then shifts from the reasons why social insurance might be desired to the effects of both private and public insurance institutions on individual behavior. In particular, each of the next three papers utilizes a dataset which links the precise details of a workers' pension plans to his or her household data in a nationally representative household survey of wealth. The third paper documents the considerable heterogeneity in both the generosity of pension plans and the accrual of pension wealth as the worker ages. In particular, most pension plans are shown to have discrete jumps in the accrual rate of pension wealth at specific ages. The magnitude of these jumps and their timing are shown to be dependent on the worker's age, tenure, and wage level. The pervasiveness of particular types of pensions, such as those that are integrated with Social Security, is also discussed.

Using this unique dataset and the variation in retirement incentives it allows to be identified, the fourth paper examines the effect of both Social Security and employer pensions on the timing of retirement. There is a large literature that attempts to relate the decline in labor force participation of older workers since the early 1970s to the contemporaneous growth in Social Security. The general results are that although there is a statistically significant relationship between the generosity of Social Security and the probability of retirement, the size of the effect accounts for at most a third of the observed decline. I demonstrate that although formally including pensions in this analysis does not fundamentally change this conclusion, pension benefits have an effect on retirement that is roughly the same size. A comparison of policy simulations that do and do not make use of the pension plan details reveals that previous studies have missed more in the timing than the magnitude of the economic response to Social Security changes. When the incentive effects of pensions are formally modeled, much of the change in retirement that results from Social Security adjustments occurs before Social Security benefits can actually be received. This suggests that one way pensions may be influencing retirement is by relaxing liquidity constraints for middle age workers.

Finally, the interaction between income uncertainty and pension formulas is addressed in the fifth paper. I show that various pension plan features--including the integration of Social Security provisions into the pension formula--can provide insurance against uncertainty in within-job wages. Using the methods from the first two papers, I calculate measures of within-job wage uncertainty and show that in some sectors of the economy, workers who face greater income uncertainty disproportionately have pensions that either compensate them in expected value for bearing it or filter it out of their retirement income. This finding establishes the existence of private social insurance arrangements that should be accounted for in the design of efficient government provided insurance.

Chapter One

The Nature of Precautionary Wealth

I. Introduction

Recent advances in the theory of precautionary saving, most notably by Kimball (1990a, 1990b, 1991) and Zeldes (1989), have sparked interest in whether precautionary saving is an empirically important phenomenon. This paper uses the *Panel Study of Income Dynamics* (PSID) to provide some of the first direct evidence that wealth is higher for consumers with greater income uncertainty. We also show, however, that the apparent pattern of precautionary wealth holding is not consistent with a standard, unconstrained life-cycle maximization problem with patient consumers. Instead, precautionary saving seems to be of the buffer stock variety described in Deaton (1991) or Carroll (1992a), in which consumers seek to buffer current consumption against near-term fluctuations in income and are far less concerned about uncertainty in lifetime (or permanent) income than in the standard model.

Our empirical approach is to construct estimates of income uncertainty using the PSID and then to investigate the relationship between uncertainty and wealth holding. We decompose income uncertainty into a variance of shocks to permanent (lifetime) income and a variance of shocks to transitory income. To be clear, an example of a shock to permanent income would be the income shock experienced by a unionized steel worker in the late 1970s, whose expected future lifetime income was permanently reduced by the restructuring and de-unionization of the U.S. steel industry in that period. An example of a transitory shock might be lottery winnings, a bad crop for a farmer, or any other change in income due to factors not expected to persist indefinitely. Controlling for simple demographic effects and for the level of income, our empirical results indicate that wealth holdings depend importantly on the degree of both transitory and permanent income uncertainty.

We then use the empirical results to discriminate between theories of optimal intertemporal consumption decisions. To address this issue, we use numerical techniques to solve a standard non-liquidity-constrained version of the life cycle model under a range of assumptions about the variances of transitory and permanent shocks to income. The fundamental result is that predicted wealth holdings are *extremely* sensitive to the degree of uncertainty in permanent income. The model predicts a regression coefficient on the variance of permanent shocks that is over 30 times greater than the typical coefficients estimated in the empirical section.

Deaton (1991) and Carroll (1992a) showed that if consumers are more impatient[1] than is usually assumed and do not borrow, they will engage in "buffer-stock" saving, holding a relatively small stock of assets for the purpose of insulating consumption against near-term shocks to income. We therefore solve a buffer-stock model and show that, in contrast with the unconstrained model with patient consumers, the buffer-stock model's theoretical predictions are roughly consistent with our empirical estimates.

The paper is organized as follows. In Section II we discuss the various methods that others have used to study precautionary saving and the relative merits of our approach. In Section III, we develop our technique for estimating the magnitude of permanent and transitory shocks to income and present these estimates by occupation and

[1]The word "impatient" is something of a misnomer here; all that is really needed is that, if income were certain, consumers would wish to borrow. Moreover, similar results can be achieved even with a low discount rate as long as future income is expected to be higher than current income, thus generating a desire to borrow even in "patient" consumers.

education groups. The econometric results relating wealth holdings to income uncertainty are given in Section IV. We find consistent evidence that both permanent and transitory shocks to income increase the level of wealth; this demonstrates that precautionary saving is an empirically relevant phenomenon. In Section V we numerically solve a standard life cycle model of consumption under a broad range of assumptions about the magnitude of permanent and transitory shocks to income and show that it is inconsistent with the empirical results in Section IV. The analogous comparisons for the buffer stock model in Section VI are then shown to be consistent with these empirical results. Section VII suggests directions for further research and concludes.

II. Empirical Evidence on the Existence of Precautionary Saving

Existing empirical evidence on the nature and magnitude of precautionary saving is scant and contradictory. In a widely cited article, Skinner (1988) finds the peculiar result that saving is lower for groups thought to have higher income uncertainty. Guiso, Jappelli, and Terlizzese (1991), using a data set in which consumers reported an expected distribution of income for the next year, find that consumption is only slightly lower and wealth is only slightly higher for consumers with a greater subjective variance of next year's income. In contrast, Dardanoni (1991) examines consumption and income data for British households and finds that average consumption across occupational and industrial groups is negatively related to the intragroup variance of income. He estimates that around 60 percent of saving is due to precautionary motives. Carroll (1992b) found that future income uncertainty is highly statistically significant and quantitatively important in regressions of current consumption on current income, future income, and uncertainty; his estimates suggest that a one standard deviation increase in income uncertainty decreases consumption by at least 8 percent.

An advantage of our approach over most of these previous studies is that all but Guiso, Jappelli, and Terlizzese relate uncertainty to current consumption rather than to the stock of wealth. The appropriate response to greater income uncertainty is to hold more wealth, not necessarily to depress the current period's consumption. In the steady state of the buffer-stock model mentioned above, for instance, average consumption will equal average income for both low and high uncertainty consumers, thus maintaining the buffer stock at the optimal level for both groups. Therefore, across consumers already holding the optimal buffer stock, there will be *no* apparent relation between current consumption and the uncertainty of income. However, consumers with greater uncertainty will desire a greater buffer stock, so there *will* be a relationship between the level of wealth and the degree of uncertainty.[2]

Measurement issues also militate for the use of wealth rather than consumption as the variable to be explained. The stock of wealth represents the household's response to all current and previous shocks to income, while current consumption may be buffeted by a host of transitory factors that have little or nothing to do with fundamental saving decisions. We therefore suspect that current consumption is likely to be a much noisier indicator of precautionary behavior than wealth holding.

---

[2]Until the optimal buffer stock is achieved there will be a relation between consumption and uncertainty--the consumer facing higher uncertainty will initially have to depress consumption more in order to build up the larger stock of wealth. But given the impossibility of determining which consumers are at the optimal buffer stock, the point remains that the relation between consumption and uncertainty is not clear.

A stronger objection to the previous studies is that none used an entirely appropriate measure of income uncertainty. The models used by Guiso, Japelli, and Terlizzese (1991) and Dardanoni (1991) imply that current consumption should depend on the uncertainty in lifetime income. However, Guiso, Japelli, and Terlizzese use the variance of households' own forecasts of income *one year ahead*; they impose strong assumptions to generate an estimate of lifetime uncertainty from this estimate of next year's variance. If the assumptions are incorrect, their measure of lifetime uncertainty could be grievously biased. Logically, there is no necessary relation between the variance of transitory shocks and the variance of lifetime shocks, so there is no necessary relation between their measure and the economically correct measure. Dardanoni's (1991) measure of uncertainty is even more troubling: he uses within-group *cross-sectional* dispersion of income as the estimate for expected variance of lifetime income. There is not necessarily any relationship *at all* between this measure of dispersion and true uncertainty. In fact, for the standard set of eight occupation groups we use, the sample correlation in our PSID sample between our estimates of permanent uncertainty (which is essentially the measure required by Dardanoni's theory) and the group cross-sectional variances (the measure Dardanoni actually uses) is -0.47. Carroll (1992b) and Carroll and Samwick (1993) use measures of uncertainty derived, as ours are, from the PSID but do not decompose the shocks into transitory and permanent components.

Using the PSID we are able to make direct estimates of the variance of innovations to permanent income--the theoretically correct measure of uncertainty for the models used by Dardanoni and Guiso, Jappelli, and Terlizzeze. Furthermore, we can make a distinction between innovations to permanent income and transitory shocks to income. In fact, as Sections V and VI of the paper demonstrate, differences between the responses to transitory and to permanent uncertainty provide an important clue about what kind of model of precautionary saving can explain the data.

III. Estimating the Variance of Transitory and Permanent Shocks

The *Panel Study of Income Dynamics* now contains income data for the years 1968 to 1989 for around 8,000 households. The PSID also contains data on household wealth in 1984. The goal of the empirical estimation below is to determine how much of the differences in wealth across households in 1984 can be explained by differences in income uncertainty.

The most natural way to link the PSID data to a model of precautionary saving is to identify the PSID "household" as the decision unit and to examine variability in household non-capital income. Changes in household status, such as divorce or marriage, present a conceptual problem because it is difficult to know how to appropriately compare post-change individual income to pre-change household income (or vice versa). To avoid these issues we restrict the sample to households which remained intact over the chosen sample period. According to our definition, a household remains intact whenever the same person is the head of household each year and the spouse, if present, is always the same person. We also impose a variety of other restrictions, which are described in detail in the Appendix. Most of these restrictions tend to reduce the amount of income variability in the sample, so that our results are not driven by a few outliers but by the general tendency of the data. A final restriction is that we use only data from 1981-1987, on the view that income uncertainty measures constructed using data from these years will correspond reasonably well to the true uncertainty for the year in which wealth is measured.[3]

The income measure considered, total family noncapital income, includes labor income of the head, spouse, and other family members; disability payments, welfare payments, and other forms of transfer income, including food stamps and in-kind transfers; unemployment insurance and social security payments; and almost all other kinds of income except direct capital income from assets like stocks, bonds, and real estate. For our purposes this measure is a considerable improvement over the measure most commonly used in the labor literature, labor income of the household head, because the head's labor income may be implicitly insured by the spouse's income, by unemployment and disability benefits, and by a variety of other sources of income. Discretionary precautionary wealth holding should be motivated by the degree of income uncertainty *after* all these income insurance arrangements are taken into account.

The assumptions necessary to decompose household income into transitory and permanent components are uncomplicated.[4] The log of permanent income yp_t is assumed to follow a random walk with drift:

---

[3]We also performed the analysis using data from 1976-1987, and the results were not generally affected by this change.

[4]Our methodology here is a generalization of that of Hall and Mishkin (1982).

$$yp_t = g_t + yp_{t-1} + \eta_t \qquad (1)$$

where g_t represents predictable growth due to life cycle aging and due to overall productivity growth and η_t is the shock to permanent income in period t.

Current income is given by permanent income plus a transitory error term:

$$y_t = yp_t + \epsilon_t \qquad (2)$$

The errors ϵ and η are assumed to be white noise and uncorrelated with each other at all leads and lags. In order to remove the predictable component of income changes over time, g_t, we normalize observed income by dividing it by the predicted value from a regression of income on age, occupation, education, household demographic variables, and interaction terms.[5] The predicted values are also adjusted for economy-wide growth in income so that the average normalized income variable shows no trend over the sample period. Thus, equation (1) can be rewritten without the drift term g_t as:

$$yp_t = yp_{t-1} + \eta_t \qquad (1')$$

Define a d-year income difference as

$$\begin{aligned} r_d &= y_{t+d} - y_t \\ &= yp_{t+d} + \epsilon_{t+d} - yp_t - \epsilon_t \end{aligned} \qquad (3)$$

where the second equality is obtained by substituting (2) into the first equality. Substituting (1') into (3) recursively yields

[5]Other sensible methods of removing the predictable changes in income had little effect on measures of uncertainty presented below.

(4) $$r_d = \left\{\eta_{t+1} + \eta_{t+2} + \ldots + \eta_{t+d}\right\} + \epsilon_{t+d} - \epsilon_t$$

Finally, we obtain the d-year variance as the second moment of the right hand side of equation (4):

(5) $$Var(r_d) = d\cdot\sigma_\eta^2 + 2\cdot\sigma_\epsilon^2$$

where σ_η^2 and σ_ϵ^2 are the variances of the permanent and transitory shocks to income, respectively. Note that equation (5) makes use of the assumption that permanent and transitory shocks are white noise and uncorrelated with each other at all lags. We estimate $Var(r_d)$ for each household i by $v_{id} = r_{id}^2$; v_{id} is an unbiased estimator of $Var(r_{id})$ if the mean of r_{id} is zero, which corresponds to an assumption that there are no individual-specific growth rates for income (other than those predictable by occupation, education, and other personal characteristics). This assumption was tested on the PSID by MaCurdy (1982), who found it to be a good characterization of the data.

For each household, we can use any two v_d's of different lengths to solve for σ_η^2 and σ_ϵ^2. For example, for any $d > 1$, one way of estimating σ_η^2 for a given household is simply to take (we suppress the subscript i for clearer notation):

(6) $$s_\eta^2 = v_d - v_{d-1}$$

This works because

(7) $$\begin{aligned} E(s_\eta^2) &= Var(r_d) - Var(r_{d-1}) \\ &= d\cdot\sigma_\eta^2 + 2\cdot\sigma_\epsilon^2 - \left[(d-1)\cdot\sigma_\eta^2 + 2\cdot\sigma_\epsilon^2\right] \\ &= \sigma_\eta^2 \end{aligned}$$

We can now set

(8) $$s_\epsilon^2 = \frac{v_{d-1} - (d-1)\cdot s_\eta^2}{2}$$

and it is easy to show that $E(s_\epsilon^2) = \sigma_\epsilon^2$. Our actual procedure for estimating σ_ϵ^2 and σ_η^2 is a generalization of this method; rather than only using two differences v_d and v_{d-1}, we use a least squares projection to optimally combine all available differences. We use only those differences that make the procedure robust to MA(2) serial correlation in the transitory component, which limits the differences used to those spanning three or more years.[6] The methodology for this projection and the robustness to MA serial correlation are developed in the Appendix.

A brief justification for the foregoing method of decomposing income shocks is in order. Quah (1990) shows that with a finite data set it is impossible to distinguish econometrically between a nonstationary series and a stationary series in which shocks simply have high persistence. We have more than econometrics to guide us, however. If we were to characterize all shocks to income as stationary, the interpretation would be that the displaced steel worker of the late 1970s would *eventually* regain a job in which his wages were as high as before. This scenario is unlikely to occur; instead, it seems much more plausible that there are sometimes shocks to individuals' earnings profiles that will simply never be reversed. Hence, we believe that the case for allowing a nonstationary component in individuals' earnings process is compelling.

---

[6]Several studies of the time-series properties of labor earnings, including MaCurdy (1982) and Abowd and Card (1989), have asserted that there is no evidence of serial correlation beyond lag 2 in the transitory component of labor income.

The case for allowing a transitory component is equally compelling. Lottery winnings, bad crops, gifts, and other transitory shocks to income are phenomena that undeniably exist and are sometimes completely unrelated to innovations in permanent income. Some previous work has attempted to reduce the dimensionality of the income process by assuming that there is only one kind of shock which has both transitory and permanent elements, but this assumption is both restrictive and, for our purposes, unnecessary.[7]

Given that our decomposition procedure is robust to MA(2) correlation in the transitory innovation component, the main circumstance under which our methodology would be invalid is if transitory innovations to income (transitory in the sense of having no implications for income far in the future) had effects that lasted three years or longer. But MaCurdy (1982) and others who have examined the labor income process found no evidence of serial correlation beyond MA(2) in the income innovations. Furthermore, there are few, if any, examples of shocks that are explicitly transitory in nature but have effects that last longer than three years. Certainly lottery winnings, crop failures, and most other examples of transitory shocks that spring to mind would not have implications for income three years hence.

Table 1 presents our estimates of the variance of transitory and permanent income shocks by group for the eight occupational groups and six educational groups we distinguish in the PSID. The patterns in the table are unsurprising: Farmers and the

[7]Another reason to decompose income shocks into separate transitory and permanent components is that such a procedure facilitates the solution of the theoretical models Sections V and VI. The models would be far less tractable if the income process followed a stationary but persistent income process.

Self-Employed have by far the greatest amounts of transitory income uncertainty; transitory uncertainty decreases with education; and there are greater differences across groups in the degree of transitory uncertainty than in the degree of permanent uncertainty. F-tests overwhelmingly reject the proposition that either transitory or permanent uncertainty is the same across people in different occupational or educational groups (with the exception of permanent uncertainty across educational groups); in other words, the degree of income uncertainty is at least partly predictable on the basis of occupation and education.[8]

The estimates in Table 1 must be treated with some caution, since the underlying income data are undoubtedly subject to measurement error. For our purposes, however, this may not cause important problems. Suppose measurement error is of the commonly specified type:

$$y_t^r = y_t^{true} + u_t \tag{9}$$

where y_t^r signifies reported income and y_t^{true} signifies true income. Then if u_t is white noise and the variance of the error term, $Var(u_t) = \sigma_u^2$, is uncorrelated with personal

[8]It is important that occupation and education have explanatory power for uncertainty because these two variables and their interactions with age will subsequently be used as instruments for income uncertainty.

characteristics, measurement error will have no effect on our estimate of permanent uncertainty. This can be seen by noting that:

$$
\begin{aligned}
(10) \qquad E(s_\eta^2) &= Var(r_d) - Var(r_{d-1}) \\
&= d\cdot\sigma_\eta^2 + 2\cdot(\sigma_\epsilon^2 + \sigma_u^2) - \left[(d-1)\cdot\sigma_\eta^2 + 2\cdot(\sigma_\epsilon^2 + \sigma_u^2)\right] \\
&= \sigma_\eta^2
\end{aligned}
$$

Equation (10) makes clear that measurement error in current income is treated by our procedure in the same way as transitory income shocks. Thus, the proof in the Appendix that our procedure is robust to MA(2) serial correlation in the transitory shock also demonstrates that, in order for the estimate of permanent variance to be affected, the measurement error term would have to be more persistent than MA(2). Naturally, the estimate of the transitory variance would be increased:

$$
(11) \qquad E(s_\epsilon^2) = \sigma_\epsilon^2 + \sigma_u^2
$$

but if σ_u^2 is the same for everyone the only effect of measurement error on our estimates would be to raise the estimate of everyone's transitory variance by the same amount. The impact of such measurement error on the coefficients of the regressions we run is explored in the Appendix.

One unfavorable property of our estimates of permanent and transitory variance is that they are negatively correlated across households. That is, households which are estimated to have high permanent uncertainty tend to have low estimates of transitory uncertainty. Although it may be the case that the true income process has some negative correlation, this can also be induced by the technique used to estimate the two components. Suppose now that v_{id} is estimated with error, $v_{id} = v_d + z_i$, while v_{id-1} is measured without error.[9] Then we will obtain (omitting the household subscript):

(12)
$$\begin{aligned} s_\eta^2 &= v_d + z - v_{d-1} \\ &= \sigma_\eta^2 + z \end{aligned}$$

(13)
$$\begin{aligned} s_\epsilon^2 &= \frac{\left[v_{d-1} - (d-1)\cdot s_\eta^2\right]}{2} \\ &= \frac{\left[v_{d-1} - (d-1)\cdot(\sigma_\eta^2 + z)\right]}{2} \\ &= \sigma_\epsilon^2 - \frac{(d-1)\cdot z}{2} \end{aligned}$$

So even if $\mathrm{Cov}(\sigma_\eta^2, \sigma_\epsilon^2) = 0$, $\mathrm{Cov}(s_\eta^2, s_\epsilon^2) = -(d-1)\sigma_z^2/2 < 0$. Although this is a severe condition at the level of the individual household, the Appendix shows that the negative correlation diminishes when the correlation is taken among aggregates of households such as occupational or educational groups. Given a large enough sample for each group, the spurious negative correlation would disappear completely as the estimates of the v's approached their true values. For the groupings presented in Table 1 the correlation coefficient is -0.2, which is much closer to zero than the -0.8 which obtains at the individual level.

IV. How Does Wealth Depend on Uncertainty?

We will consider three measures of wealth in our empirical estimates. Very liquid assets (VLA) are those assets which could be liquidated immediately: bank

[9] Note that this is a different issue from the income measurement error just discussed. Measurement error in income would cause the same measurement error in v_{id} and v_{id-1}, and would therefore cancel out; here we are supposing that v_{id} and v_{id-1} are mismeasured by different amounts.

account balances, money market funds, CDs, government savings bonds, and directly owned stocks and mutual funds. Non-housing, non-business wealth (NHNBW) adds to VLA the net value of all assets and liabilities not related to the primary residence or personally owned businesses. Such assets include non-government bonds, the cash value of whole life insurance, cars and other vehicles, secondary real estate, and other investments. In addition to loans or mortgages on any of these assets, this NHNBW also deducts the balances on credit cards, student loans, medical and legal bills, and loans from relatives. Most of the components of this measure could be liquidated (or at least have the amount of equity altered) within a matter of weeks or months. Finally, total net worth (NW) adds to NHNBW what are generally the most illiquid assets owned by households: equity in the primary residence and the net value of personally owned businesses.

It is necessary to consider several measures of wealth because, for tractability, our theoretical model has only one asset and because we do not yet have an adequate model of portfolio choice across assets with differing liquidity characteristics in the presence of precautionary saving motives. If we chose to focus only on very liquid assets as our measure of wealth, and consumers, for example, put their precautionary savings into less liquid assets, we could erroneously conclude that no precautionary saving existed. This possibility suggests using the most inclusive wealth measure, total net worth, as the dependent variable. However, some assets may be held primarily for reasons that have nothing to do with uncertainty. Home equity, for instance, may be accumulated because the tax advantages to homeownership are strong. The large transactions costs associated with readjusting the stock of housing suggest that housing assets would be an inconvenient vehicle for buffering consumption against income fluctuations. These arguments suggest that if we are searching for wealth that is held for precautionary reasons we might wish to exclude the most illiquid forms of wealth. Based on this rationale, we believe that the intermediate measure of non-housing, non-business wealth is the most appropriate measure for examining the impact of uncertainty on wealth holdings.

Means and medians of the three measures of wealth, and of their ratios to estimated permanent income,[10] are presented in Table 2, along with comparable measures constructed using the *Survey of Consumer Finances 1983* of the Federal Reserve Board. The SCF results are presented in order to verify that our PSID sample, even after making all the various sample restrictions, is approximately representative of true wealth holding behavior in the U.S. in 1983. If the two surveys had substantially different means and distributions for wealth, then there would be reason to believe that the sample might not be representative of the population or that the PSID measures of wealth were flawed. Instead, the results in the PSID are remarkably close to those from the SCF, thereby lending confidence to the PSID measures of wealth. This confirms the basic results of Curtin, Juster, and Morgan (1989) who carefully compared the SCF and the PSID and concluded that the PSID provided nearly as good a measure of wealth as the SCF for most of the population (although the SCF is far superior for examining the

[10]Our definition of permanent income here is simply the income that is predicted by our wage equation. This corresponds approximately to the definition of yp_t in equation (1).

behavior of the wealthiest families).[11] Additionally, our predicted income measure (in real 1982 dollars) was very close in the two samples: 32,900 in the SCF and 30,600 in the PSID. Comparisons of age, family size, and other demographic variables were also reasonably close in the two surveys.

With estimates of wealth and of the components of uncertainty in hand, the next step is to regress household wealth on the measures of uncertainty by household. The proper econometric technique for these regressions is instrumental variables. It is clear that at the level of the individual household our measures of the variance of transitory and permanent shocks to income will be subject to gross measurement error: equations (7) and (8) show that *in expectation* our measures of uncertainty correspond to the true measures of uncertainty, but for any *individual* household our measures are a noisy approximation to the actual uncertainty faced. In other words, for every individual household our estimates of the variances can be represented as the true variances plus some individual-specific error term. This is a standard example of the errors-in-variables problem, so that if we were to estimate an ordinary least squares regression of wealth on uncertainty, the estimated coefficients on the uncertainty term would be attenuated toward zero with the magnitude of the bias corresponding to the variance of the measurement error. Since at the individual household level the variance of the error must be considerable, the coefficients from an OLS regression would be inappropriate.[12] We

---

[11]Avery, Elliehausen, and Canner (1984a, 1984b) discuss the correspondence between the 1983 SCF and other sources of wealth data such as the *Flow of Funds* accounts.

[12]In the OLS regressions we estimated, the coefficients on the uncertainty terms were in fact smaller and less statistically significant than in the instrumental variables specification.

adopt the standard solution to such errors-in-variables problems, which is to estimate the equation using instrumental variables.

Because wealth is also influenced by factors other than uncertainty, the regressions include a variety of demographic variables as well as a measure of household permanent income. The regressions are of the form:

$$\frac{w}{y^p} = \alpha_0 + \alpha_1 s_\eta^{\ 2} + \alpha_2 s_\epsilon^{\ 2} + \alpha_3 y^p + X\beta + \varepsilon \tag{14}$$

where y^p is an estimate of permanent income and X contains age and other demographic variables to control for (predictable) life cycle effects on wealth.

In order to estimate this equation we must decide what variables to use as instruments and what variables to include in our demographic variables X. To achieve identification, the instrument set must include at least some variables that are not also in X. Our choice was to define X to include age, marital status, race, sex, and the number of children in the household, but to leave occupation and education dummies out of X so that they could provide identification for the income uncertainty terms. The necessity for making such choices makes it clear how unfortunate it is that we must estimate these equations using instrumental variables, since it is admittedly plausible that the wealth/income ratio might be correlated with some of the instruments not included in X, even absent a precautionary saving motive. Still, some identifying assumption must be made and we believe that using occupation and education to identify uncertainty is probably the most defensible identifying assumption. And if we estimate positive and significant coefficients on the uncertainty terms, it will still be correct to say that "those with higher income uncertainty tend to hold more wealth," controlling for age, marital status, race, sex, and children.

Turning again to equation (14), income is included as a regressor because empirically there seems to be a strong positive relationship between the wealth-to-permanent-income ratio and the level of permanent income; in other words, high permanent income people hold disproportionately great wealth. If this were because high permanent income consumers also have relatively greater income uncertainty it would do no harm to include income as a regressor along with the measures of uncertainty; in that case, income would get an insignificant coefficient and there would be a slight loss in the efficiency of the other estimates. However, it is plausible that high permanent income people simply save more, even independently of their degree of income uncertainty.

Furthermore, even if true wealth holdings were homogeneous with respect to income, it is likely that wealth, appropriately defined, is undermeasured for lower income households. For instance, in the *Survey of Consumer Finances 1986* most households indicated that, in times of emergency, they could borrow at least $3,000 from friends or relatives. This "precautionary liquidity" should lessen the need to hold actual assets in reserve against bad income shocks. If it is not accounted for, then the precautionary wealth of lower income households will be understated relative to that of higher income households. More generally, at lower levels of wealth unmeasured durable goods (television sets, VCR's, etc.) presumably constitute a greater fraction of total assets than for wealthy households.

Table 3a presents the results from instrumental variables regressions of the ratio (wealth/permanent income) on our measures of transitory and permanent income uncertainty, the level of permanent income, age, and other demographic variables using the entire sample of households. Focusing on the regressions which use all occupations (rows 1, 3, and 5), the coefficients on both the permanent and the transitory components of uncertainty (columns 2 and 3) are positive and significant (generally) at the 5 percent level in regressions using each type of wealth. These regressions indicate that, controlling for age, income, and demographic variables, higher uncertainty is associated with higher wealth holdings.

The coefficient on permanent variance is larger than that on transitory variance. This is what any reasonable precautionary saving model ought to predict, because an increase in the variance of permanent innovations has a vastly larger impact on the variability of the total future stream of income than an increase in the variance transitory innovations, which will eventually dissipate. It is important to note, however, that the coefficient on permanent variance is typically only about two to four times the size of the coefficient on transitory variance.

Another feature of these regressions is that the magnitude and significance of the coefficients increases as the measure of wealth becomes more comprehensive. This suggests that the reaction to greater uncertainty is to hold more of all forms of wealth, but that more illiquid forms of wealth are increased more. Although it is perhaps less intuitive to think of home equity as a buffer (at the margin) against uncertainty than to think of, say, stock ownership in such a way, the right question is presumably whether people consider the possibility of selling their homes and other illiquid assets in the event of an extremely bad income draw. It does not seem implausible that one purpose of accumulating home equity is as insurance against disaster. This is particularly feasible as an explanation for why the elderly do not decumulate housing equity (Venti and Wise (1990)) as the simple life cycle model suggests they should.

There is a venerable tradition in the precautionary saving literature of asserting that the self-employed and farmers may behave differently than other consumers, either because they are particularly risk-loving (Friedman (1957) and Skinner (1988)) or because measurement problems may be more severe for these groups, particularly for consumption and income (Skinner (1988)). This tradition has been motivated by an inability to find higher saving for these two groups even though intuition (and data) strongly suggest that they have greater income uncertainty. We do not find these arguments a compelling reason to challenge our regressions of wealth (which is not as susceptible to the measurement criticisms of Skinner and Friedman) on estimated uncertainty. If the self-employed and farmers face greater income uncertainty, they should hold more wealth. Nevertheless, we also present the results when these two groups are excluded in rows 2, 4, and 6 of Table 3a.

The results excluding the self-employed and farmers are substantially different: neither transitory nor permanent variance is useful in predicting the level of either very liquid assets or total net worth (rows 2 and 6). On the other hand, the variance of permanent income uncertainty is still significant in predicting NHNBW. Thus, for the most reasonable measure of wealth, the results are robust to the exclusion of the households with very high income uncertainty.

A reasonable criticism of Table 3a is that the results may be driven by the behavior of the very wealthiest households. If, for instance, families which inherited large amounts of wealth from previous generations also happen to choose occupations with a high degree of income uncertainty (such as self-employment or farming) there will be a positive association between wealth and uncertainty, but the causality will run from wealth to uncertainty rather than from uncertainty to wealth. It is unlikely that the principal reason that the very wealthy hold so much wealth is to insure themselves against future labor income risk.

Table 3b attempts to address this problem by excluding the wealthiest one percent of households from the sample, on the view that their wealth holding is motivated factors other than labor income uncertainty. The results change notably. Neither measure of uncertainty is significant in predicting holdings of very liquid assets, although the coefficients remain positive. For non-housing non-business wealth, both the transitory and permanent variances remain statistically significant at the 5 percent level, although the coefficients are around half of their values in Table 3a. Reassuringly, the coefficients are now virtually unchanged (although they are now significant only at the 10 percent level) when the self-employed and farmers are excluded from the regressions. Finally, for total net worth, the coefficients on both transitory and permanent variance fall somewhat, but both remain highly significant for the overall sample.[13] On balance, we

[13]Results were very similar when the top 5 percent, rather than the top 1 percent, were excluded.

believe that the results of Table 3b, and particularly the results for the intermediate measure of wealth, NHNBW, provide the most plausible and defensible estimates of the effects of uncertainty on wealth.[14]

In sum, the message of Table 3 can be stated as follows. Both kinds of income uncertainty matter separately in determining the amount of wealth households own. Permanent uncertainty matters more, both in the sense that the coefficients are generally more statistically significant and in the sense that they are larger; however, the transitory term is also usually statistically significant in the overall sample. Typically the coefficient on the permanent term is two to four times greater than the coefficient on the transitory term, and this statement is still approximately true even when the coefficients are measured imprecisely. That is, even though the coefficients are not statistically different from zero, the ratio of the coefficients is generally below 10 over the range of the 95 percent confidence intervals. The results are most robust for non-housing, non-business wealth, which includes primarily assets which could be liquidated within a matter of weeks or months. As we noted above, non-housing non-business wealth is the most appropriate place to look for precautionary saving, so it is not surprising to see that it is the measure that performs most credibly in the regressions.

---

[14]In subsequent work (Carroll and Samwick 1993), we are more systematic in guarding against misspecification. The equations estimated therein are in logs rather than levels, which tends to reduce the impact of outlying wealth observations on the parameter estimates. Rather than excluding occupations like farmers and self-employed, we shift the occupation terms from the instrument set to the independent variables in equation (14) so that the model passes formal specification tests. The qualitative empirical results are similar to those in Table 3. The cost of this finer treatment of the empirical model is the inability to focus separately on permanent and transitory shocks to income. Since the purpose here is to use different responses to permanent and transitory shocks to income to identify a plausible model of consumption and this identification can take place even with caveats regarding the estimates in Table 3, those readers concerned about the rough treatment of outliers here are encouraged to review the later paper as well.

V. Simulated Results of a Standard Life Cycle Model

In Section I we asserted that the empirical results just described are not consistent with a simple non-liquidity-constrained life cycle model because such a model predicts that wealth holdings should be very sensitive to permanent (lifetime) income uncertainty and relatively insensitive to transitory uncertainty. We now demonstrate this point by calibrating and solving a life cycle model under standard assumptions and examining the sensitivity of wealth holdings to transitory and permanent uncertainty. Wealth holdings are shown to be relatively insensitive to transitory uncertainty and extremely sensitive to permanent uncertainty.

In a standard life cycle optimization model, consumers are assumed to solve an optimization of the following form:

$$
\begin{aligned}
&\max_{C_t,\ldots,C_T} u(C_t) + E\sum_{i=t+1}^{T} \beta^{i-t}u(C_i) \\
(15)\qquad & s.t.\ \forall i\quad A_{i+1} = R\left[A_i + YL_i - C_i\right] \\
& \qquad\qquad YL_i = YLP_i \mathrm{E}_i \\
& \qquad\qquad YLP_{i+1} = G_{i+1} YLP_i \mathrm{N}_{i+1}
\end{aligned}
$$

in each period of life t, where C_t is current consumption, YL_t is current labor income, YLP_t is the level of permanent income, A_t is non-human wealth, R is the gross interest rate, β is the discount factor, G_t is the expected income growth factor, E_t is a multiplicative transitory shock, and N_t is the multiplicative shock to permanent income. Note that the specification of the stochastic process corresponds to that used to decompose income in Section III. For the simulations, the two shocks are assumed to be distributed lognormally as well as independently.

We assume a 50 year length of life (think of consumers beginning life at 25 and dying at 75) and set $R = \beta = 1$; a constant real interest rate of zero and no discounting of future utility. We further assume that the expected growth rate of income is zero throughout the whole of life; thus, there is no retirement and no need for retirement saving. Consumers start and end life with zero assets.

The specification of the utility function dictates the complexity and relevance of the solution. If utility is Constant Absolute Risk Aversion (CARA), it is possible to obtain an analytical solution to the problem. For this reason Guiso, Japelli, and Terlizzese (1991), Caballero (1991), and Dardanoni (1991) use CARA utility. However, there are at least three major objections to the assumption of CARA utility. First, under CARA utility the *optimal* consumption plan may include periods of negative consumption. It is hard to know how to interpret this when studying actual wealth holdings. Second, CARA utility predicts that a person consuming a million dollars a year would react in exactly the same way as a person consuming $10,000 a year to, say, a 50 percent chance of having consumption reduced by $5,000. Intuition suggests that the less wealthy consumer might react more strongly to such a risk. Finally, *ceteris paribus* CARA utility implies that wealthy people will save *less* as a fraction of their income than the non-wealthy. This also seems counter-intuitive.[15]

Instead of CARA utility we use a Constant Relative Risk Aversion (CRRA) utility function, $u(c) = c^{1-\rho}/(1-\rho)$. CRRA utility is not subject to any of the criticisms of the

[15]The last two shortcomings reflect the constant absolute prudence property of the CARA utility function defined by Kimball (1990a). Similarly, CRRA utility functions exhibit constant relative prudence, whereby the reaction of the two consumers will be proportional to their consumption. Such behavior seems more plausible.

CARA utility function. There is a vigorous debate in the macroeconomics literature about what constitutes a plausible value for the coefficient of relative risk aversion ρ, with empirical estimates generally falling well above 5. Many find these empirical estimates unreasonable, however, and argue for lower figures. For the simulations below we will assume $\rho = 3$; most of our essential points would be strengthened if the value of ρ were larger.

Under these assumptions, if there were no income uncertainty (i.e. $E_t = N_t = 1$) the optimal lifetime consumption profile would be flat (as in Modigliani's original life cycle model). Since income is assumed to have no trend growth ($G_t = 1$), if there were no uncertainty consumers in this model would neither borrow nor save, setting (constant) consumption equal to (constant) income over their entire lifetime. The assumptions for the standard model are dictated by the desire to have a base case where saving is zero under certainty so that the effects of uncertainty can be clearly illustrated.

When uncertainty is added to the income process, however, there is no closed-form, analytical solution to the optimal consumption problem under CRRA utility. It is therefore necessary to utilize numerical methods to solve the model. The solution method employed is the same as in Carroll (1992a) and similar to that of Deaton (1991): backwards recursion on the Euler equation from the end of life. The solution to a model of this sort consists of a series of consumption rules relating consumption and current cash-on-hand (wealth plus current income). There is a rule for the last period of life, one for the second to last period, and so on back to the first period. These rules indicate how much consumption is optimal for a given amount of cash-on-hand at a given age.

They do not directly reveal what level of cash-on-hand will prevail on average. Given a distribution of wealth at the beginning of any period and the distribution for income shocks, however, it is possible to use the consumption rule for that period to derive numerically the distribution for cash-on-hand in the next period. Starting with a wealth of zero in the first period of life, recursive application of this technique can generate the wealth distributions in successive periods up to the end of life.[16]

The results of performing such calculations for the first 15 years of the consumer's life, assuming zero assets initially, are shown in Figure 1. For our baseline parameters, we assume $\sigma_\eta = 0.075$ and $\sigma_\epsilon = 0.10$, or roughly half of the sample means of 0.14 and 0.21, in order to compensate for the likelihood that measurement error inflates the average estimates of income uncertainty. Assuming the consumer begins "life" at the age of 25, this should correspond to the distribution of wealth across life cycle consumers around 40 years old.

Table 4 summarizes the mean amount of wealth at age 40 predicted by the model under other uncertainty parameter values. It is immediately apparent from this table that the model predicts that wealth should be vastly more sensitive to the degree of uncertainty in permanent income than to the degree of uncertainty in transitory income. Table 5 reports the slope coefficients that would be estimated in a regression like those of Section IV if the standard model were representative of the population. The numbers in each row of the top panel of Table 5 are obtained by subtracting numbers in adjacent columns of that same row in Table 4 and dividing by the difference in the transitory

[16]Details of the solution method and the technique for generating the wealth distributions in each period are given in Carroll (1992a).

Figure 1

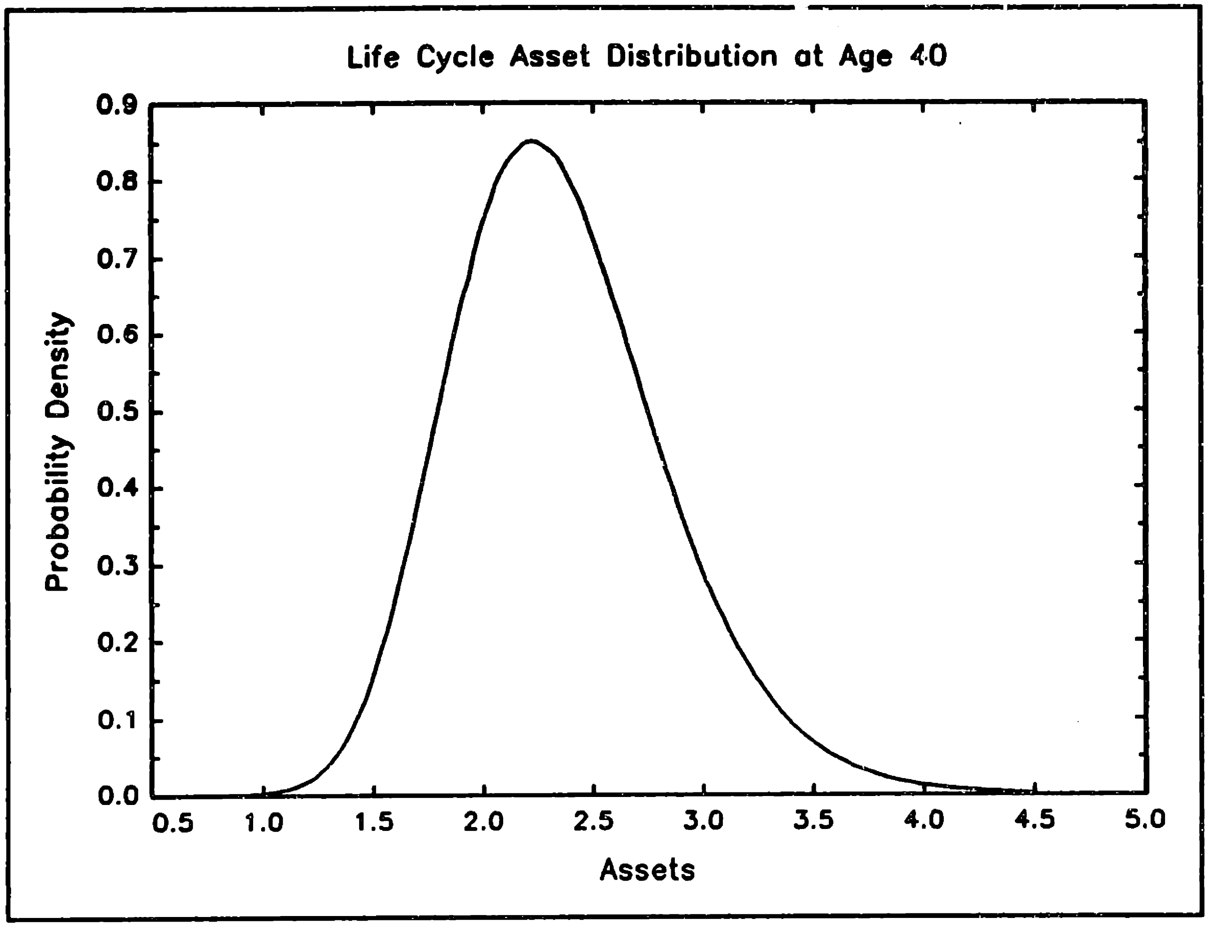

variances given at the top of those columns. For example, the figure 5.200 in the upper left corner of the table is equal to (0.392-0.353)/(0.01-0.0025). Similarly, the numbers in each row of the bottom panel of Table 5 are obtained by subtracting numbers in adjacent rows of the corresponding column in Table 4 and dividing by the difference in permanent variances given at the left of those rows. For example, the value 429.9 is equal to (1.159-0.353)/(0.0025-0.000625).

The predicted coefficients on the permanent component of uncertainty are about 30 times greater than the empirical coefficients estimated in Table 3, although the predicted coefficients for the transitory component are of the same magnitude as the empirical estimates. Furthermore, the coefficients on the permanent component are more than 100 times greater than the coefficients on the transitory component, compared to the ratio of two to four that is found in the empirical estimates. Thus, the life cycle model with the standard parameter values is deemed to be inconsistent with the empirical results of Section IV because the predicted coefficient on permanent variance is too high relative to its estimated value and the predicted ratio of the coefficients is much larger than that which is found in the data.

VI. Simulated Results of a Buffer Stock Model

Although the simple life cycle model just described appears to be a poor vehicle for explaining the precautionary saving behavior documented in Section IV, this same intertemporal optimization framework can generate behavior much closer to the empirical estimates if two previous assumptions are modified. Following Carroll (1992a), the first modification is to assume that consumers are impatient, so that (absent uncertainty) they would wish to borrow against future income in order to consume more today. We model this by choosing a discount rate of 10 percent instead of zero.[17] The second modification is to assume that, in addition to the usual transitory and permanent shocks, there is some probability in each period that income will go all the way to zero. Carroll (1992a) shows that this assumption is supported by the PSID data; he estimates the probability of one of these zero-income events to be between 0.3 and 0.6 percent per year, depending on sample restrictions. In this version of the model consumers engage in "buffer-stock" saving, holding a relatively small stock of assets to shield their consumption against either transitory or permanent shocks to income. We now show that

[17]Similar results could be achieved by leaving the discount rate at zero and assuming an income growth rate of around 3 percent.

the behavior implied by this "buffer-stock" version of the model is far more consistent with our basic empirical results than is the standard version of the model.

The presence of zero-income events has an effect similar to imposing liquidity constraints directly: it prevents consumers from borrowing against future income.[18] To understand why, consider a consumer in the second to last period of life deciding how much to consume. If such a consumer were to spend all her current income and entered the last period of life with zero assets, she might have to consume zero (if she were unlucky enough to experience a zero-income event in the last period). In a model with CRRA utility, consuming zero yields infinitely negative utility, so the consumer will never choose to enter the last period of life with zero assets. She will also never enter the second-to-last period with zero assets, for the same reason. Recursive application of this logic demonstrates that a consumer facing the possibility of a zero income event in each future period will never borrow against future income or even allow assets to fall all the way to zero.

Although the model might be more conceptually appealing if direct liquidity constraints rather than zero-income events generated the no borrowing behavior, the consumption functions are very similar in the two cases, with the primary effect of the zero-income events being to smooth the consumption function around the kink caused by a liquidity constraint. As the probability of zero-income events goes to zero, this model approaches a model with liquidity constraints directly imposed. This can be seen graphically in Figure 2. The smoothed consumption functions correspond to the optimal

[18]Zeldes (1989) made this point as well.

Figure 2

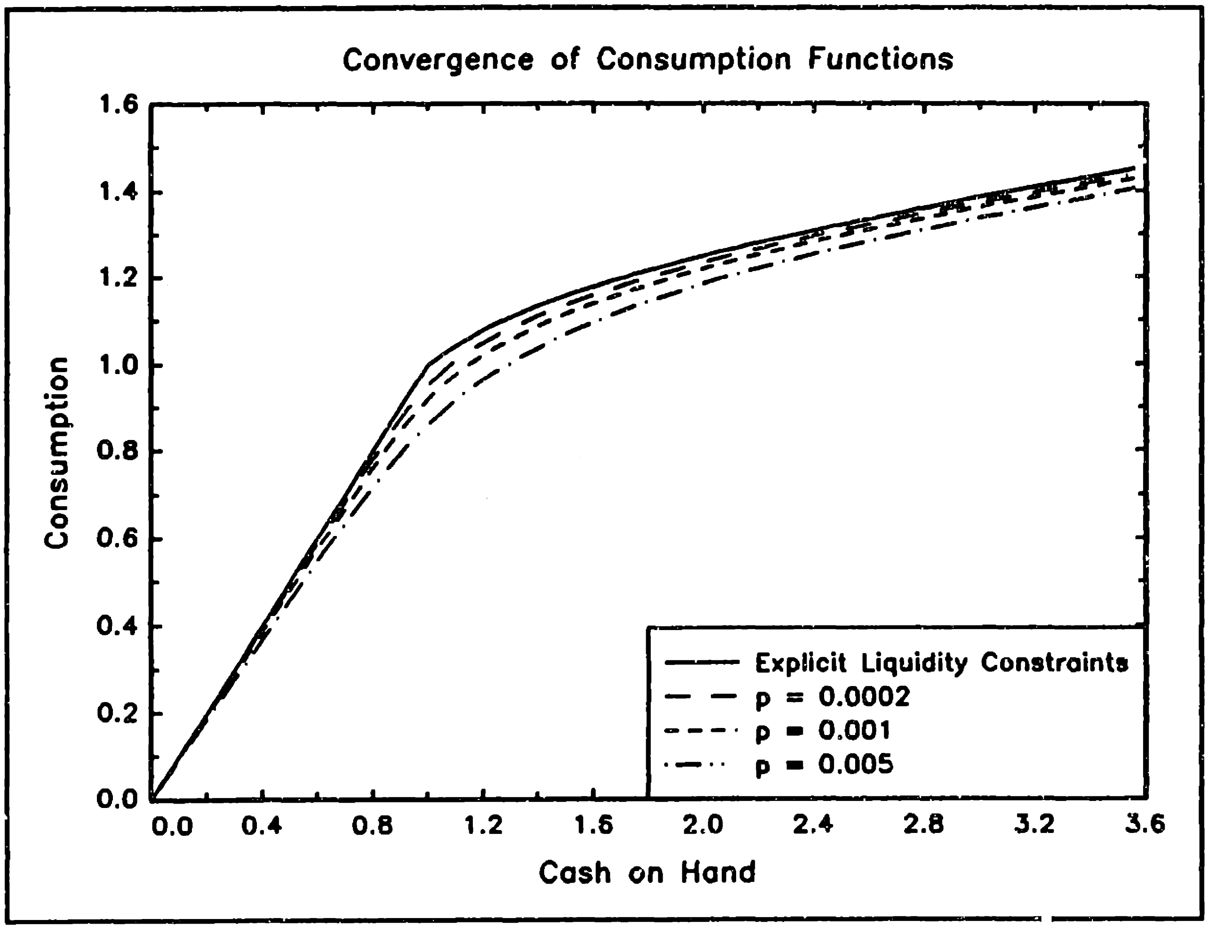

rule as the probability of zero-income events decreases from 0.5 to 0.1 to 0.02 percent. As the probability decreases, the optimal consumption rule approaches the kinked consumption rule which imposes the liquidity constraint directly. The model which imposes liquidity constraints directly is computationally intractable for our purposes (due to the lack of differentiability of the consumption function at the kink), so we use instead the model with a small probability of zero-income events. We choose a value of 0.3 percent for our simulations, the lowest value consistent with the PSID data in Carroll (1992a).

The behavior that arises from the model can be described as "buffer-stock" saving. Consumers have a target wealth-to-income ratio which they desire in order to have a sufficient cushion to protect consumption from income uncertainty. When wealth is below the target, the precautionary saving motive will be strong enough that wealth will be expected to grow; if wealth is above the target, consumers' impatience outweighs their prudence and their wealth will be expected to shrink.

In the long run, the actual level of wealth will be smoothly distributed around this target ratio, with a distribution that can be calculated numerically. Figure 3 shows the steady-state distribution of the wealth-to-income ratio under our baseline parameters. This figure is the buffer-stock analogue to Figure 1 for the unconstrained life cycle model. The most notable difference between the two figures is that the buffer stock model implies a much more restricted range for the distribution of assets: they tend to be fairly tightly distributed between 0.2 and 0.8 times permanent income. In the life cycle model the probability density function for assets was far more diffuse.

The true test of the buffer stock model is whether it can do a better job explaining the empirical results from Section IV than the standard model. Table 6 presents the average amount of wealth holdings predicted by the buffer stock model at various combinations of transitory and permanent uncertainty. The first thing to note about this table is that, for the values of the discount rate, interest rate, and risk aversion we have chosen, the predicted ratio of wealth-to-income is roughly the same as the median for NHNBW in Table 2. Comparing the numbers in Table 6 with those in Table 4 (the analogous table for the unconstrained problem), it is clear that the buffer stock model

Figure 3

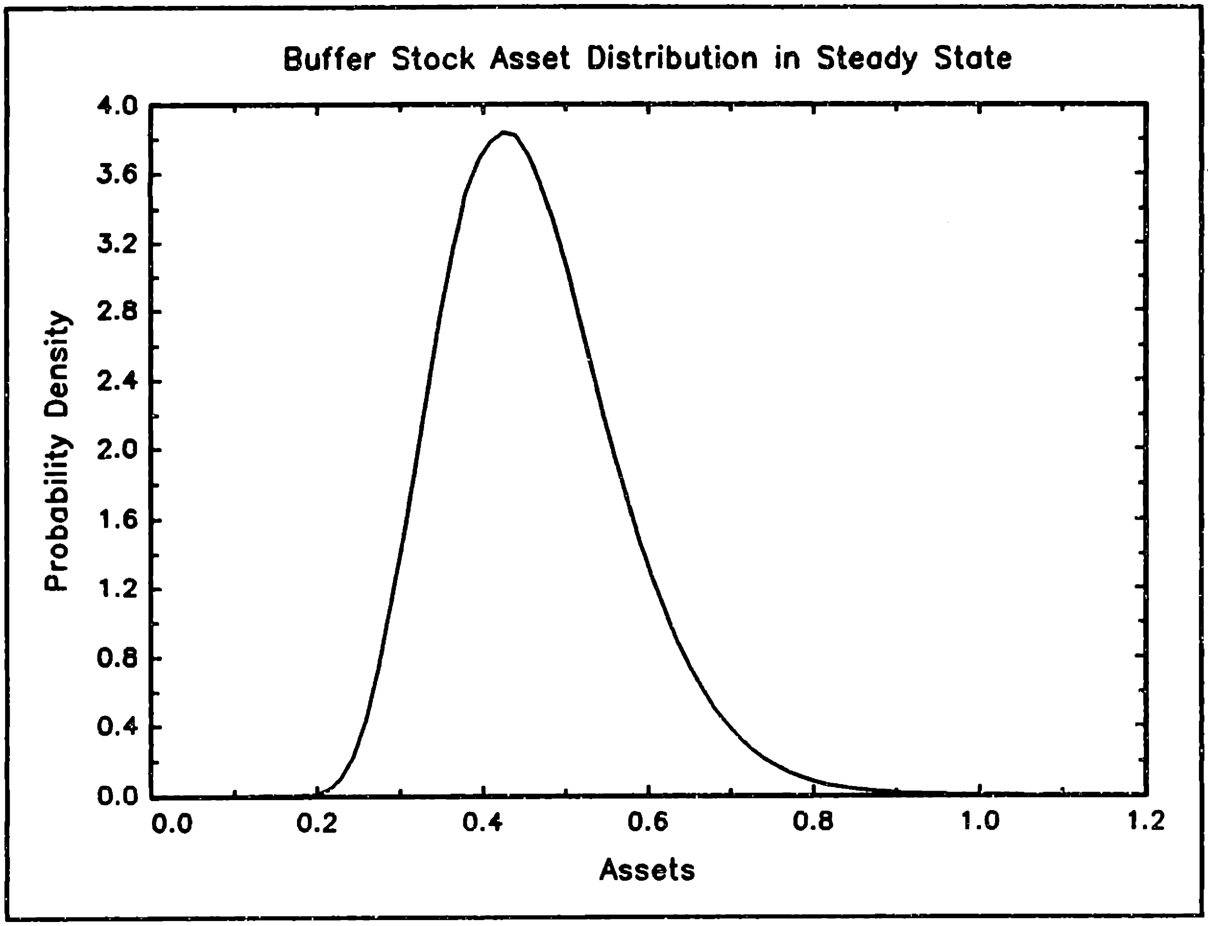

generally predicts much lower levels of wealth than the unconstrained model.[19]

Table 7 uses the data in Table 6 to calculate what the regression coefficients of Section IV would look like if similar regressions were estimated on a population consisting entirely of buffer stock consumers. The numbers in Table 7 are derived from those in Table 6 in the same manner as Table 5 was derived from Table 4. The greatest

---

[19]The exception occurs at low levels of permanent uncertainty, such as $\sigma_n = 0.025$ in our simulations. The possibility of experiencing a zero-income event (job loss) in the next period forces households in a buffer stock model to keep some small amount of assets on hand even if income uncertainty is low. In the standard problem, bad income draws can be offset by going into debt next period and repaying the debt over the remaining lifetime. There is some evidence to support the buffer stock view on this question: very few households (only 5 percent) have negative net worth in the U.S., and among those households the median amount of net worth is about -$1,000, which is well within the debt limit on one credit card. (Calculations from 1983 *Survey of Consumer Finances*).

difference between these results and those for the life cycle model in Table 5 is that the responsiveness of wealth to the variance of the permanent uncertainty term is greatly reduced. In fact, at our baseline parameter value of $\sigma_n = 0.075$ the range of coefficients of [10.7, 18.6] compares with an estimated coefficient of 10.5 for our preferred measure of wealth NHNBW in Table 3b, and is below the estimated coefficient of 37.5 for total net worth in that same table.

The coefficient on the transitory variance is more stable within a range of [2.9, 4.8] for all values of the permanent variance and ranges for the transitory variance. This range is generally close to the estimates found in Table 3, especially to the statistically significant estimates of 4.9 and 4.5 for NHNBW in Table 3b. As in the case of the standard model, the predicted coefficients are somewhat lower than the estimated coefficients on NW in Table 3. Furthermore, the ratio of the permanent coefficient to the transitory coefficient is between 2 and 5 for the buffer stock model, which also matches up with the range of two to four found in the empirical estimates in Table 3.

Based on the predicted magnitude of the coefficient on permanent variance and its ratio to the predicted coefficient on transitory variance, we conclude that, in marked contrast to the unconstrained life cycle model with patient consumers discussed above, a buffer stock model is roughly consistent with our empirical evidence on the nature of precautionary saving.

VII. Conclusion

This paper has found significant evidence of a precautionary saving motive. However, that motive does not take the form implied by an unconstrained life cycle optimization problem with patient consumers. Instead, saving is of the buffer stock variety, in which consumers maintain a comparatively small stock of assets which is effective at insulating consumption against near-term income shocks but is too small to importantly insulate consumption against permanent shocks to income.

The empirical part of the paper decomposes income uncertainty into a variance for transitory shocks and a variance for permanent shocks. For the most plausible regressions summarized in Table 3b, both transitory and permanent uncertainty are statistically significant in predicting the levels of the two more comprehensive measures of wealth. The results are even stronger when the wealthiest households are included in the sample in Table 3a.

We further show that these empirical results are inconsistent with a simple version of the standard intertemporal optimization model with patient consumers. Such a model predicts an extremely large response of wealth holdings to the magnitude of permanent shocks to income: the regression coefficient on the variance of permanent income predicted by the model is on the order of 20-50 times greater than the coefficient estimated in our empirical work. The model also predicts that the permanent coefficient should be on the order of 100 times the value of the transitory coefficient, yet our empirical results typically found ratios in the range of two to four. These extreme results are not due to unusual assumptions about consumers' prudence (the intensity of their precautionary saving motive) or about the magnitude of income shocks (which, if anything, are conservatively estimated). Rather, it arises because even modestly prudent consumers realize that permanent shocks to income will last the rest of their lives; if consumers are patient, so that future utility is discounted only slightly, the present discounted utility from a possible negative shock to permanent income is very large and therefore justifies a large saving response.

The inability of what we term a "standard" life cycle model to explain these results does not imply that the traditional intertemporal optimization framework should be abandoned. Instead, we show that the empirical results are consistent with a "buffer stock" saving configuration of the intertemporal optimization model. In the buffer stock version, consumers are impatient but do not borrow. (This reason can either be direct liquidity constraints, as in Deaton (1991), or fear of extreme negative income shocks, as in our model). Assets are prevented from getting too low by consumers' prudence; they are prevented from getting too large by impatience. We show that the responsiveness of wealth holdings to both transitory and permanent uncertainty, as well as the ratio of these coefficients, in such a buffer stock model is fundamentally consistent with our empirical results.

Of course, the buffer stock framework cannot explain all consumer saving behavior. For instance, it is difficult to explain why consumers as impatient as ours would choose to participate in a pension plan. Explanations of holdings of pension assets, housing equity, and personal business wealth require a more elaborate model. The buffer stock model is most plausible as an explanation of consumers' direct holdings of assets which could be liquidated relatively quickly in the case of unlucky income draws.

If even this more modest claim is true, however, the buffer stock model may be the appropriate model for understanding the behavior of consumption, income, and wealth at business cycle frequencies. If most consumers behave according to a buffer stock model, many of the intuitions formed with more traditional versions of the intertemporal model will be inappropriate. For example, in a buffer stock model predictable changes in income can affect contemporaneous consumption, providing a potential explanation for Campbell and Mankiw's (1989) finding that consumption responds strongly to predictable changes in income without postulating, as they do, the existence of mindless consumers whose consumption exactly equals their income. The buffer stock model is at least qualitatively consistent with the "excess sensitivity" and "excess smoothness" of aggregate consumption found by Campbell and Deaton (1989). It can also explain the results of Carroll and Summers (1991) showing the existence of an aggregate "consumption/income parallel."

Given the broad congruity of the buffer stock model and both macroeconomic and microeconomic findings, the next step is to apply it to more practical economic issues. In subsequent work, Carroll and Samwick (1993), we derive a relationship between the target wealth to income ratio in the buffer stock model and a theoretical measure of income uncertainty. Simulations based on estimates of this relationship suggest that between 25 and 40 percent of aggregate wealth is held in response to labor income uncertainty.

Appendix

1) Sample Restrictions Imposed

As described in Section III, we make several sample restrictions in order to prevent the results from being affected by income changes that may not correspond to uncertainty in the sense required by the model. In addition to those households which experienced a major change in composition, we exclude all households in which the head was younger than 25 or older than 62 at any time during the sample period. The rationale is that younger individuals are not likely to be in stable households and the precautionary saving decisions of older households are likely to be confounded by retirement issues. We also exclude those households whose income in any year was less than 20 percent of its average over the period. This exclusion is necessary in order to calculate our measures of uncertainty; if these households are included our results tend to be almost entirely dominated by a few observations. We also restrict the sample by excluding those households who were included in the PSID's poverty sample. All of these restrictions reduce the amount of variability in the data.

Finally, we use income data only from the years 1981-1987, in the hope that choosing a period centered around 1984 will give us the best possible estimates of the degree of income uncertainty near that year. Results were basically similar when the sample was extended to run from 1976-1987. Although it would seem that more years of data would yield superior estimates, the other sample restrictions we impose--most notably on the age of the household head--sharply reduce the number of households we can study when the sample period is extended. The resulting point estimates are typically close, but the standard errors are larger due to the smaller number of observations in the sample.

2) Estimating Permanent and Transitory Variances in the Sample

A) The Projection Methodology

As described in the text, the variance of an income difference of length d is given by (Equation (5) from Section III):

(A.1) $$Var(r_d) = d \cdot \sigma_\eta^2 + 2 \cdot \sigma_\epsilon^2$$

where σ_η^2 and σ_ϵ^2 are the variances of the permanent and transitory shocks to income, respectively. We estimate $Var(r_d)$ for each *household* by

(A.2) $$v_d = r_d^2 = Var(r_d) + \mu_d$$

where μ_d is a mean-zero disturbance. Using the previous equation to substitute for $Var(r_d)$ yields the following regression equation:

(A.3) $$v_d = d \cdot \sigma_\eta^2 + 2 \cdot \sigma_\epsilon^2 + \mu_d$$

where observations are distinguished by the length of the difference d. Our method simply does OLS household by household of $v = \{v_d(1),\ldots,v_d(n)\}'$ on [d 2], where $d = \{d(1),\ldots,d(n)\}'$, $2 = \{2,\ldots,2\}'$, and d(t) is the t^{th} household difference. As discussed below, we use n = 10 for each household. The coefficients obtained for this regression give household estimates of σ_η^2 and σ_ϵ^2, which we have denoted s_η^2 and s_ϵ^2 in the text.

We have assumed that μ_d and $\mu_{d'}$ are i.i.d. for each household, making OLS the efficient way to conduct the estimation. There does not appear to be any reason to believe that the noise in observing r_d's should vary with d for a given observation, as the data is collected each year and r_d represents a *difference* in annual incomes. For the purposes of this regression, we could also assume without loss of generality that the μ's are i.i.d. across households. The reason is that the regressors for each household are identical, so that a Seemingly Unrelated Regression model would reduce to OLS household by household.

B) Robustness to MA(q) Serial Correlation

We also asserted in the text that the methodology we have adopted can be made robust to serial correlation in the transitory shock of the form MA(q). This can be seen in Equation (4) from Section III:

$$r_d = \left\{\eta_{t+1} + \eta_{t+2} + \dots + \eta_{t+d}\right\} + \epsilon_{t+d} - \epsilon_t \tag{A.4}$$

by noting that $\text{Var}(\epsilon_{t+d} - \epsilon_t) = 2{*}\sigma_\epsilon^2 - 2{*}\text{Cov}(\epsilon_{t+d},\epsilon_t)$. This will yield an unbiased estimate of σ_ϵ^2 whenever $\text{Cov}(\epsilon_{t+d},\epsilon_t) = 0$. Thus, as long as we restrict our choices of d to those greater than q, the procedure will yield the unbiased estimate. Results from MaCurdy (1982) and Abowd and Card (1989) suggest that there is no evidence of serial correlation beyond order 2.

Thus, we choose all possible income differences in which the years are at least three apart. This leaves us with 10 differences for the period 1981-1987:

Length (d)	Years used
6	1981-1987
5	1981-1986, 1982-1987
4	1981-1985, 1982-1986, 1983-1987
3	1981-1984, 1982-1985, 1983-1986, 1984-1987

3) The Impact of Measurement Error in Income on Regression Coefficients

In the text we demonstrated that measurement error in income would not affect the estimate of permanent variance but would increase the estimated variance of transitory income. We argued that the most common form of measurement error would result in our estimating $s_{ik}^2 = \sigma_{ik}^2 + \sigma_u^2$ with the same σ_u^2 for all households. In a regression of the form (see equation 14):

$$\frac{w}{y^P} = \alpha_0 + \alpha_1 s_\eta^2 + \alpha_2 s_\epsilon^2 + \alpha_3 y^P + X\beta + \varepsilon \qquad \text{(A.5)}$$

the σ_u^2 will simply be treated as part of the constant term. In this case the estimate of coefficient α_2 should be unbiased; only α_0 will be biased, and we draw no inferences from the estimated value of α_0.

Consider a second, more sophisticated kind of measurement error. Suppose σ_u^2 is correlated with σ_{ik}^2; indeed, suppose $\sigma_{iu}^2 = \Gamma * \sigma_{ik}^2$, where $\Gamma > 1$ (that is, measurement error is worse for those who have higher transitory variance anyway--a plausible assumption). Then $s_{ik}^2 = (1 + \Gamma)\ \sigma_{ik}^2$. But the effect of this would be to bias our estimate of α_2, the coefficient on transitory income, *down* by increasing the standard deviation of the variable that coefficient is multiplying. Since the primary focus of our attention is on the coefficients on permanent uncertainty, and, that they are already too *small* compared to the coefficients on transitory uncertainty to be explained by the standard life cycle model, this point does not undermine the main results of the paper.

4) Effect of Larger Groups on Negative Correlation in Uncertainty Estimates

In Section III we assert that the negative correlation in estimates of transitory and permanent variances by group would approach zero if the group size approached infinity. In our simplified example using only two differences to estimate s_η^2 and s_ϵ^2, we had the following equations:

(A.6)
$$s_{i\eta}^2 = \sigma_{i\eta}^2 + z_i$$
$$s_{i\epsilon}^2 = \sigma_{i\epsilon}^2 - \frac{(d-1)z_i}{2}$$

where the subscript i refers to household i and the z_i are mean-zero and independent across households and $\mathrm{Var}(z_i) = \sigma_{iz}^2 < \infty$. Divide the population into K groups of N observations each. The covariance across groups is given by:

(A.7)
$$K^{-1}\sum_{k=1}^{K}[N^{-1}\sum_{i=1}^{N}(s_{i\eta}^2 - \sigma_{i\eta}^2)][N^{-1}\sum_{i=1}^{N}(s_{i\epsilon}^2 - \sigma_{i\epsilon}^2)]$$

which is equivalent to

(A.8)
$$K^{-1}\sum_{k=1}^{K}[N^{-1}\sum_{i=1}^{N} z_i][N^{-1}\sum_{i=1}^{N} -\frac{(d-1)}{2}z_i]$$

and can be simplified to

(A.9)
$$K^{-1}\sum_{k=1}^{K}[N^{-1}\sum_{i=1}^{N} -\frac{(d-1)}{2}\frac{z_i^2}{N} + N^{-2}\sum_{i=1}^{N}\sum_{j\neq i} -\frac{(d-1)}{2}z_i z_j]$$

The first summation within the brackets converges to zero as N gets large because $E(z_i^2) = \sigma_{iz}^2 < \infty$. The second summation converges to zero as N gets large because the z_i's are independently distributed. This proves the conjecture. More generally, the result will obtain whenever $\{z_i\}$ satisfies a Law of Large Numbers.

References

Abowd, John M. and David Card (1989). "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, Vol. 57, 411-445.

Avery, Robert B., Gregory E. Elliehausen, and Glenn B. Canner (1984a). "Survey of Consumer Finances, 1983," *Federal Reserve Bulletin*, Vol. 70, 679-692.

Avery, Robert B., Gregory E. Elliehausen, and Glenn B. Canner (1984b). "Survey of Consumer Finances, 1983: A Second Report," *Federal Reserve Bulletin*, Vol. 70, 857-868.

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

Campbell, John Y. and Angus Deaton (1989). "Why Is Consumption So Smooth?" *Review of Economic Studies*, Vol. 56, 357-374.

Campbell, John Y. and N. Gregory Mankiw (1989). "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in Blanchard and Fischer (eds.) *NBER Macroeconomics Annual*. Cambridge: MIT Press, 185-216.

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

Carroll, Christopher D. (1992b). "How Does Future Income Affect Current Consumption?" Board of Governors of the Federal Reserve System, Economic Activity Section Working Paper No. 126.

Carroll, Christopher D. and Lawrence H. Summers (1991). "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven (eds.) *National Saving and Economic Performance*. Chicago: University of Chicago Press.

Carroll, Christopher D. and Andrew A. Samwick (1993). "How Important Is Precautionary Saving?" mimeo, Massachusetts Institute of Technology (also Chapter Two of this thesis).

Curtin, Richard T., F. Thomas Juster and James N. Morgan (1989). "Survey Estimates of Wealth: An Assessment of Quality," in Lipsey and Tice (eds.) *The Measurement of Saving, Investment, and Wealth*. Chicago: University of Chicago Press, 473-552.

Dardanoni, Valentino (1991). "Precautionary savings under income uncertainty: a cross-sectional analysis," *Applied Economics*, Vol. 23, 153-160.

Deaton, Angus (1991). "Saving and Liquidity Constraints," *Econometrica*, Vol. 59, 1121-1142.

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

Guiso, Luigi, Tullio Jappelli and Danielle Terlizzesse (1991). "Earnings Uncertainty and Precautionary Saving," mimeo, Bank of Italy, April 1991.

Hall, Robert E. and Frederic S. Mishkin (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, Vol. 50, 461-481.

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

Kimball, Miles (1990b). "Precautionary Saving and the Marginal Propensity to Consume," NBER Working Paper, No. 3403.

Kimball, Miles (1991). "Precautionary Motives for Holding Assets," NBER Working Paper, No. 3586.

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

Quah, Danny (1990). "Permanent and Transitory Movements in Labor Income: An Explanation for 'Excess Smoothness' in Consumption," *Journal of Political Economy*, Vol. 98, 449-475.

Skinner, Jonathan (1988). "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics*, Vol. 22, 237-255.

Venti, Steven F. and David A. Wise (1990). "But They Don't Want To Reduce Housing Equity," in David A. Wise (ed.) *Issues in the Economics of Aging*, Chicago: University of Chicago Press, 13-29.

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

Table 1

Variance of Permanent and Transitory Income Shocks by Occupation and Education			
Group	Estimated Variance of Permanent Component (Standard Error)	Estimated Variance of Transitory Component (Standard Error)	Number of Obs.
Occupation			
Professional and Technical Workers	0.0183 (0.00573)	0.0303 (0.0114)	305
Managers (not self-employed)	0.0198 (0.00739)	0.0305 (0.0147)	182
Self-employed Managers	0.0147 (0.0134)	0.0962 (0.0266)	56
Clerical and Sales Workers	0.0266 (0.00714)	0.0289 (0.0142)	196
Craftsmen	0.0126 (0.00605)	0.0552 (0.0121)	273
Operatives and Laborers	0.0245 (0.00656)	0.0543 (0.0131)	231
Farmers and Farm Managers	0.0296 (0.0180)	0.129 (0.0358)	31
Service Workers	0.0318 (0.0107)	0.0384 (0.0214)	87
Education			
0-8 grades	0.00588 (0.0126)	0.119 (0.0251)	63
9-12 grades	0.0243 (0.00834)	0.0599 (0.0166)	144
High School Diploma	0.0238 (0.00643)	0.0516 (0.0128)	242
Some College, No degree	0.0232 (0.00433)	0.0340 (0.00863)	534
College Degree	0.0155 (0.00625)	0.0359 (0.0125)	256
Advanced Degree	0.0166 (0.00906)	0.0411 (0.0180)	122

Note: Estimated variances of each group are the averages of individual variances for all members of that group.

Table 2

Descriptive Statistics of Wealth Measures in the PSID and SCF						
	Very Liquid Assets		Non-Housing, Non-Business Wealth		Total Net Worth	
	PSID	SCF	PSID	SCF	PSID	SCF
Level of Wealth						
Mean	19,106	19,193	40,704	31,533	98,081	109,679
Median	4,350	3,749	12,000	11,039	48,635	47,216
Ratio of Wealth to Permanent Income						
Mean	0.4535	0.5504	1.0083	0.9469	2.6125	3.1050
Median	0.1359	0.1385	0.4230	0.3939	1.6204	1.6034
Number of Observations	1204	1364	1080	1364	1045	1364

Note: The data for the SCF are weighted to represent the population in the March 1983 CPS; the PSID data are unweighted. The number of observations differs by wealth measure in the PSID due to missing values for some observations in a particular asset type. Precise definitions of asset categories are given in the text.

Table 3a

Instrumental Variables Regressions of Wealth on Uncertainty
Entire Sample

Wealth Measure	Constant	Permanent Variance	Transitory Variance	Permanent Income	Age	Married	Race	Female	Kids
Very Liquid Assets									
All Occupations	-0.898 (0.211)	6.473** (2.731)	1.462 (1.727)	0.016* (0.084)	0.015** (0.005)	0.058 (0.166)	0.189* (0.102)	-0.112 (0.078)	-0.087** (0.028)
No S.E. or Farmers	-0.605** (0.231)	1.544 (3.996)	-2.063 (3.334)	0.008 (0.013)	0.017** (0.006)	0.136 (0.190)	0.218* (0.126)	-0.083 (0.079)	-0.108** (0.027)
Non-housing, Non-business Wealth									
All Occupations	-2.427** (0.531)	25.205** (10.030)	8.460* (4.874)	0.029 (0.023)	0.027* (0.016)	0.352 (0.419)	0.430 (0.300)	-0.195 (0.267)	-0.023 (0.095)
No S.E. or Farmers	-1.517 (0.369)	11.477* (6.837)	1.037 (6.066)	0.016 (0.023)	0.030** (0.014)	0.402 (0.402)	0.360 (0.289)	-0.140 (0.163)	-0.147** (0.056)
Total Net Worth									
All Occupations	-5.364** (0.955)	49.432** (16.904)	25.248** (7.741)	0.055 (0.036)	0.072** (0.016)	0.944 (0.664)	0.767 (0.478)	-0.431 (0.532)	0.018 (0.151)
No S.E. or Farmers	-2.338** (0.541)	4.194 (9.090)	-2.798 (6.793)	0.020 (0.027)	0.077** (0.016)	0.674 (0.484)	0.658* (0.354)	-0.210 (0.221)	-0.118 (0.081)

Notes:

(1) Heteroscedasticity-robust standard errors are reported in parentheses.

(2) ** denotes significance at the 5% level; * denotes significance at the 10% level.

(3) Instrumental Variables used for the two income variances and permanent income are 8 occupational dummies, 6 educational dummies, and these dummies interacted with age and age-squared.

(4) Demographic variables: Married = 1 if married, 0 otherwise; Race = 1 if white, 0 if nonwhite; Female = 1 if female head of household, 0 otherwise; Kids = number of children less than 18 in household.

Table 3b

Instrumental Variables Regressions of Wealth on Uncertainty
Omitting Top 1 Percent of Wealth Distribution

Wealth Measure	Constant	Permanent Variance	Transitory Variance	Permanent Income	Age	Married	Race	Female	Kids
Very Liquid Assets									
All Occupations	-0.547** (0.122)	2.732 (1.811)	0.560 (0.774)	0.012** (0.004)	0.012** (0.003)	0.012 (0.067)	0.148** (0.052)	-0.037 (0.057)	-0.070** (0.015)
No S.E. or Farmers	-0.482** (0.141)	0.849 (2.477)	-0.667 (1.182)	0.011** (0.004)	0.013** (0.003)	0.000 (0.064)	0.134** (0.050)	-0.039 (0.052)	-0.071** (0.015)
Non-housing, Non-business Wealth									
All Occupations	-1.301** (0.255)	10.549** (3.716)	4.890** (1.692)	0.026 (0.008)	0.019** (0.006)	0.137 (0.168)	0.252 (0.164)	-0.087 (0.137)	-0.108** (0.039)
No S.E. or Farmers	-1.256** (0.282)	9.687* (5.403)	4.542* (2.629)	0.031** (0.009)	0.018** (0.007)	0.074 (0.163)	0.183 (0.175)	-0.120 (0.135)	-0.091** (0.038)
Total Net Worth									
All Occupations	-3.778** (0.675)	37.550** (12.231)	20.107** (5.889)	0.070** (0.021)	0.044** (0.017)	0.522 (0.424)	0.417 (0.368)	-0.386 (0.401)	-0.055 (0.098)
No S.E. or Farmers	-2.182** (0.477)	3.807 (8.152)	1.346 (4.153)	0.043** (0.016)	0.063** (0.010)	0.182 (0.255)	0.438* (0.254)	-0.213 (0.193)	-0.060 (0.064)

Notes:

(1) Heteroscedasticity-robust standard errors are reported in parentheses.

(2) ** denotes significance at the 5% level; * denotes significance at the 10% level.

(3) Instrumental Variables used for the two income variances and permanent income are 8 occupational dummies, 6 educational dummies, and these dummies interacted with age and age-squared.

(4) Demographic variables: Married = 1 if married, 0 otherwise; Race = 1 if white, 0 if nonwhite; Female = 1 if female head of household, 0 otherwise; Kids = number of children less than 18 in household.

Table 4

Mean Asset Holdings at age 40 in Life Cycle Model				
	Transitory Shock Variance (Standard Deviation)			
Permanent Shock Variance (Standard Deviation)	0.0025 (0.05)	0.01 (0.10)	0.0225 (0.15)	0.04 (0.20)
0.000625 (0.025)	0.353	0.392	0.437	0.494
0.0025 (0.05)	1.159	1.184	1.216	1.255
0.005625 (0.075)	2.331	2.353	2.379	2.410
0.01 (0.10)	3.579	3.593	3.617	3.647

Note: Numbers in this table are the theoretically predicted means of the wealth to income ratio from an intertemporal optimization described in the text using the given uncertainty parameters. Coefficient of Relative Risk Aversion is 3. Discount rate and interest rate are 0. Probability of zero-income event is zero.

Table 5

"Coefficients" from Life Cycle Model			
	Permanent Fixed, Transitory Free		
	Coefficient when Transitory Shock Variance goes from-to:		
Permanent Shock Variance Fixed At:	0.0025-0.01	0.01-0.0225	0.0225-0.04
0.000625	5.200	3.600	3.257
0.0025	3.333	2.560	2.229
0.005625	2.933	2.080	1.771
0.01	1.867	1.920	1.714
	Transitory Fixed, Permanent Free		
	Coefficient when Permanent Shock Variance goes from-to:		
Transitory Shock Variance Fixed At:	0.000625-0.0025	0.0025-0.005625	0.005625-0.01
0.0025	429.9	375.0	285.3
0.01	422.4	374.1	283.5
0.0225	415.5	372.2	283.0
0.04	405.9	369.6	282.7

Note: The numbers in this table are the slopes between the means listed in Table 5. Further details are given in the text. Coefficient of Relative Risk Aversion is 3. Discount rate and interest rate are 0. Probability of a zero-income event is zero.

Table 6

Steady-State Mean Asset Holdings in Buffer-Stock Model				
	Transitory Shock Variance (Standard Deviation)			
Permanent Shock Variance (Standard Deviation)	0.0025 (0.05)	0.01 (0.1)	0.0225 (0.15)	0.04 (0.20)
0.000625 (0.025)	0.388	0.412	0.449	0.502
0.0025 (0.05)	0.400	0.426	0.469	0.528
0.005625 (0.075)	0.430	0.459	0.507	0.574
0.01 (0.10)	0.493	0.528	0.588	0.670

Note: Numbers in this table are the theoretically predicted means of the wealth to income ratio from an intertemporal optimization described in the text using the given uncertainty parameters. Coefficient of Relative Risk Aversion is 3. Discount rate is 0.1. Interest rate is 0. Probability of zero-income event is 0.003.

Table 7

"Coefficients" from Buffer Stock Model

Permanent Fixed, Transitory Free

Permanent Shock Variance Fixed At:	Coefficient when Transitory Shock Variance goes from-to: 0.0025-0.01	0.01-0.0225	0.0225-0.04
0.000625	3.320	2.896	3.063
0.0025	3.373	3.448	3.380
0.005625	3.880	3.816	3.820
0.01	4.733	4.768	4.686

Transitory Fixed, Permanent Free

Transitory Shock Variance Fixed At:	Coefficient when Permanent Shock Variance goes from-to: 0.000625-0.0025	0.0025-0.005625	0.005625-0.01
0.0025	6.880	9.408	14.423
0.01	7.093	10.624	15.886
0.0225	10.733	12.096	18.606
0.04	13.706	14.592	22.057

Note: The numbers in this table are the slopes between the means listed in Table 5. Further details are given in the text. Coefficient of Relative Risk Aversion is 3. Discount rate is 0.1. Interest rate is 0. Probability of zero-income event is 0.003.

Chapter Two

How Important Is Precautionary Saving?

I. Introduction

Empirical estimates of the magnitude and even existence of precautionary saving have been remarkably inconsistent: Dynan (1992) found no evidence at all of precautionary saving; Guiso, Jappelli, and Terlizzese (1991) found precautionary saving to be statistically significant but economically unimportant; Kazarosian (1990), Dardanoni (1991), and Carroll (1992a) all found statistically significant evidence of important amounts of precautionary saving. In earlier work, Carroll and Samwick (1992), we have argued that the discrepancies reflect different (and often inappropriate) methodologies, uninformative data, and implausible models. That paper argued that a "buffer-stock" model of saving like those developed in Deaton (1991) or Carroll (1992a), in which households try to accumulate a target wealth stock to self-insure against labor income risk, provides a better explanation than the traditional life cycle model for the empirical relationship between uncertainty and wealth.

Building on the results of our earlier work, this paper attempts to estimate the proportion of wealth that is due to precautionary saving against labor income risk, both in the aggregate and for interesting subgroups of the population.[1] Section II of the paper uses the buffer-stock model to derive equations which capture the theoretical relationship between wealth and two measures of income uncertainty. The first measure

[1] The theory of precautionary saving as developed through Kimball (1990) can apply to any stochastic process--not just that of labor income--which affects a household's budget constraint. For example, in simulations of a life-cycle consumption model with no income uncertainty, Kotlikoff (1986) finds that introducing health expenditure uncertainty can increase long term savings by up to one third, depending on the types of insurance policies available. More recently, Palumbo (1991) has shown that incorporating health and mortality risk can improve the life-cycle model's ability to explain the slow dissaving rates of retired households. For the purposes of this paper, however, we focus only on the fraction of wealth that is due to uncertainty in non-capital income.

is familiar, the coefficient of variation (CV), while the second is based on a theoretical construct called the Equivalent Precautionary Premium (EPP), developed by Kimball (1990). This section of the paper describes how we construct an estimate of the CV and the EPP for each consumer in our PSID sample.

Section III estimates an econometric model of wealth holdings based on the specification derived from the buffer-stock model in Section II. Because the initial model fails a number of specification tests, we develop an augmented model which passes the specification tests. Section IV then applies the augmented model to the question of how much wealth is due to precautionary saving. We find that setting everyone's income uncertainty to the smallest predicted uncertainty in the sample would reduce total net worth by 25 percent and net worth exclusive of housing and business equity by 40 percent. We also conduct some simple distributional analyses and find that precautionary saving accounts for a greater proportion of the wealth of people in the bottom and top deciles of the income distribution than for people in the middle. Section V discusses directions for future research and concludes.

II. A Model of Precautionary Saving

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

The particular version of the buffer-stock model considered here imposes liquidity constraints directly, as in Deaton (1991).[2] Specifically, the consumer is assumed to solve the following problem:

$$
\begin{aligned}
&\max_{\{C_t\}} \quad E_t \sum_{i=t}^{T} \beta^{i-t} u(C_i) \\
(1) \qquad & s.t.\ W_{t+1} = R[W_t + YL_t - C_t] \\
& YL_t = P + e_t \\
& W_s \geq 0 \quad \forall s \\
& u(c) = \frac{c^{1-\rho}}{1-\rho}
\end{aligned}
$$

where ρ is the coefficient of relative risk aversion, W is the stock of physical net wealth, $R = (1+r)$ is the gross interest rate, $\beta = 1/(1+\delta)$ is the discount factor (δ is the discount rate), YL_t is the total noncapital income of the household in period t, P is the permanent noncapital income of the household (that is, the income that would be earned if there were no transitory shocks), and e_t is the transitory shock to income in period t.

[2]This model produces qualitatively similar results to one in which consumers choose not to borrow purely for precautionary reasons, as in Carroll and Samwick (1992). The liquidity constrained formulation is used here because its solution is easier to obtain in the case of an arbitrary income distribution function.

We will assume that the e_t are i.i.d. and will use empirical distribution functions calculated from the *PSID* income data to estimate the distribution of the e_t.[3]

In this model, the optimal consumption rule can be written as a relationship between the ratio of consumption to permanent income, C/P, and the ratio of gross wealth X to permanent income, X/P, where gross wealth X is defined as assets plus current income: $X_t = W_t + YL_t$. If lower case variables denote the upper case counterpart normalized by permanent income, this means that the optimal consumption rule can be written as $c_t(x_t)$--the optimal consumption ratio is a function of the gross wealth ratio and nothing else. (See Deaton (1991) for a proof).

This model and the method of solving it are described in Carroll (1992a) and Deaton (1991), so we will confine the discussion here to an overview of our parametric assumptions and the broad implications of the model. In order to generate "buffer-stock" saving behavior, it must be the case that $\rho > 0$ (prudence) and $\delta > R-1$ (impatience). Plausible values that satisfy these conditions are $\rho = 2$, $\delta = 0.1$, and $R = 1$. The assumptions about the stochastic process for income are described below. The end result of the solution method is an optimal consumption rule which indicates how much should be consumed at each possible level of the wealth-to-income ratio. As described in

---

[3]This is a very restrictive assumption to make about the nature of income uncertainty. In Carroll and Samwick (1992), we allow a more general specification in which there are both permanent and transitory shocks to income, and we find that permanent shocks are quite important. The distinction between permanent and transitory shocks made it possible to distinguish between the buffer-stock and more traditional models of precautionary saving but hampered our efforts to reliably estimate a fraction of aggregate wealth due to labor income risk. Since we are now using only a buffer-stock model, little is gained by the more general error structure. Additionally, the assumption that all shocks to income are transitory allows us to use more flexible, nonparametric kernel estimates of the distribution of income shocks derived directly from the PSID data rather than assuming that the distribution of the permanent and transitory shocks was bivariate lognormal.

Carroll (1992a), this consumption rule implies a "target" wealth-to-income ratio such that if wealth is greater than the target, impatience will outweigh prudence and the consumer will plan to dissave, whereas if wealth is less than the target, prudence will overcome impatience and the consumer will attempt to save.

We must now determine what the empirical implications of this model are and, in particular, what the model implies about the relationship between income uncertainty and wealth holdings. Our method is to solve the model under a range of assumptions about income uncertainty and to derive from the results an empirical formulation that can be tested with the PSID data. The income distributions we use to solve the model are estimated directly using the PSID.

The PSID extract used in this paper contains labor income data for the years 1981 through 1987 for a panel of about 4000 households.[4] We restrict the sample to households whose marital status did not change over the period, who were not part of the poverty subsample of the PSID, who were the head of their household and between 25 and 63 years old over the whole period, and who reported valid occupational, educational, and wealth information. In addition to improving the average quality of data in the sample, these restrictions tend to remove households for whom the variation in reported income may be due to (possibly planned) changes in household composition or transitions into retirement. The resulting estimates of how much income uncertainty is present in the data are thereby reduced, but since the analysis below suggests that income uncertainty is responsible for a sizable portion of wealth holdings, being conservative

[4]Labor income here includes the wage and salary income of all earners in the household, as well as all transfer income from welfare programs or social insurance collected by any member of the household.

about sample selection does not weaken the conclusions. Imposing these sample restrictions also narrows the sample size to a little over 1000.

Our next task was to compute a measure of income uncertainty for these households. After removing the predictable life cycle component of income changes[5] (which presumably do not contribute to income uncertainty), we computed the coefficient of variation of income for each household. The coefficient of variation is an atheoretical measure of uncertainty, however, and without some further theoretical exploration there is no particular reason to believe that it is a good or sufficient summary statistic for the income uncertainty facing households.[6] In order to have a measure of uncertainty that is more solidly grounded in the theory we are testing, we also constructed a measure inspired by a theoretical construct called the "Equivalent Precautionary Premium," (EPP) developed by Kimball (1990).

Kimball's EPP is given by the amount ψ such that $u'(c_1 - \psi) = E\, u'(c_1 + \xi)$, where ξ is the random element in consumption. Since we do not measure total

---

[5]The detrending procedure was as follows. In each year, mean household income was calculated for the whole sample, and "detrended" income for each household was defined as the ratio of actual household income to the mean. This scaling removes common movements in income across all households, including both cyclical and secular trend effects. The next step was to remove predictable life cycle movements from income. An equation was estimated relating detrended income to age, occupation, education, industry, and interactions of these terms. This equation was then used to predict income in each year for each household. Detrended income from the first procedure was then divided by predicted income for each household, generating the yl series used in the calculations above. This second procedure removes changes in income due purely to the predictable effects of aging. Results were basically the same as those reported using a variety of other procedures for removing aggregate and life cycle trends.

[6]It would be a sufficient statistic only if the distribution of income were lognormal. Even in this case, however, the function relating the coefficient of variation to wealth could be highly nonlinear, so the proper functional form for empirical estimation would not be self-evident. The results in Table 1 suggest that CV^2 is an appropriate measure. The estimates of the variances of permanent and transitory shocks to income used in Carroll and Samwick (1992), in which the shocks are assumed to be lognormally distributed, are akin to such a measure.

consumption in the PSID, we cannot measure Kimball's EPP directly.[7] Following Carroll (1992b), however, we can construct a measure of income uncertainty that is analogous to the EPP. Suppose each household had to consume exactly its income each year. Consider a household *i* who in each period *t* consumed amount $YL_{i,t}$ which was distributed i.i.d. with mean P_i, $YL_{i,t} = P_i + e_{i,t}$.[8] Define the mean of $YL_{i,t}$ over the observed period for this consumer as

$$\hat{Y}_i = \frac{1}{7}\sum_{t=1981}^{1987} YL_{i,t}$$

If households really had to consume exactly their income, actual marginal utility for household i in period t would be $u'(YL_{i,t})$. Assuming that, on average, expected marginal utility equals actual marginal utility (rational expectations), and further assuming that the distribution of $e_{i,t}$ does not change over time, we have:

$$E[u'(YL_i)] = \frac{1}{7}\sum_{t=1981}^{1987} u'(YL_{i,t}) + \epsilon_i$$

$$= \bar{u}_i + \epsilon_i$$

where ϵ_i represents the deviation of actual experience over the 1981-87 period from true underlying expectations. Rational expectations implies that on average ϵ_i is zero so that $E[u'(YL_i)] = \bar{u}_i$.

---

[7]Food consumption and scattered other components of consumption are measured, but we consider the quality of these data to be too poor to produce credible results about the extent of uncertainty.

[8]Because the household in this model builds up a buffer stock of assets precisely to insure consumption against large negative shocks to income, the random element in consumption will be less variable than $YL_{i,t}$. The relationship between the two, however, is monotonic; increases in the variability of $e_{i,t}$ will increase the variability of the corresponding ξ.

Suppose again that this consumer has a utility function of the form $u(c) = c^{1-\rho}/1-\rho$, where ρ is the coefficient of relative risk aversion, implying $u'(c) = c^{-\rho}$. Under the assumption that consumption must equal income in each period, it is possible to solve for an estimate of ψ. The analogue to Kimball's theoretical equation in this context is:

$$u'(P_i - \psi_i) = E[u'(YL_i)]$$

Substituting estimated $\hat{Y}$ for P and the constructed value of $\bar{u}$ for E[u'(YLi)] gives:

$$\hat{U}_i = \hat{Y}_i - \left[\frac{1}{7}\sum_{t=1981}^{1987} YL_{i,t}^{-\rho}\right]^{-\frac{1}{\rho}}$$

In this case $\hat{U}_i$ will be a consistent estimator for the ψ_i which corresponds to the true distribution of YL_i. The only further assumption needed in order to compute $\hat{U}_i$ for each household in the PSID is an assumption about ρ. For the purposes of our calculations we assume $\rho = 2$, although we also experimented with other values and found empirical results similar to those reported below. Note that $\hat{U}_i$ is not a scaleless variable: if $YL_{i,t}$ is doubled in every period, $\hat{U}_i$ will double as well. To convert $\hat{U}_i$ into a scaleless variable, we divide by $\hat{Y}_i$ yielding our measure of relative income uncertainty, which we will henceforth designate

$$EPP_i \equiv \frac{\hat{U}_i}{\hat{Y}_i}$$

Kimball shows that the equivalent precautionary premium (constructed, as he does, using consumption uncertainty rather than, as we do, income uncertainty) is, in essence, a direct measure of the intensity of the precautionary saving motive *at the point of zero precautionary saving*. Although there is no closed-form relationship between this concept and the target wealth-to-income ratio which governs consumption decisions, given the explicit model of saving we can solve the model numerically under a variety of possible income distributions and graph the relationship between the EPP and the model's predicted wealth holdings.

To obtain representative observations of the income distribution confronting the households in the data, we divided our sample into subsamples corresponding to the eight occupational and six educational categories to which the head of household could belong. Next, we calculated the empirical cumulative distribution function for the observed ratios of $YL_{i,t}/\hat{Y}_i$ for all the households in each of the categories. The empirical CDFs were approximated by ten-point kernel estimators. We then solved the buffer-stock model of consumption described above and calculated the implied target wealth-to-income ratio for each of these distributions. We also calculated the average value of the equivalent precautionary premium and the coefficient of variation for the households in each of the categories.

The end result was fourteen observations, one for each of the eight occupational groups and the six educational groups. Each observation consisted of values for three variables: 1) the target wealth-to-income ratio predicted by the buffer-stock model when the income distribution function is given by the kernel estimate of the distribution for that group; 2) the empirical EPP for that group; and 3) the empirical CV for that group.

Using these observations, we can ascertain how the income-based EPP is related to wealth, according to the buffer-stock model.

The relationship is given by Figure 1a, which plots the constructed EPP against the model's predicted wealth-to-income ratio. The dashed line in the figure represents the fitted linear regression line of W/P on EPP. It is clear that the linear regression provides an excellent fit for the model relationship between wealth and the EPP. Somewhat surprisingly, the coefficient of variation also appears to be an excellent summary statistic for the impact of uncertainty on wealth holding, at least for the income distribution functions that were estimated for the fourteen groups. Figure 1b plots the wealth-to-income ratio against the CV; this relationship turns out to be close to linear as well, although the fit is not quite as tight as with the EPP. However, over a wider range of uncertainty values the estimated linear relationship depicted in the diagram clearly could not hold, because the intercept is negative and large, implying that if the CV were zero consumers would hold negative wealth, an outcome which clearly could not arise from a model like the one that generated this data in which wealth is constrained to be greater than or equal to zero.

Table 1 presents the results when a variety of regression specifications are estimated for the relationship between the model's predicted target (W/P) ratio and EPP or CV. The first two lines of the table just represent the regressions whose fitted lines are plotted in Figures 1a and 1b. The next regression shows that, between EPP and CV, EPP is more closely related to the theoretical W/P ratio, as should be expected since EPP is based on an economic theory and CV is not.[9] The next two regressions perform a simple test of the linearity of the relationships by including the squares of EPP and CV. Although the EPP^2 term is not significant, the CV^2 term is, and in fact when CV^2 is included the linear term CV is no longer statistically significant. This confirms the theoretical proposition that the true relationship between these two variables is not linear, even though the linear specification fits fairly well over the limited range of values of CV estimated here. This would presumably be a bigger problem for a broader range of income distribution functions.

The final two regressions consider the relationship between the log of (W/P) and the logs of the two measures of uncertainty. This specification corresponds to our empirical work in the next section, in which we find that wealth and the wealth-to-income ratio appear to be much better characterized by a log rather than a linear specification. Both specifications fit the simulated data quite well, but once again the EPP specification fits slightly better.

Our first conclusion from this exercise is that, if income uncertainty has the extremely simple structure we assumed here, the buffer-stock model generates the implication that the relationship between income uncertainty and the wealth-to-income ratio should be well characterized by some equation based on:

[9]In particular, EPP is more sensitive to negative shocks than to positive shocks of equal magnitude, whereas CV treats positive and negative shocks symmetrically. This property is due to EPP's dependence on $u'(\bullet)$ being convex, the same condition necessary to generate precautionary saving.

(2) $$\frac{W}{P} = a_0 + a_1 * EPP$$

Alternatively, the log specification fits approximately as well and generates an equation similar to:

(3) $$\log\left(\frac{W}{P}\right) = a_0 + a_1 * \log(EPP)$$

Our second conclusion is that, at least over a limited range of values of uncertainty, the coefficient of variation performs nearly as well as the equivalent precautionary premium as a predictor of the amount of precautionary wealth.

III. Empirical Estimation of the Model

Our first empirical efforts were directed at estimating equations based on equation (2):

(4) $$\frac{W}{P} = a_0 + a_1 * EPP + a_2 * P + a_3' Z + v$$

We included the level of permanent income on the right hand side because it is plausible to suspect that the wealth-to-income ratio might increase as the level of permanent income increases. We included the demographic variables Z because it is likely that these demographic variables affect the marginal utility of consumption, the coefficient of relative prudence, or other legitimate theoretical determinants of the wealth-to-income ratio.

When we estimated equations like (4) using IV techniques described below, we generally found that the coefficient a_1 was overwhelmingly significant. However, closer examination of the results turned up several problems, mostly stemming from the extreme nonnormality of the distribution of the W/P ratio. The top few observations were many, many standard deviations from the predicted value and therefore exercised inordinate influence on the coefficients. (This was particularly disturbing because the richest households are probably those for whom the buffer-stock model of saving is least plausible.) The results were somewhat fragile to the inclusion of outlying observations, and tests of the normality of the residuals rejected overwhelmingly.

Alternatives to ordinary least squares do provide ways of estimating equation (4) despite these problems with the specification. To temper the effect of outliers on the parameter estimates, we used the robust regression framework proposed in Huber (1964) as implemented by Li (1985). This method systematically weights observations with larger residuals less in the minimization of squared residuals. The robust method produces smaller parameter estimates which are insignificantly different from zero. The effect of nonnormal residuals can be diminished by estimating (4) as a quantile regression model. The median regression of Koenker and Bassett (1978) minimizes the sum of absolute (rather than squared) deviations and is therefore less sensitive to the magnitude of extreme residuals. Median regression estimates of equation (4) yield much smaller, but still significant, coefficients. Estimates at quantiles other than the median reveal a pattern of increasing coefficients (and significance) on EPP and P at higher quantiles. Powell (1986) shows that large differences in estimates at different quantiles is indicative of heteroscedasticity in the error terms.

Our solution to this problem is based on the empirical regularity that over most of its range, wealth is well characterized by a lognormal distribution. Figure 2 confirms this: it plots the cumulative distribution function of the log of our preferred measure of wealth, non-housing non-business wealth, along with an exact lognormal distribution whose parameters are estimated from the data.[10] Figure 2 suggests that a better specification for the regression would use the log of wealth as the dependent variable. Since the results in Table 1 found that the logarithmic specification fit the theoretical model about as well as the linear model, we adopted a new specification derived from equation (3):

(5) $$\log\left(\frac{W}{P}\right) = a_0 + a_1 * \log(EPP) + v$$

Adding log(P) to both sides gives:

(6) $$\log(W) = a_0 + a_1 * \log(EPP) + \log(P) + v$$

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

(7) $$\log(W) = a_0 + a_1 * \log(EPP) + a_2 * \log(P) + a_3' Z + v$$

where the Z variables are demographic controls for age, race, sex, marital status, and the number of children. Equation (7) does not constrain the coefficient on log(P) to be unity, thereby allowing the wealth-to-income ratio to rise with the level of permanent income.

---

[10] In order to take the log we had to drop from our sample all observations where wealth was zero or less. These observations comprised slightly less than 10 percent of the sample. The median value of wealth for these people was about -$2000, or roughly the capacity of one credit card.

The theoretical model in Section II makes two simplifying assumptions that affect the appropriate measure of wealth to be the dependent variable in equation (7). First, the model contains only one asset, but in reality there are a multitude of assets in which to invest. Apart from considerations of the covariance of the asset's return with the shocks to income, the relevant dimension along which these assets will differ is their liquidity. In principle, any asset can be used for precautionary saving, but more illiquid assets are less useful as a safeguard against adverse shocks to income because of the extra time or money that they require to be converted into cash. It would therefore not be surprising to find a weaker link between the amount of uncertainty and amount of less liquid assets. Second, the model abstracts from other economic motives for saving such as provision for retirement.

To counter this shortcoming with the model, we examine three measures of wealth (the components of which are detailed in the Appendix): very liquid assets (VLA), non-housing, non-business wealth (NHNBW), and total net worth (NW). The liquidity problem argues for a very narrow measure of wealth like VLA, as any asset included in VLA can be liquidated on short notice with nearly zero transactions costs. The narrowness of the definition of VLA must be counterbalanced by the possibility that, for reasons outside the scope of the model, some households choose to hold precautionary saving in less liquid forms than VLA. Additionally, the problem of savings not related to uncertainty suggests a very general measure such as NW should be used so that the impact of other saving is spread out over many types of assets. Our preferred measure of wealth is the intermediate category of NHNBW. It includes the net equity in all types of wealth excluding housing and businesses, which tend to be less liquid than other types of wealth.[11] Thus, NHNBW appears to be the most relevant measure to use to study the effect of labor income risk on wealth holding.[12]

The final issue to be addressed before presenting the estimation results is the nature and justification of the instrumenting procedure used. Both uncertainty and permanent income are unquestionably measured with error in our data, so they must be instrumented in order to get consistent estimates of the coefficients on these terms. The instrument set contains dummies for the occupation, education, and industry of the head of household in 1981, along with the demographic variables already contained in Z. We also interacted the occupation and education variables with the age and age^2 terms in order to allow for different lifetime profiles of income and uncertainty. (The set of instruments is described fully in the Appendix). The R^2 from the first-stage regression of log(P) on the instruments was 0.45, while the R^2 from the first-stage regression of both of the uncertainty variables on the instruments was 0.15.

Results for our basic specification for all three measures of wealth, and for both measures of uncertainty, are presented in Table 2. The coefficients on the EPP and CV terms are overwhelmingly significant for all three measures of wealth and are largest for

[11]This statement is more true of 1984, the year for which wealth is measured in the PSID, than of today, given the proliferation of home equity loans since the Tax Reform Act of 1986.

[12]Carroll and Samwick (1992) show that the sample means and medians of these wealth measures match up quite closely with those of the *Survey of Consumer Finances 1983*, a survey with much more detail than the PSID. This finding confirms the basic results of Curtin, Juster, and Morgan (1989) who carefully compared the SCF and PSID and concluded that the latter provided nearly as good a measure of wealth as the former for most of the population.

NHNBW.[13] Note that the coefficient on the level of permanent income is always well in excess of one, the value that would be expected if the wealth-to-income ratio were homothetic. Although this is, strictly speaking, a violation of the homothetic model, it is not particularly surprising to find that people with high permanent incomes hold disproportionately higher amounts of wealth.[14]

Another notable result is that the coefficients on the EPP and CV terms are of approximately equal significance. These variables are also highly correlated with each other, and regressions (not reported) in which both measures are included on the RHS find that neither measure is individually significant.

The uniformity of statistical significance of the coefficients on the uncertainty terms seems to strongly support the precautionary saving model. However, the next table reveals some problems with the basic specification. In using occupation, education, and industry variables as instruments, it is implicitly assumed that the only reason why occupation, education, and industry are correlated with wealth is because they have predictive power for the independent variables in our model. This is an overidentifying assumption which can be tested using standard tests of overidentification. The first column of Table 3 presents the p-value from one such test. The model fails the

[13]Heteroscedasticity tests rejected the null hypothesis of homoscedasticity, so all standard errors are heteroscedasticity-robust.

[14]This robust finding is not an implication of traditional intertemporal optimization frameworks. Hubbard, Skinner, and Zeldes (1992) do generate such a relationship endogenously due to the presence of means-tested transfer programs.

specification test because the 1.9 percent p-value is well below the usual 5 percent significance level.[15]

Further investigation revealed that the reason the model fails is that occupational variables have explanatory value for wealth independent of their ability to predict the current level of income or uncertainty. This is not surprising, because a more traditional life cycle model of saving would predict that, if lifetime profiles of income differ by occupational group, lifetime wealth profiles should also depend on occupational group. In such a case, occupational variables should be included as RHS variables in the main regression, and not just as instruments in the first-stage regression.

To correct this problem we estimate a second model, which we will refer to as the "augmented" model, in which the occupational variables are added to the set of Z variables when the equation is estimated. This leaves primarily education and industry as the excluded variables from the model specification. It remains possible, of course, that education and/or industry are also correlated with wealth in ways independent of their ability to predict income or uncertainty. The second row of Table 4, however, shows that an overidentification test of the augmented model has a p-value of 19 percent, which does not reject the specification as correct at conventional significance levels.

The second specification test presented in this table is a test for the normality of the residuals of the model. Although the theoretical model does not necessarily imply that the residuals from the estimation equation should be normal, a strong rejection of

---

[15]The test statistic in this case is the sample size times the R^2 from the regression of the residuals from OLS on equation (7) on the full instrument set. The test statistic is asymptotically distributed χ^2 with degrees of freedom equal to the number of instruments not included in Z. The derivation of this test is given in Godfrey (1988).

normality would suggest that the model specification was a poor one. Indeed, overwhelming rejections of normality of residuals were one of the main reasons we rejected the linear specification of the model which we tried first, as discussed above. The p-value for the normality test is between 5 and 10 percent for both the basic and the augmented models, indicating that while the residuals may be slightly nonnormal, the implications are not as severe (as they were in the estimation of equation (4) above).[16]

Given that the augmented model passes the specification tests, the crucial question is whether uncertainty is related to wealth in the augmented model. Table 4 presents the basic results from estimating the augmented model.[17] The coefficients on uncertainty remain positive for all three measures of wealth and both measures of uncertainty. However, the coefficients on the uncertainty terms are considerably smaller for all three measures of wealth, and are not significantly different from zero for very liquid assets or for total net worth. Nevertheless, the coefficient on NHNBW, the most reasonable measure of wealth to be the repository of precautionary saving, does remain statistically significant at the 5 percent level.[18]

---

[16]In general, the coefficients and standard errors were far less sensitive to sample selection and specification issues in this log-linear formulation than in the linear specification originally tried. In the case of NHNBW, the parameter estimates (and their statistical significance) in the basic specification obtained using both robust and median regression were quite close to the ordinary and two-stage least squares estimates. Additionally, the pattern of increasing coefficients at higher quantiles is much less severe than in the linear specification.

[17]Coefficients for the occupation variables are estimated but are not of direct interest and therefore not reported. Note, however, that the inclusion of these variables changes the interpretation of the coefficients for the intercept, the age, and the age^2 variables in that they now represent the values for the omitted occupation group rather than the whole sample.

[18] These results are broadly consistent with those in Carroll and Samwick (1992), in which very liquid assets were not closely related to income uncertainty but non-housing, non-business wealth was. In that paper, however, there was a significant effect of uncertainty on the holdings of total net worth.

The regressions reported in Table 4 and the specification tests in Table 3 suggest that the augmented model is an adequate representation of the relationship between income uncertainty and wealth holding. The marginal effect of income uncertainty on wealth holding is always estimated to be positive, and the coefficient is statistically significant under the most reasonable specification.

IV. How Would Wealth Change If There Were Less Uncertainty?

Having estimated a suitable empirical model of the effect of uncertainty on individual households, the next step is to use that model to determine the importance of income uncertainty in the aggregate wealth distribution. The basic experiment will be to simulate the distribution of wealth under the assumption that all households faced the same, small, amount of uncertainty, instead of the amount of uncertainty they actually faced.

In the second-stage regressions presented in Table 4, each household has a corresponding wealth equation:

$$\text{(8)} \qquad \log(W_i) = a_0 + a_1 * \log(EPP_i) + a_2 * \log(P_i) + a_3' Z_i + v_i$$

This equation fits exactly because the v_i are *defined* as the difference between the predicted value of log wealth and the actual value.[19] The procedure is merely to substitute some constant value EPP^* for EPP_i while leaving the values of all other RHS variables the same. This procedure indicates how much wealth this household would own if its uncertainty were given by EPP^* but it was otherwise identical (including the

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

value of the error term v_i). That is, we construct a new measure of wealth, W_i^*, such that:

(9) $$\log(W_i^*) = \log(W_i) + a_1\left[\log(EPP^*) - \log(EPP_i)\right]$$

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

Because there is no single level of EPP^* that is the obvious one to use, a range of values is tried. In particular, EPP^* is first set to the minimum predicted value of EPP_i in the sample and then to the values at the 10th and 25th percentiles. As a reference point, the calculations are also done setting EPP^* equal to the median predicted value of EPP_i.

The "predicted" value of EPP refers to the predicted value of log(EPP) from the first-stage regression of $\log(EPP_i)$ on the instrument set. It would be inappropriate to set EPP^* to the minimum *actual* value in the sample because EPP_i is measured with error. The minimum actual value in the sample may therefore correspond not to someone who actually had such a small amount of uncertainty but to someone for whom the measurement error for uncertainty was especially large and negative. By using the predicted value, the conceptual experiment is to set everyone's uncertainty to the value that would be predicted for someone with a particularly low-uncertainty combination of characteristics such as a stable industry, high education, a secure occupation, etc. This will presumably be a considerable overestimate of the minimum amount of uncertainty faced by anyone in the sample, and for this reason the estimate of the effect setting everyone's uncertainty to the minimum predicted value is almost surely an underestimate of the true effect of setting uncertainty to the minimum *possible* value.

Figures 3a-d plot the simulated values of the log(NHNBW) measure of wealth against the true distribution of log(NHNBW) for each of our four values of EPP^*. To be specific, two points are plotted for each household in the sample: (w_i, p_i) and (w_i^*, p_i) where pi indicates the percentile ranking for household i in the true wealth distribution, w_i indicates the log of actual wealth W_i, and w_i^* indicates the simulated value of the log of wealth for that household when EPP_i is set to the chosen value of EPP^* (i.e. $\log(W_i^*)$ in equation (9)).[20]

In the first figure, where EPP^* is set equal to the minimum predicted value of EPP_i in the sample, all simulated wealth points lie to the left of the actual wealth distribution. Everyone's uncertainty has been reduced (except the household which had the minimum predicted value of EPP, which has an unchanged value); consequently, everyone's wealth holdings are reduced. In the second figure, 10 percent of the points lie to the right of the distribution; in the third figure, 25 percent are to the right; and in the final figure, half of the points are to the right and half are to the left.

The figures give a graphical depiction of how much the log(wealth) of individual households is changed by the reduction of uncertainty. The question at hand, however, is how much aggregate wealth is affected. The procedure for calculating the effect on

[20]The simulations conducted here are partial equilibrium in the sense that interest rates and wage rates are not permitted to adjust in response to the reduction in income uncertainty. Aiyagari (1992) investigates a general equilibrium model of precautionary saving and concludes that for sufficiently high variability and persistence in earnings, aggregate saving can be 7-14 percentage points higher than in the model with no income uncertainty.

aggregate wealth is straightforward: if each household's simulated log(wealth) is given by w_i^*, then its simulated level of wealth is given by $\exp(w_i^*)$. Simulated aggregate wealth is then given by the sum of the simulated wealth of the individual households.

The results are shown in Table 5 for each of the three categories of wealth. The first row shows that reducing every household's uncertainty to the minimum predicted value reduces aggregate liquid assets by 25 percent, aggregate non-housing, non-business wealth by nearly 40 percent, and total net worth by about 25 percent. Recall that these results underestimate the total amount of wealth due to precautionary saving, because the actual minimum possible value uncertainty is presumably less than the minimum predicted value of EPP. This experiment is in effect examining how much wealth would be reduced if everyone's uncertainty were set equal to the *average* uncertainty faced by a person in the lowest-uncertainty group. The rest of the table presents the results for the other experiments.

Our estimates of 25 to 40 percent of wealth being held in response to income uncertainty are considerably less than other estimates found in the literature in different contexts. Dardanoni (1991) finds that "60% of savings in the sample arise as a precaution against future income risk," and Skinner (1988) finds that "[u]sing empirical parameter values, precautionary savings are estimated to be 56 percent of total life cycle savings." Also, Caballero (1991) calibrates a theoretical model which suggests that for plausible parameters, income uncertainty may be responsible for "in excess of 60 percent of aggregate wealth accumulation." The latter two papers base their conclusions on simulations from purely theoretical models; that their estimates are so high reflects that despite claims to the contrary, their models are poorly parameterized. As discussed in Carroll and Samwick (1992), Dardanoni's measure of income uncertainty and emphasis on current consumption rather than accumulated wealth are inappropriate to the study of precautionary saving.

The final simulations pertain to the effects of income uncertainty on wealth for households in different permanent income and wealth categories. Intuition suggests (and regressions confirm) that income uncertainty is less for people in higher income groups. If this were so, we would expect to find that less of the wealth of such people is due to uncertainty. The first column of Table 6 calculates the amount of wealth that is estimated to be due to precautionary saving for people in each decile of the permanent income distribution. As expected, precautionary saving accounts for a larger portion of wealth for people in the lower income deciles, accounting for more than half of the wealth of people in the bottom decile. Proceeding up the income distribution, the proportion of wealth due to precautionary motives steadily declines--until the top decile, where it suddenly jumps again. Further investigation revealed that this jump was due to the presence in this decile of a disproportionate number of the self-employed, who had high income uncertainty and therefore a high fraction of their wealth attributed to precautionary motives.

A similar question is addressed by the second column of the table: How much wealth of the people in each *wealth* decile is due to precautionary saving? This column shows that the bottom two and top two deciles hold disproportionately more of their wealth for precautionary reasons, while the fraction of wealth due to precautionary motives is virtually flat from the 30th to the 80th percentiles.[21] Note, however, the peculiar nature of the experiment being performed here. If there is in fact a positive effect of uncertainty on wealth holdings, then households in the top wealth decile will be disproportionately those with a lot of income uncertainty. Households with much income uncertainty are of course the households whose wealth is reduced the most by reducing everybody's uncertainty to the same, small number. From this perspective it is no surprise to find that the people in the top deciles appear to hold more of their wealth for precautionary reasons.

However, the same principle should be working in reverse at the bottom of the wealth distribution: households with very little wealth should be disproportionately people with very little income uncertainty. Thus, reducing uncertainty to the minimum amount should have relatively little effect on their wealth than on that of others. Instead, the actual result is that the bottom two deciles are estimated to hold more of their wealth for precautionary reasons. This is a strong indicator that they in fact have much greater than average income uncertainty.

V. Conclusion

This paper improves upon the existing literature on precautionary saving by developing an econometric specification that is more firmly rooted in economic theory. The buffer-stock model used is the same as that in Deaton (1991) and Carroll (1992a), which has been found to be consistent with both macroeconomic (see Carroll (1992a))

---

[21]Households with zero or negative wealth were excluded from the sample because the logarithmic specification ruled them out. Since about 10 percent of the sample had zero or negative wealth, the people in the 10th decile of the wealth distribution of this sample are around the 20th percentile of the full sample.

and microeconomic (see Carroll and Samwick (1992)) data. When confronted with actual income distributions found in the PSID, the model is consistent with a log-linear relationship between wealth and a theoretical measure of income uncertainty based on Kimball's (1990) equivalent precautionary premium. Our augmented econometric specification based on these simulations is not rejected by a standard overidentification test.

Our principal finding is that wealth holdings are positively and significantly related to income uncertainty for various measures of both wealth and uncertainty. Simulations reveal that a conservative estimate of the amount of non-housing, non-business wealth held for precautionary reasons is 40 percent, with approximately 25 percent of both the more general (total net worth) and less general (very liquid assets) wealth measures attributable to income uncertainty. That the largest effect is measured for the intermediate measure of wealth suggests that there may be important consequences of income uncertainty on portfolio composition as well. The importance of precautionary saving in determining the aggregate wealth stock indicates that policies designed to encourage saving or redistribute wealth must also be evaluated with regard to their effects on income uncertainty.

We also find that precautionary saving is more important in the tails of the permanent income distribution than in the middle. That the wealth holdings of those with very high permanent incomes should be more sensitive than those with average permanent incomes to the degree of income uncertainty is not a prediction of the buffer-stock model. One explanation for this finding is that the behavior of the very wealthiest households is in accordance with a different model in which both the level of wealth and amount of labor income uncertainty confronted rises with permanent income. We searched for a specification which would allow the relationship between income uncertainty and wealth to be different for the households at the top end of the income distribution from the relationship for all other households, but we were unable to find evidence of any such structural break in the equations. The failure to find evidence of such a break in the PSID data suggests that if it does exist, it is near, or higher than, the top of the wealth distribution found in the PSID.

This finding, along with the non-homotheticity of wealth with respect to income, suggests the need for more investigation of the very poor and very rich in order to better understand the behavior of small but important groups within the population. This has prompted us to begin investigating the wealth data in the Federal Reserve Board's *Surveys of Consumer Finances*, which oversample the wealthiest households. This dataset provides a better opportunity to find some evidence of the structural break in the relationship between income uncertainty and wealth holdings which is not evident in the PSID data.

References

Aiyagari, S. Rao (1992). "Uninsured Idiosyncratic Risk and Aggregate Saving," Federal Reserve Bank of Minneapolis, Research Department Working Paper No. 502.

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

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

Carroll, Christopher D. (1992b). "How Does Future Income Affect Current Consumption?" Board of Governors of the Federal Reserve System, Economic Activity Section Working Paper No. 126.

Carroll, Christopher D. and Andrew A. Samwick (1992). "The Nature and Magnitude of Precautionary Wealth," Board of Governors of the Federal Reserve System, Economic Activity Section Working Paper No. 124 (slightly revised as Chapter One of this thesis).

Curtin, Richard T., F. Thomas Juster, and James N. Morgan (1989). "Survey Estimates of Wealth: An Assessment of Quality," in Lipsey and Tice (eds.) *The Measurement of Saving, Investment, and Wealth*. Chicago: University of Chicago Press, 473-552.

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

Deaton, Angus (1991). "Saving and Liquidity Constraints," *Econometrica*, Vol. 59, 1121-1142.

Dynan, Karen E. (1992). "How Prudent are Consumer?" mimeo, Harvard University.

Godfrey, L. G. (1988). *Misspecification Tests in Econometrics: The Lagrange Multiplier Principle and Other Approaches*. Cambridge: Cambridge University Press, 191-198.

Guiso, Luigi, Tullio Jappelli and Danielle Terlizzesse (1991). "Earnings Uncertainty and Precautionary Saving," mimeo, Bank of Italy, April 1991.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1992). "Precautionary Saving and Social Insurance," mimeo, University of Pennsylvania.

Huber, Peter J. (1964). "Robust Estimation of a Location Parameter," *Annals of Mathematical Statistics*, Vol. 35, 73-101.

Kazarosian, Mark (1990). "Precautionary Saving--A Panel Study," mimeo, Boston College.

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

Koenker, Roger and G. Bassett (1978). "Regression Quantiles," *Econometrica*, Vol. 46, 33-50.

Kotlikoff, Laurence J. (1986). "Health Expenditures and Precautionary Savings," National Bureau of Economic Research Working Paper No. 2008.

Li, G. (1985). "Robust Regression," in David C. Hoaglin, Frederick Mosteller, and John W. Tukey (eds.) *Exploring Data Tables, Trends, and Shapes*. New York: John Wiley and Sons.

Palumbo, Michael G. (1991). "Health Uncertainty and Optimal Consumption Near the End of the Life Cycle," Ph.D. Dissertation, University of Virginia.

Powell, James L. (1986). "Censored Regression Quantiles," *Journal of Econometrics*, Vol. 32, 143-155.

Appendix

1) The set of instrumental variables

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

A) Household composition

Variable	Description	Comment
Age	Age in years	Sample mean is 38 years.
Age2	Square of age	
Married	Marital status indicator	86.2% of household heads are married.
White	Race indicator	92.5% of household heads are white.
Female	Gender indicator	9.2% of household heads are female.
Kids	# of children under age 18	Sample mean is 1.3 children.

B) Indicator variables for occupation

Occupation group	Percent of sample
Professional and Technical Workers	24.39
Managers (not self-employed)	13.11
Managers (self-employed)	4.00
Clerical and Sales Workers	13.73
Craftsmen	21.00
Operatives and Laborers	16.09
Farmers and Farm Laborers	1.95
Service Workers	5.74

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

C) Indicator variables for education

Education group	Percent of sample
0-8 Grades	4.51
9-12 Grades	9.03
High School Diploma	17.74
Some College, No Degree	37.95
College Degree	20.92
Some Postgraduate	9.85

The education variables are also interacted with Age and Age^2 to allow for education-specific age-income and age-uncertainty profiles.

D) Indicator variables for industry

Industry Group	Percent of sample
Agriculture, Forestry, Fishing	3.06
Mining	1.27
Construction	6.34
Manufacturing	29.99
Transportation, Communications, and Utilities	11.40
Wholesale and Retail Trade	13.73
Finance, Insurance, and Real Estate	4.86
Business and Repair Services	3.80
Personal Services	1.48
Entertainment and Recreation Services	0.63
Professional and Related Services	15.42
Public Administration	8.08

2) The wealth regressions

A) Basic specification

The basic specification (see Table 2) has the household composition variables (A) as independent variables, along with permanent income and income uncertainty. The instrument set consists of all the variables given in (A) - (D) above.

B) Augmented specification

The augmented specification (see Table 4) has the household composition variables (A) and occupation variables (B) as independent variables, along with permanent income and income uncertainty. The instrument set consists of all the variables given in (A) - (D) above.

3) Wealth measures

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

A) Very Liquid Assets (VLA)

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

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

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

C) Total Net Worth (NW)

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

Wealth in the PSID				
Wealth Measure	Number of Observations	Mean	Median	Standard Deviation
VLA				
Full Sample	1204	19,106	4,350	58,708
VLA > 0	1073	21,439	5,500	61,787
NHNBW				
Full Sample	1080	40,704	12,000	120,800
NHNBW > 0	976	45,974	15,000	125,505
NW				
Full Sample	1045	98,081	48,635	229,347
NW > 0	997	103,233	53,450	233,461
All dollar amounts are *not* weighted.				

Table 1

Regressions of Simulated Target Wealth Ratios on Uncertainty Measures

Dependent Variable	Constant	EPP	CV	EPP^2	CV^2	Log(EPP)	Log(CV)	R^2
W/P	-0.014 (0.004)	2.473 (0.042)						0.9965
W/P	-0.159 (0.015)		1.548 (0.065)					0.9802
W/P	-0.021 (0.020)	2.368 (0.328)	0.068 (0.211)					0.9965
W/P	-0.024 (0.013)	2.672 (0.243)		-0.867 (1.041)				0.9967
W/P	-0.025 (0.070)		0.507 (0.545)		2.032 (1.027)			0.9854
Log(W/P)	0.440 (0.029)					1.073 (0.027)		0.9925
Log(W/P)	0.374 (0.047)						1.685 (0.073)	0.9779

Notes:

1) W/P - Target ratio of wealth to permanent income; EPP - Equivalent Precautionary Premium; CV - Coefficient of Variation
2) Standard Errors in parentheses.
3) Uncertainty values are the group means of the uncertainty variable for the eight occupational and six educational groups in the PSID.
4) Wealth values are the target wealth to income ratios generated by a liquidity constrained household facing the earnings distribution for each occupational or educational group.

Table 2

Instrumental Variables Regressions of Wealth on Income Uncertainty -- Basic Specification

Uncertainty Wealth	Constant	Uncertainty	Income	Age	Age^2*10^{-3}	Married	White	Female	Kids	Obs.
EPP										
VLA	-23.680	0.600	3.161	0.077	-0.069	-0.544	0.442	-0.050	-0.281	1073
	(2.965)	(0.183)	(0.330)	(0.073)	(0.092)	(0.236)	(0.253)	(0.190)	(0.053)	
NHNBW	-14.033	0.620	2.235	0.116	-1.228	-0.035	0.363	0.142	-0.164	976
	(2.542)	(0.167)	(0.279)	(0.061)	(0.774)	(0.195)	(0.183)	(0.179)	(0.045)	
NW	-13.556	0.504	2.056	0.211	-2.285	0.329	0.400	0.097	-0.095	997
	(2.265)	(0.141)	(0.238)	(0.061)	(0.768)	(0.195)	(0.185)	(0.152)	(0.048)	
CV										
VLA	-23.232	1.186	3.138	0.074	-0.066	-0.554	0.429	-0.049	-0.277	1073
	(2.854)	(0.354)	(0.321)	(0.073)	(0.091)	(0.235)	(0.256)	(0.188)	(0.052)	
NHNBW	-13.879	1.235	2.239	0.115	-1.233	-0.061	0.363	0.129	-0.164	976
	(2.482)	(0.322)	(0.276)	(0.060)	(0.765)	(0.196)	(0.185)	(0.179)	(0.044)	
NW	-13.439	1.213	2.064	0.210	-2.291	0.298	0.406	0.082	-0.094	997
	(2.191)	(0.266)	(0.234)	(0.060)	(0.763)	(0.195)	(0.185)	(0.153)	(0.048)	

Notes:
1) VLA - Very Liquid Assets; NHNBW - Non-housing, Non-business Wealth; NW - Total Net Worth
EPP - Equivalent Precautionary Premium; CV - Coefficient of Variation.
2) The measure of income uncertainty is EPP in the first three equations and CV in the second three equations.
3) All wealth, income, and uncertainty values are in logs.
4) All observations with negative reported wealth are excluded.
5) Heteroscedasticity-robust standard errors in parentheses.
6) The Set of instrumental variables is described in the text.

Table 3

Specification Tests		
Model	Overidentification Test (p-value)	Normality of Residual Test (p-value)
Basic	0.01	0.09
Augmented	0.19	0.07

Notes:

1) The overidentification test is a χ^2 test for the exogeneity of the instruments with respect to the residuals from the second stage regression.
2) The normality test reports the p-value for the null hypothesis that the residuals exhibit zero skewness.

Table 4

Instrumental Variables Regressions of Wealth on Income Uncertainty -- Augmented Specification

Uncertainty Wealth	Constant	Uncertainty	Income	Age	Age^2*10^{-3}	Married	White	Female	Kids	Obs.
EPP										
VLA	-21.085	0.286	2.812	0.058	-0.027	-0.377	0.498	0.051	-0.275	1073
	(4.158)	(0.231)	(0.414)	(0.123)	(1.559)	(0.280)	(0.223)	(0.182)	(0.050)	
NHNBW	-14.692	0.386	1.992	0.224	-2.374	0.019	0.333	0.089	-0.162	976
	(3.634)	(0.192)	(0.348)	(0.105)	(1.334)	(0.222)	(0.166)	(0.159)	(0.043)	
NW	-15.162	0.224	1.839	0.330	-3.578	0.326	0.364	0.023	-0.082	997
	(3.189)	(0.178)	(0.291)	(0.099)	(1.243)	(0.195)	(0.173)	(0.127)	(0.044)	
CV										
VLA	-20.566	0.474	2.770	0.047	-0.013	-0.379	0.504	-0.049	-0.271	1073
	(4.096)	(0.464)	(0.411)	(0.122)	(1.544)	(0.280)	(0.223)	(0.181)	(0.050)	
NHNBW	-14.575	0.729	1.995	0.218	-2.306	-0.009	0.338	0.081	-0.160	976
	(3.661)	(0.385)	(0.355)	(0.104)	(1.323)	(0.223)	(0.081)	(0.158)	(0.042)	
NW	-14.710	0.341	1.806	0.314	-3.379	0.316	0.376	0.020	-0.078	997
	(3.169)	(0.351)	(0.293)	(0.097)	(1.233)	(0.195)	(0.172)	(0.125)	(0.044)	

Notes:

1) VLA - Very Liquid Assets; NHNBW - Non-housing, Non-business Wealth; NW - Total Net Worth
   EPP - Equivalent Precautionary Premium; CV - Coefficient of Variation.
2) The measure of income uncertainty is EPP in the first three equations and CV in the second three equations.
3) All wealth, income, and uncertainty values are in logs.
4) All observations with negative reported wealth are excluded.
5) Heteroscedasticity-robust standard errors in parentheses.
6) The set of instrumental variables is described in the text.

Table 5

Reduction in Average Wealth for Fixed Equivalent Precautionary Premiums			
EPP Fixed at Percentile:	Very Liquid Assets	Non-housing, Non-business Wealth	Total Net Worth
0	0.2414	0.3928	0.2577
10	0.1357	0.2348	0.1436
25	0.0842	0.1510	0.0857
50	0.0200	0.0441	0.0200

Each cell contains the percent reduction in average wealth holdings were each household to face an EPP equal to the value of the EPP at the specified percentile of the EPP distribution. See text for more discussion.

Table 6

Simulated Reductions in NHNBW by Income and Wealth Deciles		
	Reduction in NHNBW level by decile when EPP is fixed at its minimum predicted value:	
Percentile Range	Permanent Income	Non-housing, Non-business Wealth
0 - 10	0.5200	0.3953
10 - 20	0.4567	0.3878
20 - 30	0.4599	0.3510
30 - 40	0.4383	0.3556
40 - 50	0.3765	0.3571
50 - 60	0.4044	0.3490
60 - 70	0.4345	0.3478
70 - 80	0.4047	0.3445
80 - 90	0.2783	0.3609
90 - 100	0.3677	0.4200

Each cell contains the average reduction in NHNBW holdings by decile were each household to face the minimum predicted EPP instead of its actual EPP.

Figure 1a

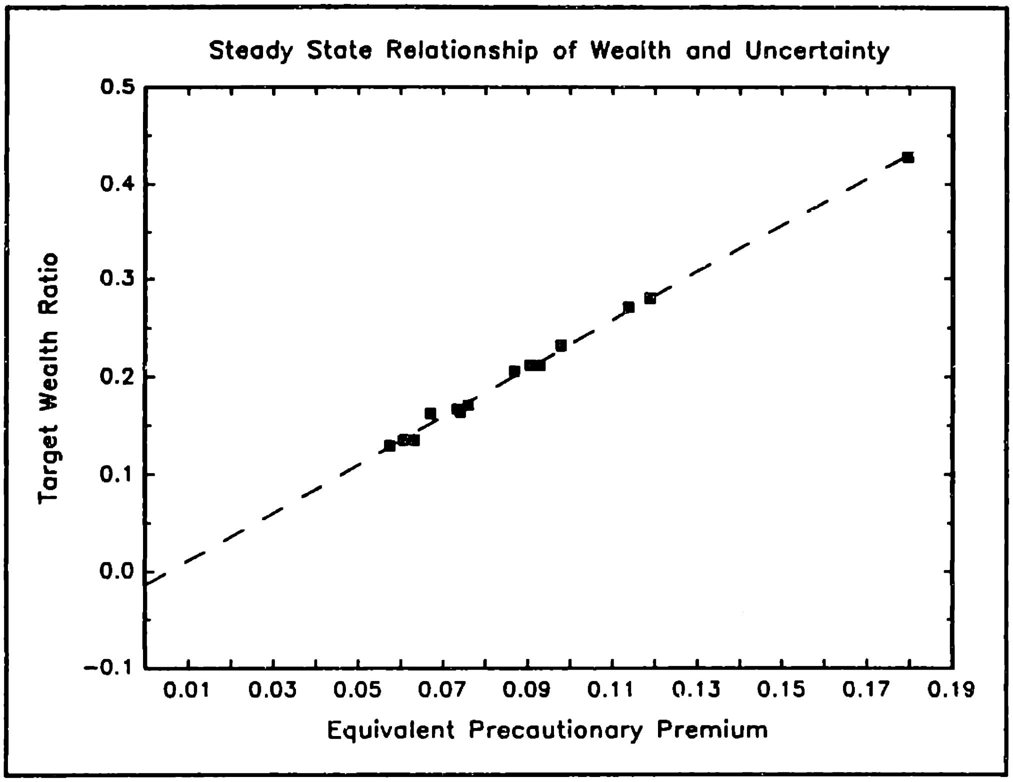

Figure 1b

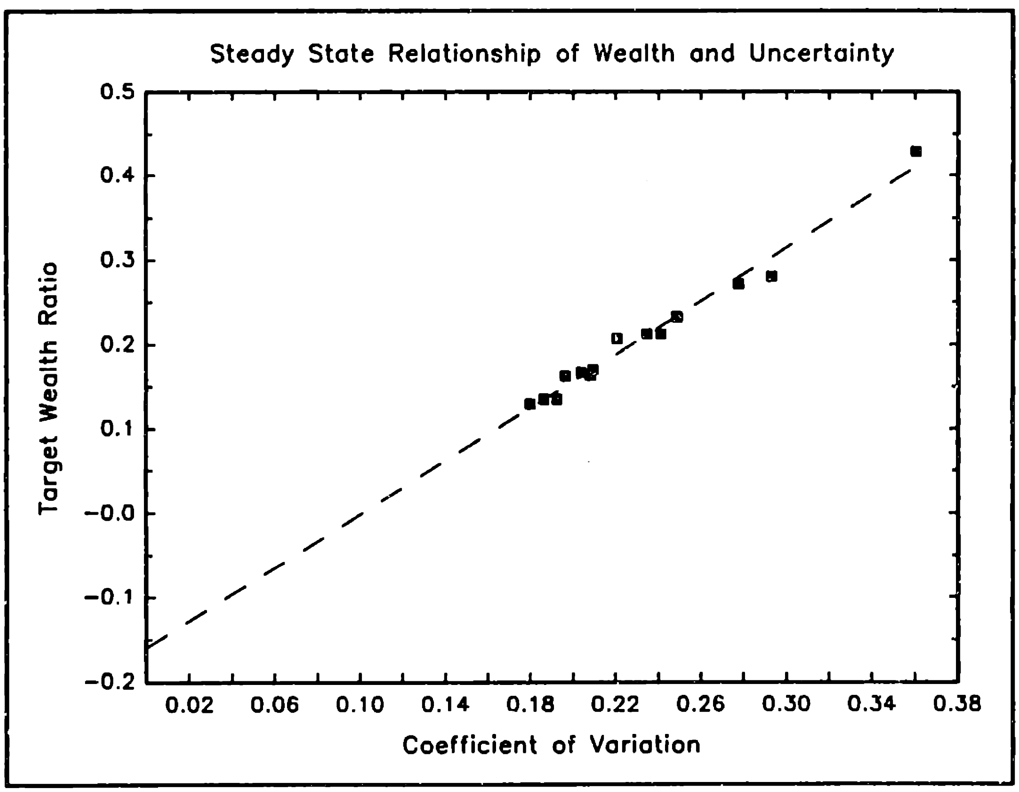

Figure 2

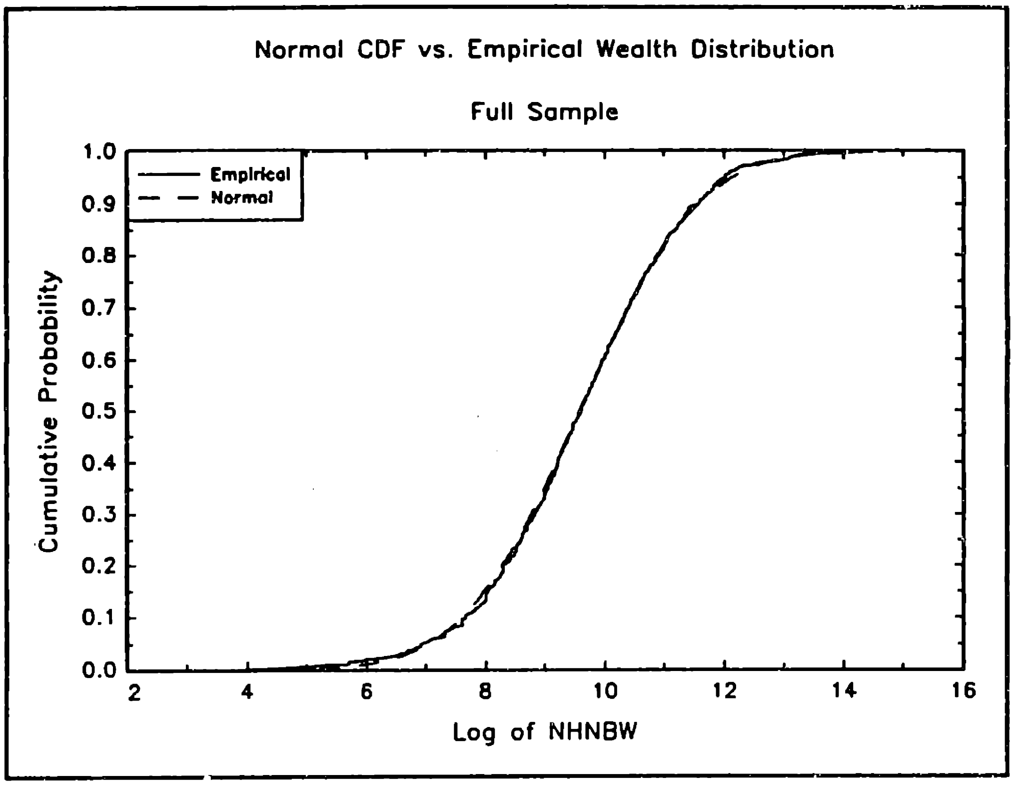
Normal CDF vs. Empirical Wealth Distribution
Full Sample
Empirical
Normal
Cumulative Probability
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
2
4
6
8
10
12
14
16
Log of NHNBW

Figure 3a

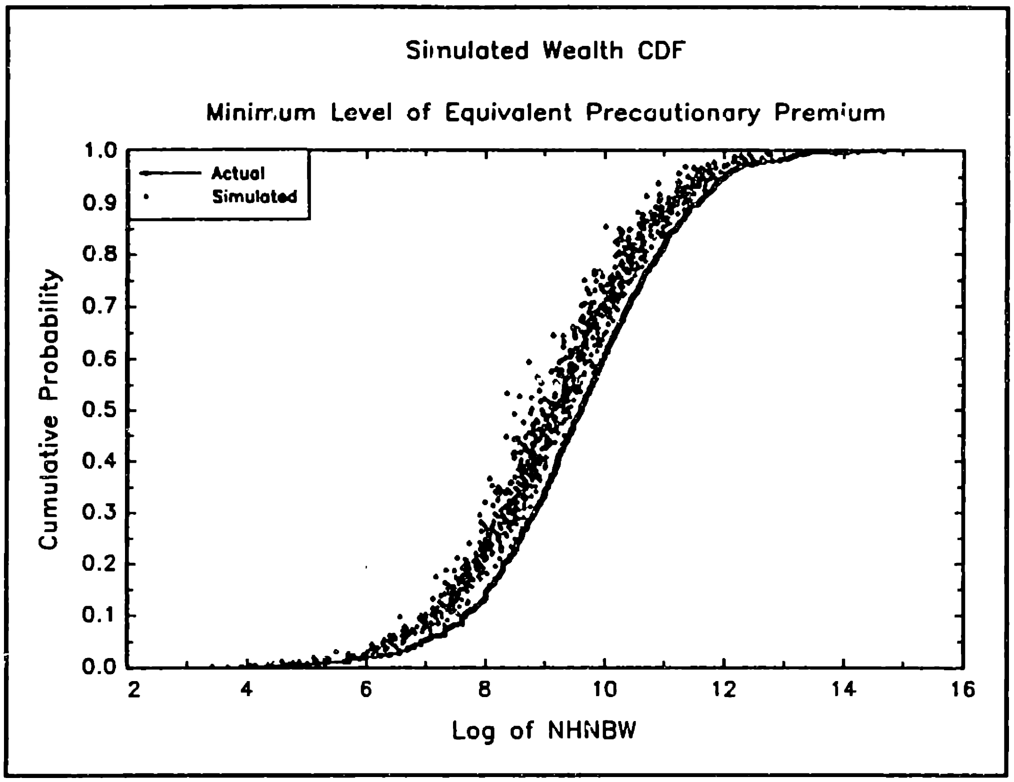

Figure 3b

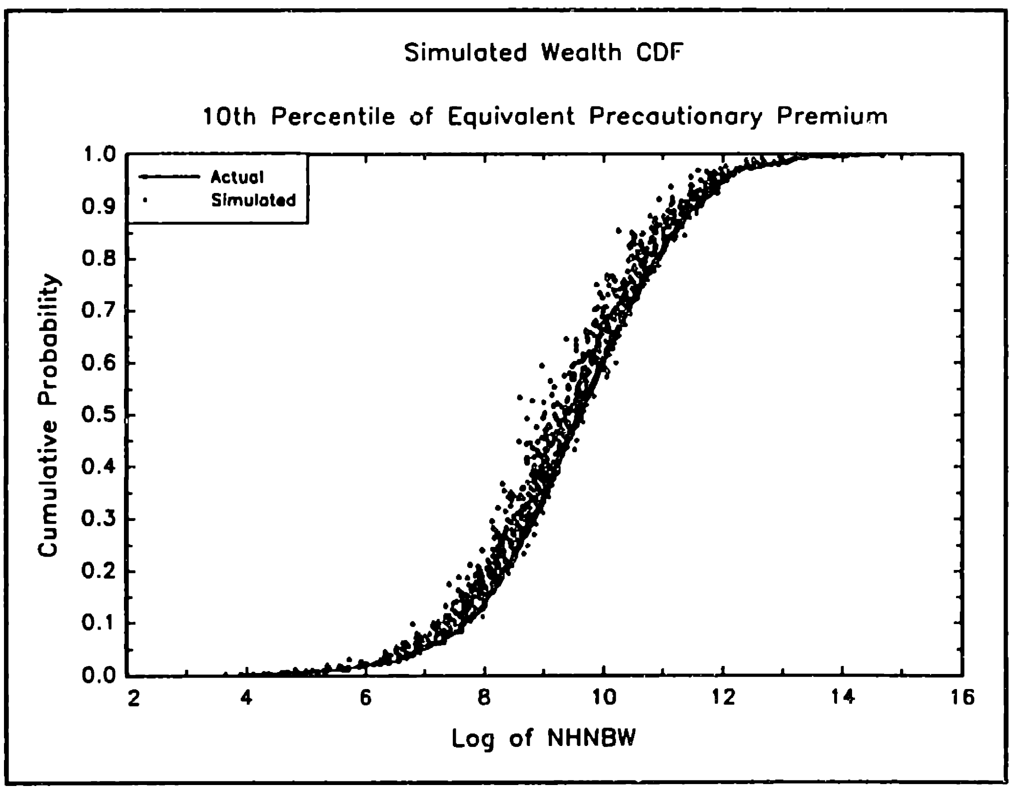

Figure 3c

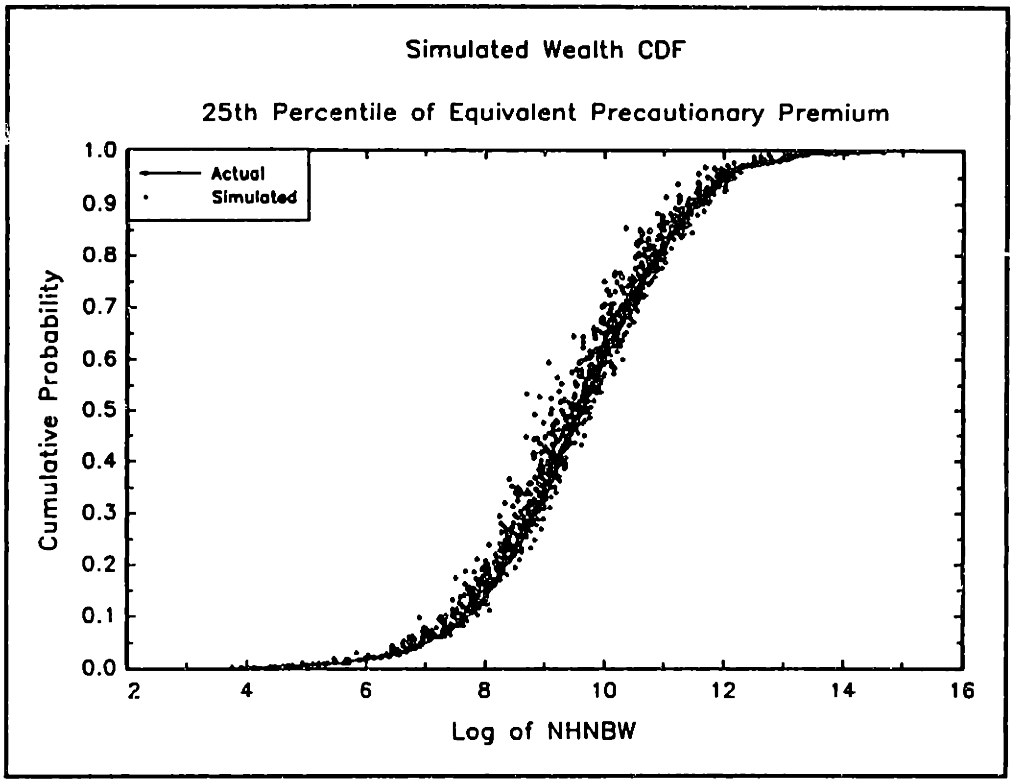

Figure 3d

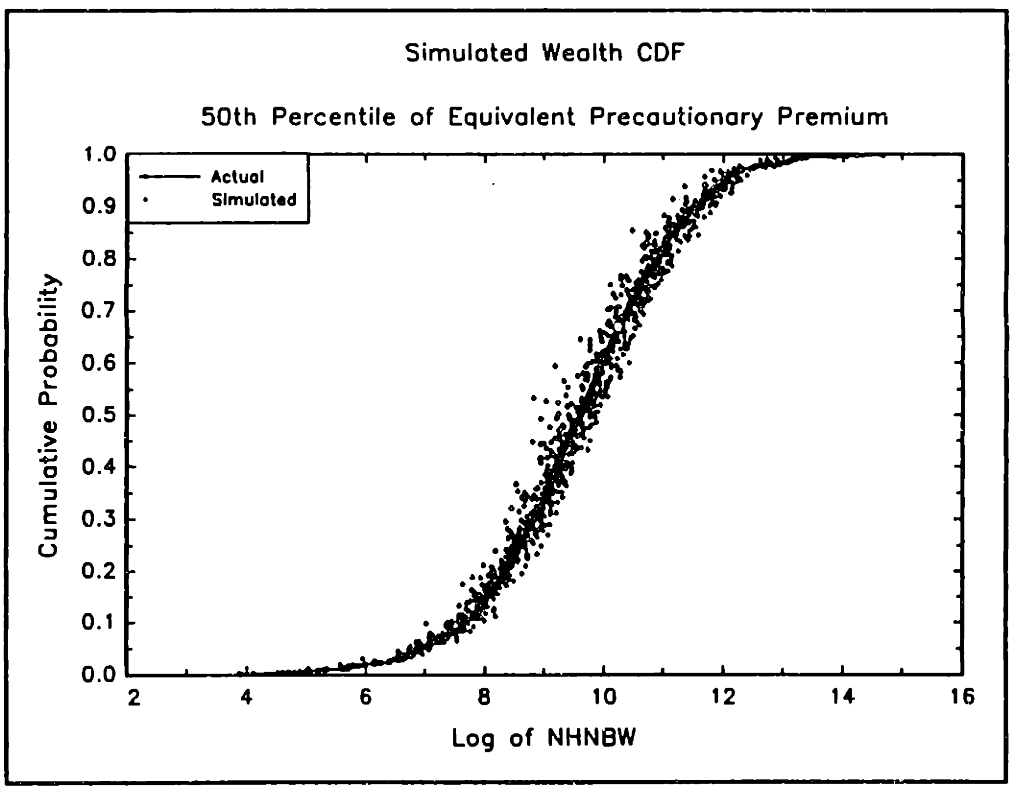

Chapter Three

Retirement Incentives in the 1983 *Pension Provider Survey*

I. Introduction

The *Pension Provider Survey* (PPS), a supplement to the Federal Reserve Board's *Survey of Consumer Finances 1983* (SCF), represents a significant advance in the scope of household economic surveys. For each worker in the SCF households who claimed to be covered by a pension plan, an attempt was made to survey the entity responsible for providing the pension benefits (usually the employer). The full summary plan descriptions for up to three plans under which each worker was covered were obtained. The level of detail in the PPS is such that the pension benefits a worker will receive conditional on a broad range of potential economic environments and retirement dates can be computed. As such, it improves upon the quality of the self-reported pension data found in the SCF and other household surveys.[1]

By linking the two surveys, it is now possible to study economic agents with more extensive knowledge of their retirement opportunity sets. The detailed information on pension formulas in the PPS allows fundamental questions of saving and labor force behavior to be revisited with more powerful tests by constructing better estimates of pension wealth. Additionally, such a survey encourages new hypotheses regarding the economic ramifications of pension coverage. As a prelude to two such papers (Samwick 1993a, 1993b) and a guide for others, this paper provides background information on the pension plans in the PPS. Section II describes the nature of censoring in the PPS; that is, the extent to which the plans in the PPS are representative of those in the population.

---

[1]The PPS for the 1983 SCF was the first attempt to match the details of workers' pension plans with their other personal and economic characteristics in a nationally representative dataset. This effort was repeated for the 1989 SCF, and similar pension supplements are currently proposed for the *National Longitudinal Survey of Mature Women* and the *Health and Retirement Survey*.

Section III reviews the literature that has utilized the PPS and argues that its most important resource--the broad diversity of pension plan features--is yet to be tapped. Section IV illustrates and quantifies this considerable heterogeneity in pension formulas found in the PPS. Section V discusses directions for further research and concludes.

II. Censoring in the PPS

The usefulness of the PPS for economic analysis will depend on how well it represents the pension plans that cover workers in the population. Of the 2,261 people in the SCF's 4,303 households who reported being covered by one or more pension plans, 1,641 or 73 percent were successfully coded in the PPS for a total of 1,011 different pension plans in the survey. The official documentation of the PPS in Curtin (1987) shows that the attrition is due in roughly equal parts to obtaining permission from the individuals to contact the provider, making contact with the provider once permission is granted, and coding the plan formula details after the interview. The bottom row of Table 1 presents population-weighted data on the nature of the sample in the PPS. Of the 86.194 million workers in 1983, 42.240 million or 49 percent reported being covered by a pension in the SCF. Of these workers, 32.057 million or 76 percent have their pension plans represented in the PPS.

Table 1 also contains breakdowns of the workforce, pension coverage, and PPS representation by employer type. The first column reports the share of the total workforce employed by each type of employer. The next column gives the fraction of workers of a given type that report being covered by a pension in the SCF. Pension coverage is eighty percent or more among workers employed by governments and public schools or colleges. The third column shows that roughly the same shares have their pensions represented in the PPS for these employer types, making for take-up rates of about ninety percent or more.[2] Although they comprise a very small fraction of the workforce and have a coverage rate of only about forty percent, private school (including elementary, secondary, college, and university) workers also have a very high take-up rate in the PPS. Private sector workers, especially those in smaller firms, have lower rates of pension coverage and lower take-up rates in the PPS than do public sector workers. Coverage is close to seventy percent among workers at firms with more than 100 employees, and over seventy percent of them are picked up in the PPS. Among workers in firms with less than 100 employees, however, only about a fifth are covered and slightly more than half of them are picked up in the PPS.

These differences can be attributed to the inability of the SCF sponsors to locate the pension provider or to code in the pension plan details from the information given by the provider. For public sector workers, the plans are fairly similar and the employer is easy to locate. Private sector workers may have been unwilling to give the name of the employer and the employers themselves may have had less information on hand to aid in the interview. This would be especially true for workers at smaller firms.[3]

Thus, the effect of attrition between the SCF and the PPS is to remove about one quarter of the plans from consideration and to over-represent public sector workers at the

[2]The take-up rate is simply one minus the attrition rate between the SCF 1983 and the PPS.

[3]For military personnel, the take-up rate is only about 10 percent despite the high coverage rate and the limited number of military plans. One possible explanation for this is that the PPS can only include information on pensions from one job, which will be a private sector job if the individual has ever had one.

expense of private sector workers, especially those at small firms. This caveat should be borne in mind whenever the PPS is used.

III. Existing Research Using the PPS

Despite the richness of the data provided by the PPS, only a handful of papers have used it to study economic issues. The first papers to make use of the PPS compared the self-reported pension entitlements in the household data of the SCF with the entitlements given by the full plan description in the PPS to determine how accurately individuals forecast their own pension benefits. Curtin, Juster, and Morgan (1989) find that the median values for *initial* annual benefits reported by the survey respondent are 83 percent of the corresponding amounts under the benefit rules in the PPS. They also discuss methods for imputing the value of pension entitlements from generally accessible household data on industry and occupation combined with limited information on pension eligibility. McDermed, Clark, and Allen (1989) use the PPS to compute the share of total household net worth embodied in future pension benefits. Their estimates range from 17.4 to 31.4 percent, depending on the assumed length of future employment and the method of forecasting earnings. These numbers represent considerably higher average pension entitlements than those computed by Avery, Elliehausen, and Gustafson (1985) using the self-reported data in the SCF. Thus, the calculations using the PPS suggest that most people's actual entitlements are larger than they believe.

Another set of papers uses the PPS to study the relationship between pensions and other forms of labor compensation. Montgomery, Shaw, and Benedict (1992) estimate hedonic price equations for wages to measure the implied tradeoff between pensions and wages. They find a one-for-one pension-wage tradeoff over a worker's lifetime and attribute their results, which differ from those of previous researchers, to the higher quality of the pension data and their focus on lifetime (instead of annual) wages. In later work, Montgomery and Shaw (1992) use a similar framework and the theory of compensating differentials to find evidence of wage-premia in unionized and large firms.

Two papers by Gustman and Steinmeier (1989a, 1989b) use the PPS to study labor force participation. In the former, they examine the incentives that pensions create for labor turnover and relate these incentives to plan characteristics. They show that the strongest incentives to remain with or leave the firm derive from the conditioning of early or normal retirement benefits on job tenure, a point also made by Kotlikoff and Wise (1987, 1989). In the second paper, they rely on the better data in the PPS to control for pension wealth when simulating the effects of various policy reforms of the Social Security retirement earnings test and delayed retirement credits.

A common attribute of the existing literature using the PPS is that none of the papers use any information from the survey beyond the (improved) calculations of pension wealth in any econometric analysis. If the study conducts any econometric tests, it is only on the level of pension benefits under one economic scenario, whereas if it considers pension wealth in different circumstances, it is merely descriptive. In later work, I use the ability to compute the *evolution* of pension wealth over an employee's tenure to estimate the effect of pension benefits on retirement behavior (Samwick 1993a) and the ability to compute pension wealth under different realizations of a stochastic wage process to identify patterns of risk-sharing between firms and workers (Samwick 1993b).

IV. Pension Incentives for Retirement

Recent studies have demonstrated the magnitude and heterogeneity of pension entitlements present in the workforce. Kotlikoff and Wise (1987) use the 1979 BLS *Level of Benefits Survey* to examine the incentive effects of private pension plans on retirement. They find that pension plans provide substantial incentives to retire just after early and normal retirement ages. Gustman and Steinmeier (1989a) study plans in the PPS and conclude that the most important retirement incentives derive from the requirements for normal retirement benefits. To give some background on the components of pension formulas and highlight the relationship between pension provisions and retirement incentives, the graphical method of representing pension incentives in Kotlikoff and Wise (1987) will be adopted here. Particular emphasis will be placed on the heterogeneity in pension entitlements generated by features of the plans that would be missed in a study which used less detailed pension data in economic analysis. In particular, the retirement incentives of different pension plans are shown to be sensitive to tenure, the wage level, the form of the pension (defined benefit versus defined contribution), and the level of Social Security benefits in addition to the plans' average replacement rates.

Data

The approach taken here will be to analyze the distribution of pension entitlements for all pension plans in the PPS using representative individuals from the SCF. Focusing on the differences in pension benefits for the same individual, rather than the actual individuals who are covered by the plans, removes the component of the variation in pension benefits that is due to differences in age, tenure, and wages across individuals. The residual variation is due solely to the heterogeneity in the types of pension plans covering workers.

To aid in the selecting of representative individuals, the first column of Table 2 presents descriptive characteristics relevant to calculating pension benefits for the individuals in the SCF that are between ages 21 and 70 and covered by at least one plan in the PPS. The important statistics are the average age of 42, average tenure of 11 years, and average wage of $22,880.[4] The first hypothetical worker reflects these average characteristics and is given in the fourth column. The second and third columns modify this hypothetical worker by moving his hire date 20 and 10 years into the future, respectively. By comparing the distribution of pension entitlements for these three workers, the effect of tenure on pension entitlements can be isolated. Similarly, the fifth and sixth columns modify the first hypothetical by doubling and halving his wage, respectively. By comparing the distribution of pension entitlements for the hypothetical individuals in the fourth, fifth, and sixth columns, the effect of the wage level on pension entitlements can be isolated.

[4]As described in the Appendix, sex, race, and education (along with occupation) affect the computation of pension benefits only through their use in the forecasting of future wages. Because there is no sense in which a "mean" occupation can be defined, the wage forecasts for the representative individuals will not be based on equations that control for occupation.

Replacement Rates

A useful starting point is to consider the fraction of pre-retirement wages a pension would immediately replace if an employee retired at his current age. Such replacement rates for each representative worker are presented in Table 3. For all workers, replacement rates are increasing with age, and the average rate of increase between ages 55 and 65 varies between 1.8 and 3.0 percentage points per year. Defined benefit pension plans typically increase benefits separately with both age and tenure. Defined contribution plans will naturally have replacement rates that increase only with tenure.

The increase in average replacement rates is not smooth with age, however. For the worker with average characteristics, the average replacement rate more than doubles from 7.1 percent to 15 percent at age 55. This jump reflects the prevalence of age 55 as an early retirement age at which benefits can first be received (for these calculations, the replacement rate on a defined benefit pension plan before the early retirement age is zero). The smaller jump at 56 reflects new benefit schedules that arise at 25 years of service. Some plans allow for extra benefits as soon as age plus tenure equals some value; the slight increase at age 58 is the result of new benefit schedules that arise when age plus tenure first exceeds 85. Finally, the relatively large increases at age 60 and 65 are the result of many plans first offering normal retirement benefits at those ages.

To consider the effects of tenure separately from those of age, we can compare the replacement rates for workers of a given age with different tenures. At age 55, having worked with the firm from age 41 to 51 and from 31 to 41 increase the average replacement rates by 6 and 8.4 percentage points, respectively. At age 50, the corresponding increases are 13.9 and 13.5, and at age 65 they are 13.5 and 11.7. Thus, the pure effect of tenure on the replacement rate is approximately 1 percentage point per year.

As described below, pension benefits are generally proportional to wage levels over large ranges of wages. There are some plans, however, which increase replacement rates with the wage level by integrating with Social Security as well as some which do not increase benefits with wages at all. To consider the relative importance of such plans in the aggregate, we can compare replacement rates at given ages for different levels of the wage. At all ages, average replacement rates are declining with the wage level, suggesting that Flat Rate pensions exert more of an influence on replacement rates than do integrated plans.

The standard deviations of replacement rates by age are generally as large as the replacement rates themselves, suggesting ample heterogeneity in these rates. Figure 1 displays a stacked bar graph of the distribution of replacement rates for ages 55 to 65 for the worker with average characteristics. At age 55, two thirds of the plans have replacement rates below 0.2, one quarter have rates between 0.2 and 0.4, and just under 10 percent have replacement rates higher than 0.4. At age 60, only three eighths have rates below 0.2, about one third have rates between 0.2 and 0.4, and almost 30 percent have rates over 0.4. By age 65, just over 20 percent have rates below 0.2, about 30 percent have rates between 0.2 and 0.4, about 20 percent have rates between 0.4 and 0.6, and the remaining 30 percent have rates over 0.6.

Figure 1

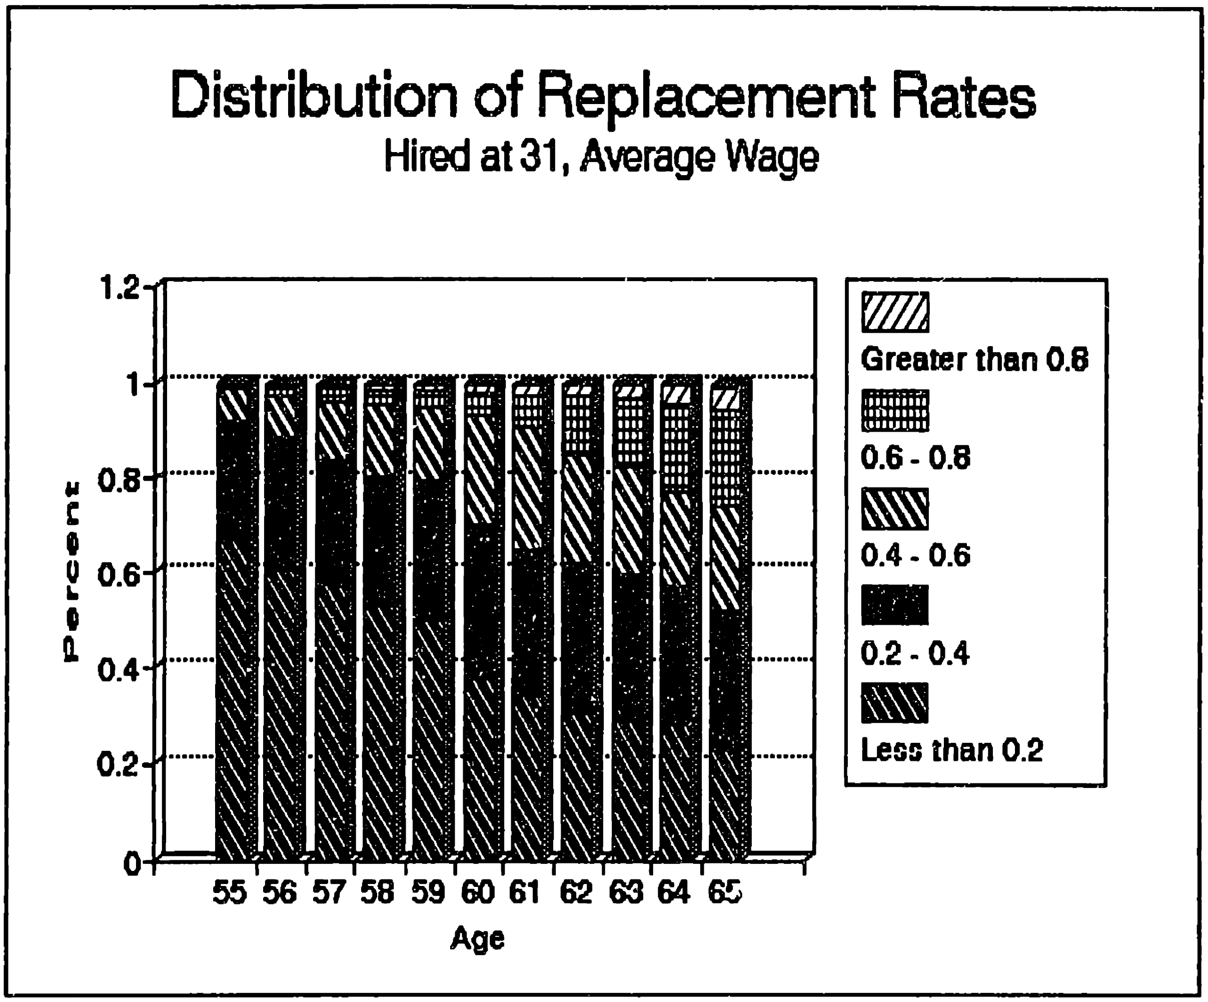

Accrual Rates

Even the distribution of the replacement rates in Figure 1 conceals some important differences in pension entitlements. For ease of exposition, the calculations in Figure 1 and Table 3 ignore the effect of aging on how valuable an initial benefit is to the worker.[5] That is, there is a tradeoff between postponing retirement for a year to obtain a higher replacement rate and losing a year of retirement in which to collect them. A

[5]They also exclude the contribution of vested benefits if they cannot be taken immediately. The present value framework adopted below does not suffer from this problem.

more general framework is to consider the actuarial present value (APV) of benefits conditional on retiring at each age.

To fix ideas, denote by I(a) the difference in the worker's APV of pension benefits at ages a and a-1, with the earlier value suitably increased by the interest rate to account for the 1 year delay in receipt. Let the ratio of I(a) to the current wage, W(a), be given by R(a). That is,

$$I(a) = APV(a) - (1+r)APV(a-1)$$

$$R(a) = \frac{I(a)}{W(a)}$$

Underlying any model of retirement behavior is the notion that a worker retire when the opportunity cost of doing so is low. A basic component of this opportunity cost is the level of the wage; if wages are high, then the price of retirement is high as well. I(a) is analogous to a wage in the sense that the worker's lifetime wealth is augmented by this amount by working at age a; it is a form of compensation for being employed at that job. When I(a) is high, then postponing retirement for one more year yields a relatively large increase in lifetime wealth per year of foregone leisure.

When comparing the value of I(a) across individuals, it is better to have a scaleless measure so that differences in the retirement incentives are not merely the result of different wage levels, as all pensions are (weakly) monotonic functions of the wage level. R(a) is a natural way to index the generosity of a pension plan to workers of age a relative to their wages. For example, if R(a) = 0.2, then the effect of the pension on total compensation at age a is to increase wages by 20 percent.

Figure 2 graphs the average accrual rates by age for the worker with average characteristics, along with schedules representing the top and bottom 10 percentiles of the distribution of accruals at each age. The solid line represents the weighted average of R(*a*) for all plans in the PPS for this worker. The large spike at age 55 is indicative of most pensions having that as the early retirement age; staying on until this age results in a pension accrual equal to 62 percent of the current wage for the average worker. The small peaks at ages 50, 58, 60, 62, and 65 represent the tendency for some plans to have these ages as the first at which benefits can be taken or the first ages at which a more generous benefit schedule is made available. The dashed and lines represent the ninetieth and tenth percentiles of pension accruals at each age. The former shows that 10 percent of the plans have accruals over two times the average at most ages, including the spike at age 55. The latter shows that after the normal retirement age of 65, at least 10 percent of the plans have negative accruals to lifetime wealth associated with the pension.

Thus, holding everything else constant about the employee, there is substantial variation in the level of pension accruals, especially after the Early Retirement Age. Figure 3 focuses on the pension accruals between ages 55 and 65 in more detail. At age 55, 46 percent have zero or negative pension accruals, 21 percent have accruals between 0 and 0.1, 25 percent have accruals between 0.1 and 0.3, and 8 percent have accruals over 0.3. At age 60, only 17 percent have zero or negative pension accruals, 20 percent have accruals between 0 and 0.1, 33 percent have accruals between 0.1 and 0.3, 24 percent have accruals between 0.3 and 0.5, and 6 percent have accruals greater than 0.5. Finally, at age 65, 10 percent have zero or negative accruals, 13 percent have accruals

Figure 2

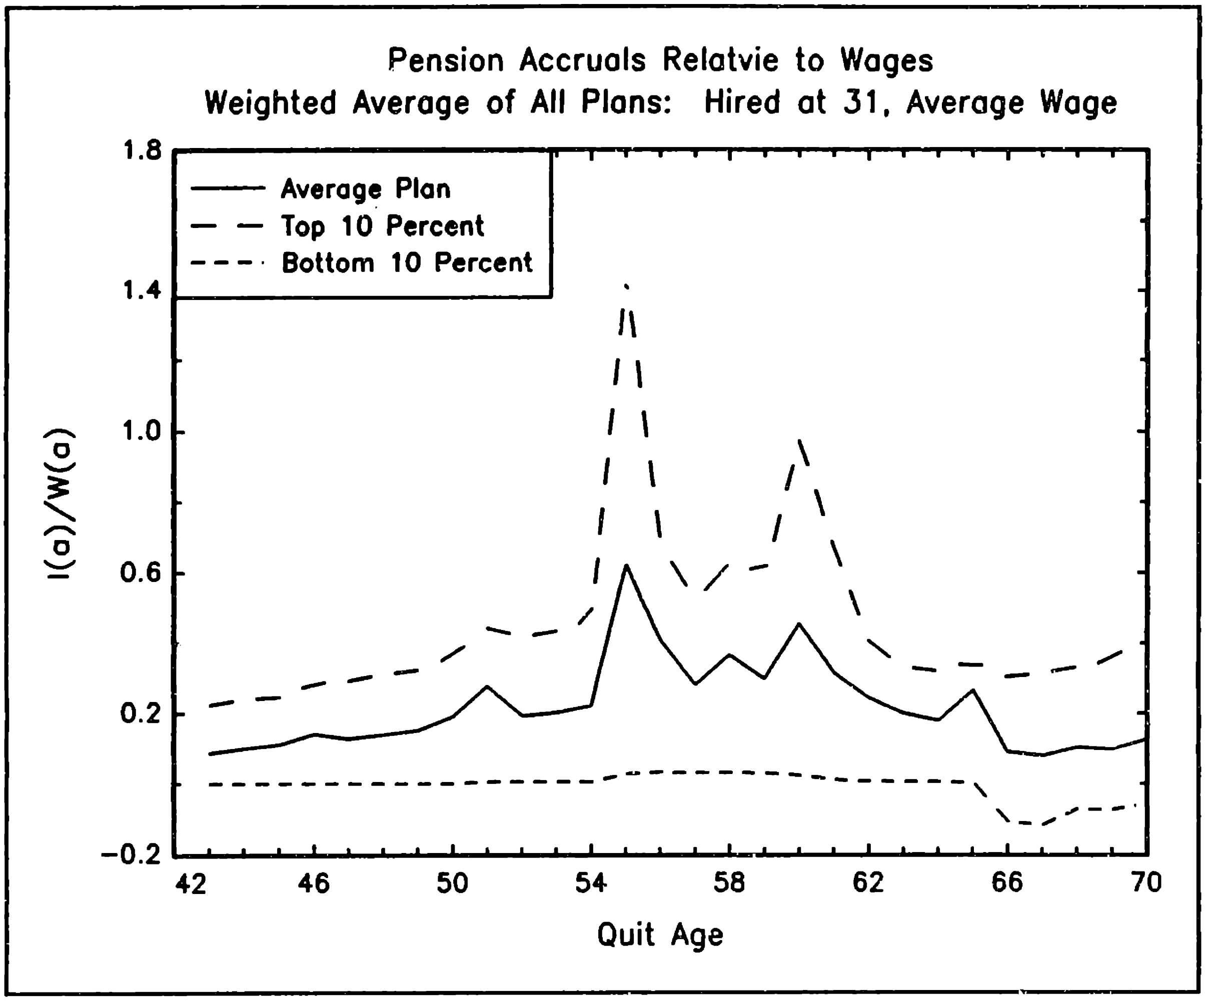

between 0 an 0.1, 29 percent have accruals between 0.1 and 0.3, 23 percent have accruals between 0.3 and 0.5, and fully 25 percent have accruals greater than 0.5.

Figures 2 and 3 demonstrate that even with regard to the same individual, the many different pension plans in the economy give rise to quite diverse incentives for retirement which differ systematically by age. The variation in incentives at the individual level over time derive from the eligibility requirements for early and normal retirement benefits. The next several subsections use typical pensions found in the PPS to show that when comparing accrual rates across persons with different economic characteristics, another nontrivial source of variation is introduced as economic

Figure 3

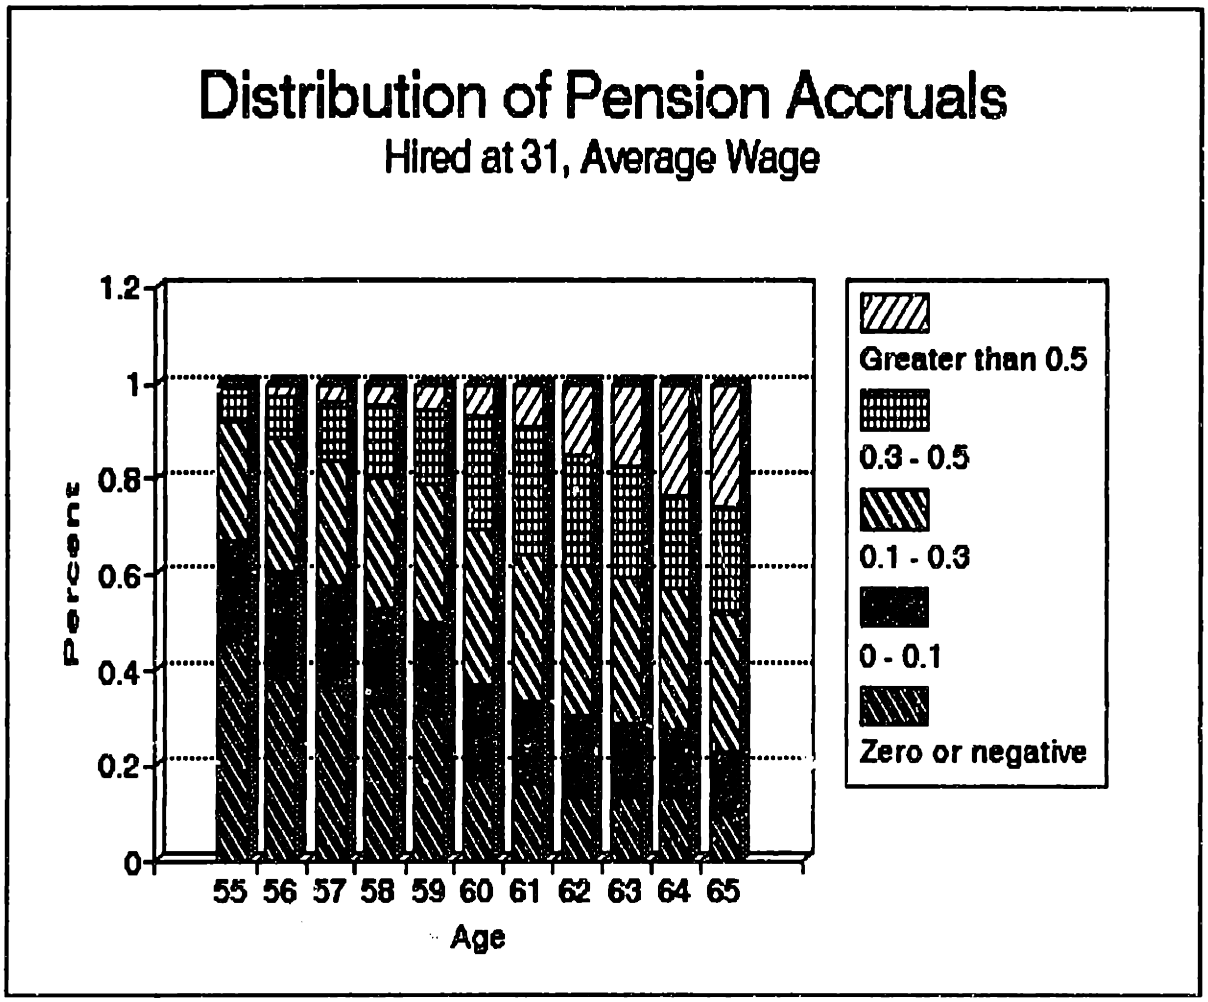

characteristics interact with specific pension plan features.

The Effect of Tenure on Pension Accrual Rates

The most important source of variation in pension accruals by age is due to years of service. Figure 4 plots the accrual schedules of the three hypothetical workers who differ in their dates of hire. The solid line represents R(a) for a worker hired at age 31 over his tenure at the firm. This pension plan is a Defined Benefit (DB) plan: the benefits the worker receives are determined by a specific formula that depends on years of service, age, and average earnings. There are three spikes in this schedule corresponding to three important features of the plan. The first jump always occurs at the age of vesting. Before a worker is vested in a plan, he is not entitled to any retirement benefits if he leaves the firm. For this plan, vesting occurs after 10 years of service.[6] At age 41, this worker is legally entitled to retirement benefits which begin at some future date; in this case, vested benefits can be received first at age 55. Thus, the accrual at this age equals the APV of the benefits beginning at that time. The magnitude of this accrual is only about 20 percent of the current wage because future benefits are not indexed to inflation and the worker must wait 14 years to receive them.

The next spike occurs at age 55, which is the early retirement age (ERA) in this plan. The significance of the ERA is twofold. First, benefits can be taken immediately upon retirement, so there is no devaluation to due inflation *before* they are received. This feature can also be seen in the upward trend in R(*a*) as the date of vested benefit receipt approaches. Second, provided the employee stays on until the ERA, the benefits are reduced at a better than actuarially fair rate (6% here) for each year they are received before the normal retirement age (NRA) of 65. Vested benefits are reduced at an actuarially fair rate (about 8% for men at this age) for each year before the NRA that they are collected. The difference in reduction rates also contributes to the sharp increase at the ERA.

Like many plans, this one is comprised of two early retirement schedules. The one above is valid for any employee who is at least 55 years old with 10 years of service. If the employee has been with the firm for 30 years and is at least 62 years old,

[6]Tax Reform Act of 1986 reduced the vesting periods under all forms of vesting. In the case of cliff vesting (entitlement changing from 0 to 100 percent in one year), the maximum period is now 5 years. Under such a change, the first spike in the accrual graph would be at age 36 and would be smaller in magnitude. Total area under the accrual graph through age 41 would be unchanged, however.

Figure 4

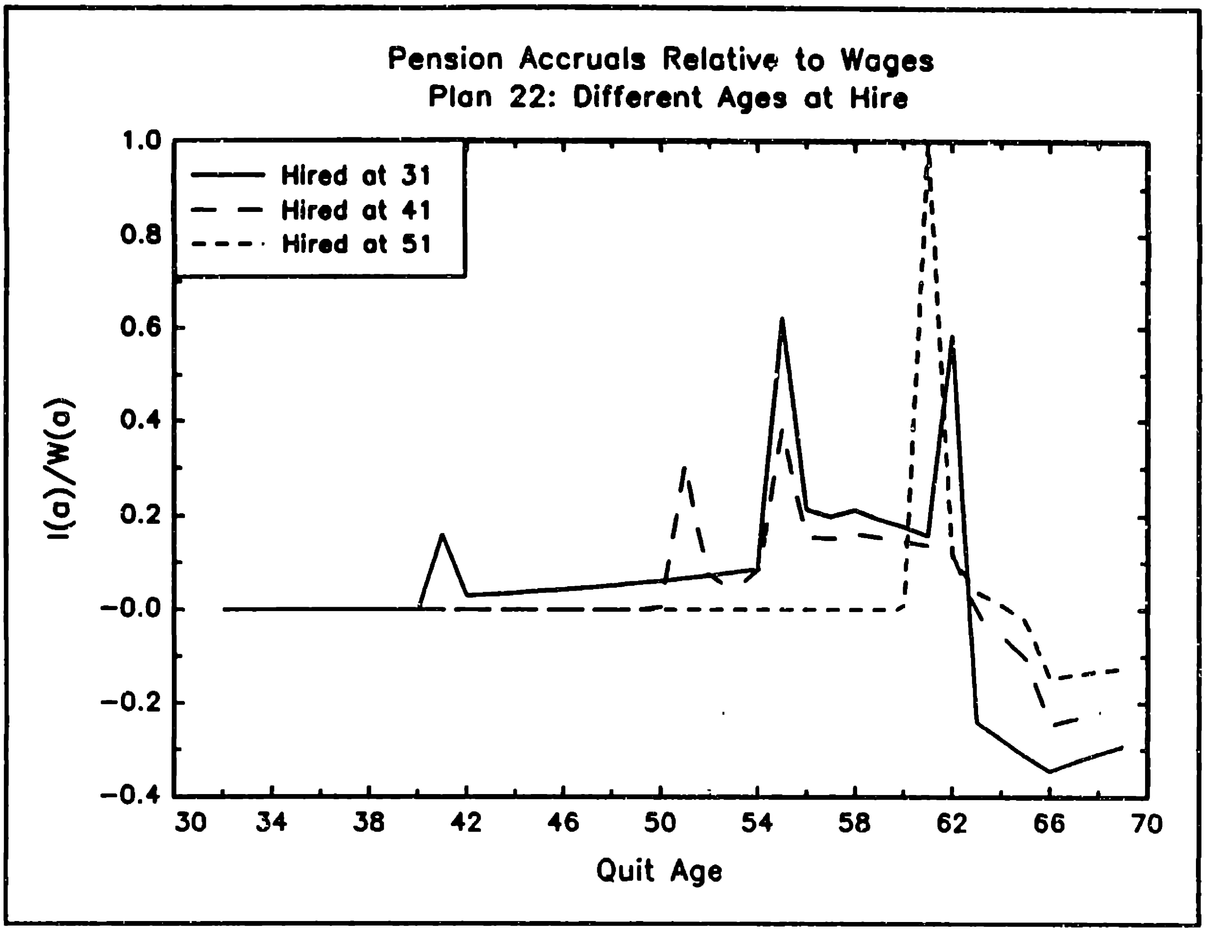

then another ER schedule becomes operative at that time. At this ERA, the reduction factor applied for years of benefit receipt before NRA is only 2 percent. The level of retirement benefits will increase dramatically when the more generous rate is applied; this accounts for the spike at age 62. Note that the wage equivalent of pension accrual at each ERA is approximately 60 percent. Also, accruals are consistently positive between the two ERA's, as the 6 percent reduction factor is applied one fewer time with each year of continued work. After the second ERA, pension accruals fall precipitously and remain negative. At this age, wages do not exhibit the same growth as they did earlier in the worker's career;[7] since Final Average Pay (FAP) is determined as the average of the last three years of wages, some of the incentive to stay with the firm is gone. Moreover, each extra year of work reduces the actuarial factor by only 2 percentage points. Thus, after the second ERA, there is no pension incentive to stay with the firm, and employees who choose to do so have the APV of their pensions reduced.

Thus, important increases in total compensation occur at specific ages in this pension plan, not all of which are perceptible in the aggregate graph depicted in Figure 2. Kotlikoff and Wise (1985) showed that these spikes were evidence for a contractual--as opposed to a spot--market of labor compensation because we do not observe declines in actual wages at those years. The dashed and dotted schedules in Figure 4 show that the timing of such spikes are related to the age at which the employee is hired, so that even workers under the same plan may have very different accrual patterns. The dashed line corresponds to an otherwise identical worker who joins the firm at age 41 rather than 31. The spike due to vesting now occurs at 51 and is larger due to the shorter delay in benefit receipt. The spike at the first ERA is smaller because the benefit formula depends on years of service, of which there are 10 fewer, and because vested benefits were not delayed (and therefore devalued in real terms) as much as before. Note that this worker will not accumulate 30 years of service until age 72, so that the second ER schedule never becomes operative. The pension accruals start to decline gradually as the 6 percent increase in initial benefit level for an extra year of work becomes less attractive to an aging worker.

---

[7]Age-earnings profiles were generated from cross-sectional wage equations estimated on individuals in the Current Population Surveys 1983-1986. This methodology is described in detail in the Appendix.

Finally, consider the worker who joins this firm at age 51. The only spike in his accrual is the one at 61 when he is first vested. Since he is immediately eligible for ER benefits, the spike is large--equal to his annual wage in this case. An interesting aspect of this plan is that the last significant spike in the accrual pattern occurs at age 55 for the worker hired at age 41, at age 61 for the worker hired at age 51, and at age 62 for the worker hired at age 31. Thus, if workers do indeed respond to these pension incentives, then retirement age will not even be a monotonic function of years of service.

Figure 4 shows that in addition to affecting the magnitude of accrual rates (by affecting the replacement rate), changes in tenure can change the timing of the largest spikes in the accrual schedule by confounding the ordering of vesting and eligibility for early and normal retirement schedules in a diverse group of workers. The first three columns of Table 4 present average accrual rates for the three hypothetical workers that differ only in the year of hire. For most ages before age 62, the worker with the highest tenure has the highest average accruals. At the years of vesting for the workers with later start dates, however, the highest tenure corresponds to the second highest accrual. Additionally, for ages after 62, there is virtually no net effect of tenure on average accrual rates, and there are at least two ages at which each worker has the highest accrual rate.

The graphs for the single plan and the aggregate figures in Table 4 suggest that there is no general relationship between tenure and accrual rates despite the strong interactions present in each plan. Knowledge of the details of the pension formula for each individual is essential to capture the effects of tenure on pension accrual rates.

The Effect of the Wage Level on Pension Accrual Rates

Figure 5

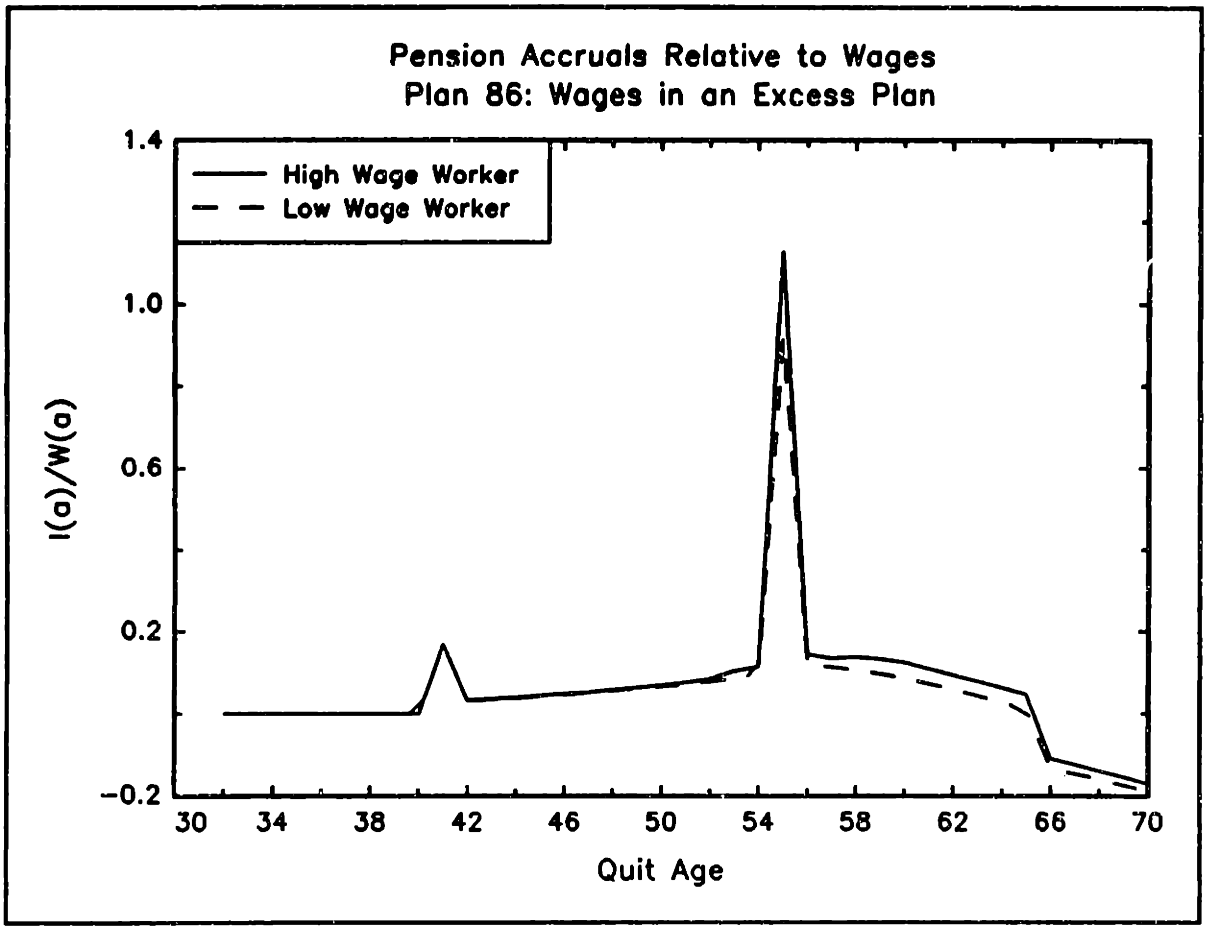

Figure 5 shows the effect of different wage levels on the accrual rates under a pension plan in which the replacement rate is a nonlinear function of the wage. First, note the usual incentive effects at the vesting age of 41 and the ERA of 55. The solid schedule is for the hypothetical worker making twice the average annual wage; the dashed line is for a worker making half the average annual wage. The nonlinearity arises at the Maximum Taxable Earnings (MTE) for Social Security in the year of benefit receipt, making this an integrated "Excess" plan. The MTE in 1983 was 35,700 and is indexed to the growth in average annual wages. The OASDI benefit formula does not depend on wages above the MTE; therefore, the Social Security replacement rate on earnings is a decreasing function of earnings.[8] To compensate for this, the pension plan credits average earnings below the MTE at 1.2% per year of service and those above the MTE at 1.5% per year of service. With 24 years of service, this differential is 7.2% of FAP. An Excess plan has the effect of smoothing out total replacement rates for high wage workers.[9] Note that there is no difference in the accrual rates at the age of vesting because the high wage worker's average earnings do not exceed the MTE until his early fifties. If there were no breakpoint in the formula (making this a simple "Percent of Earnings" plan), then the schedules would be identical.

Figure 6 demonstrates that it is also possible for low wage workers to receive higher accrual rates under a given pension than high wage workers. The plan represented here is a "Flat Rate" plan in which the benefits are a specified dollar amount multiplied by years of service. A worker's wage has no impact whatsoever on benefits; therefore, the higher wage worker necessarily receives a lower replacement rate. For these workers, the ratio of both wages and accruals is roughly 1 to 4. For the low wage worker, the spikes are substantial, but for the high wage worker, they are no more than 20 percent of current wages. Taken together, Figures 5 and 6 show that the pension incentives to retire in a sample of workers need not be monotonic in their level of wages,

[8]In general, the replacement rate for Social Security is a decreasing function of earnings, but it is not zero until the MTE is reached.

[9]Of course, the ability to pay high wage workers more may be the reason that these plans are provided. A typical article in a legal journal (Watson and Feutz, 1988) offered the following: "Historically, ... [F]or those able to untangle the maze of rules, however, the reward has been the ability to maximize deferred compensation for key employees, while minimizing the costs for providing similar benefits to the rank and file." and "Without the ability to integrate, many employers would not adopt qualified plans at all."

Figure 6

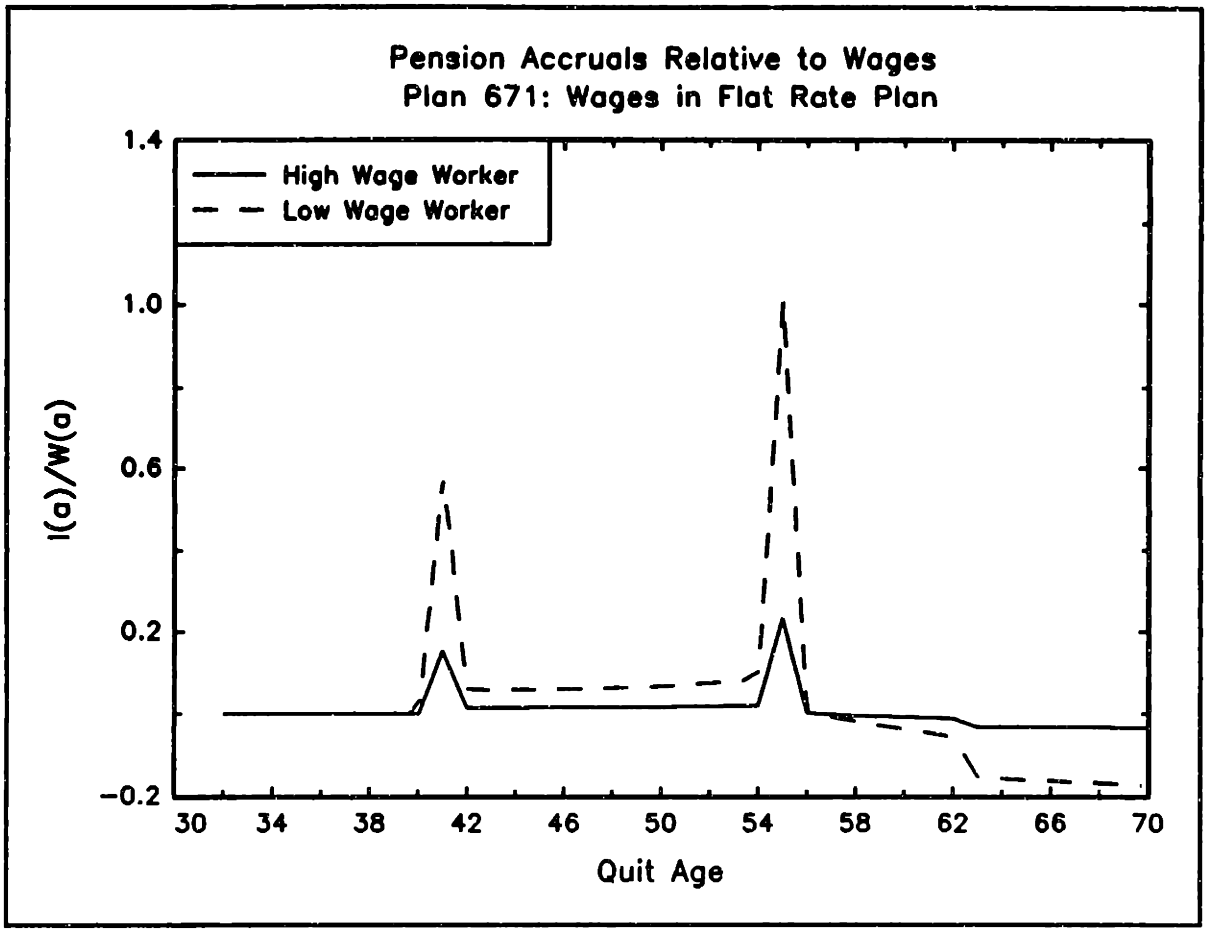

even if covered by a limited number of plans.

The last three columns of Table 4 demonstrate that the effect of wage heterogeneity on the distribution of pension benefits is mild compared to that of tenure. At all ages through 65, pension accruals are a declining function of the wage level, as the effects of Flat Rate plans outweigh those of integrated plans. At age 55, the accrual rates differ by 4.4 percentage points between the high wage and average wage workers and 3.4 percentage points between the average wage and low wage workers. At age 60, these differences shrink to 1.6 and 0.4, respectively. By age 65, they are 0.8 and 0, respectively.

The Effect of Plan Type on Pension Accrual Rates

Another major type of pension plan is a Defined Contribution (DC) plan, an example of which is shown in Figure 7. A DC plan is organized so that employers (and sometimes employees) make specified contributions to an account. Upon leaving the firm, the employee is entitled to some fraction of the accumulated assets in the account. For the workers in Figure 7, the assumed growth rate of the assets in the account is 7 percent per year. The three schedules correspond to workers hired at different ages. For each worker, there are three spikes in the accrual rates corresponding to the step-vesting schedule. After 1 year, the employee is entitled to 10% of the account; after 5 years, 50% of the account; and after 10 years, the employee can keep all of the money in the account. Thus, the only retirement incentives in a DC plan like this are related to vesting. Comparing the three schedules in the figure, only for workers who joined the firm fairly late will these spikes occur at ages when they might be considering retirement. This arrangement differs from the DB plan, under which almost all workers have large accruals at least some point in their later years. Thus, it is important to know what type of pension a worker has--not just that he is covered by a pension--when studying the economic effects of pensions. The first column of Table 5 shows that about ten percent of those with pensions have only a DC plan.

The Effect of Social Security on Pension Accrual Rates

The next graph, Figure 8, shows the effect of Social Security benefits on the accrual schedule in an integrated plan. A plan is integrated when some feature of the Social Security program affects the pension benefit formula. The Excess plan in Figure

Figure 7

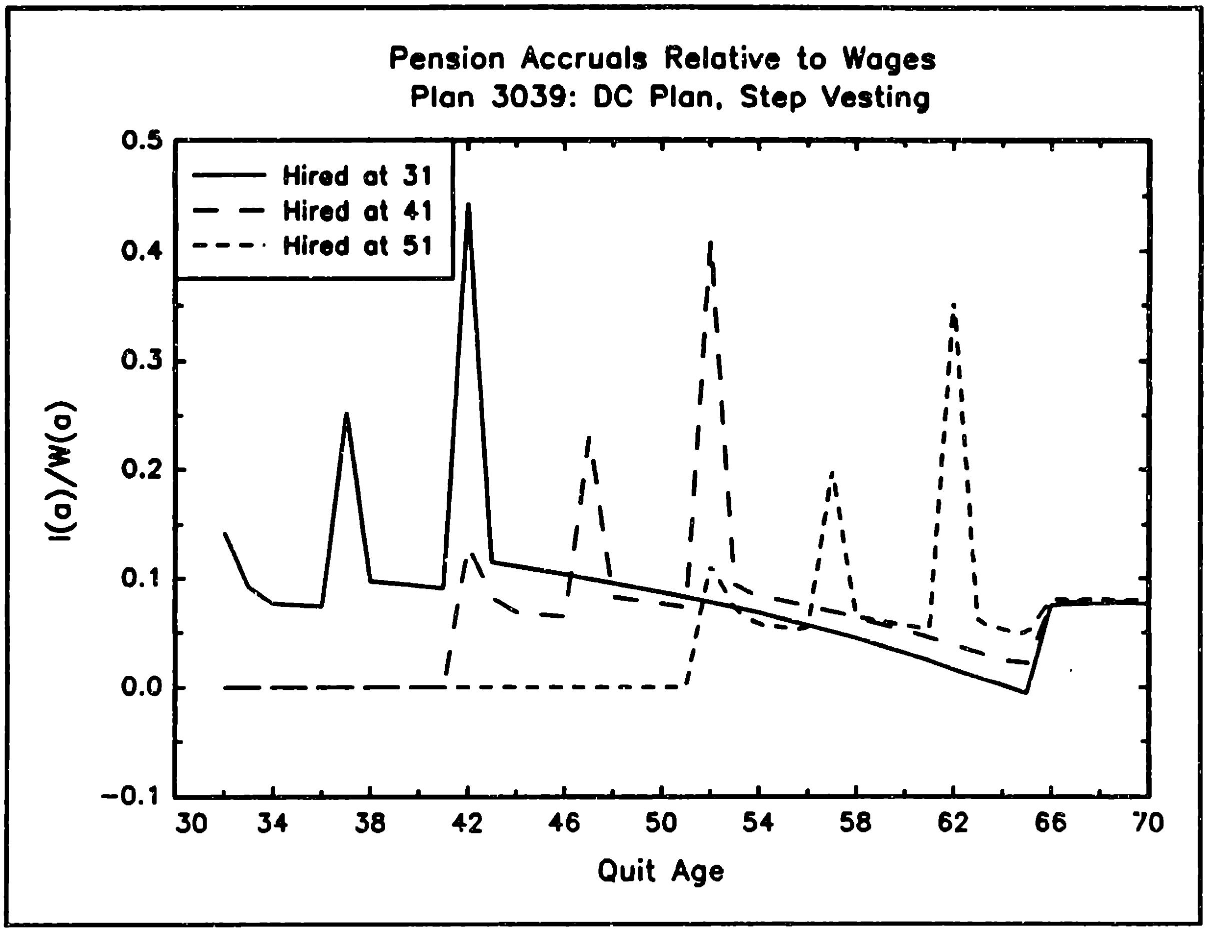

5 is one example; if the MTE were to decline, then the higher replacement rate above the MTE would apply to more earnings, thereby increasing pension benefits. Thus, integration serves to reduce the impact of changes in Social Security on the *total* level of retirement income.[10] The plan in Figure 8 integrates in a different way, by subtracting a fraction of the worker's expected Social Security benefits at age 65 from the amount paid to him by the pension. Any change in the Social Security program that influences the level of benefits will affect the *pension* accrual rates as well. Such plans

[10]Merton, Bodie, and Marcus (1987) and Bodie (1990) contain good discussions of the insurance aspects of integrating pensions with Social Security.

are called Offset plans because any change in Social Security benefits will be to son degree offset by an opposite change in the pensions.

Figure 8

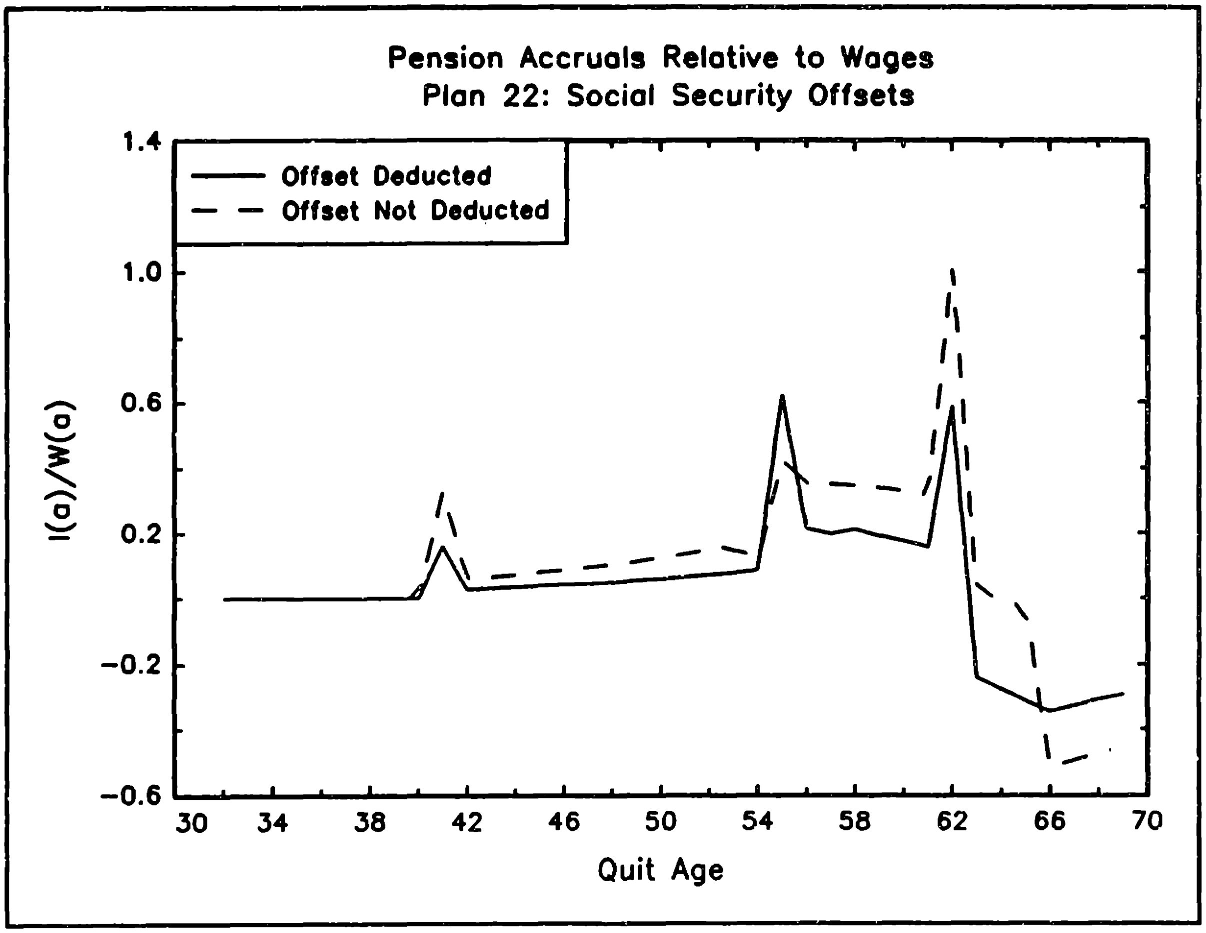

The solid schedule in Figure 8 is that of the worker hired at age 31 in the plan in Figure 4. The dashed line computes the pension accrual rates assuming that the offset component is not deducted, i.e. assuming Social Security benefits are zero. The timing of the spikes in accrual rates is not affected, but the relative magnitudes are different. Since overall benefits are higher, the spike at the vesting age of 41 and subsequent accrual rates (except at the first ERA) are larger without the offset. Because vested benefits were more generous without the offset, the opportunity to take benefits immediately at the first ERA is worth relatively less and the spike is smaller. Because the new wrinkle at the second ERA is only the better reduction rate rather than the lack of delay in benefit receipt, accrual there is also higher without the offset. For this particular plan, failing to account for the offset overstates the pension's financial incentive for the worker to stay on past the first ERA.

Figure 9

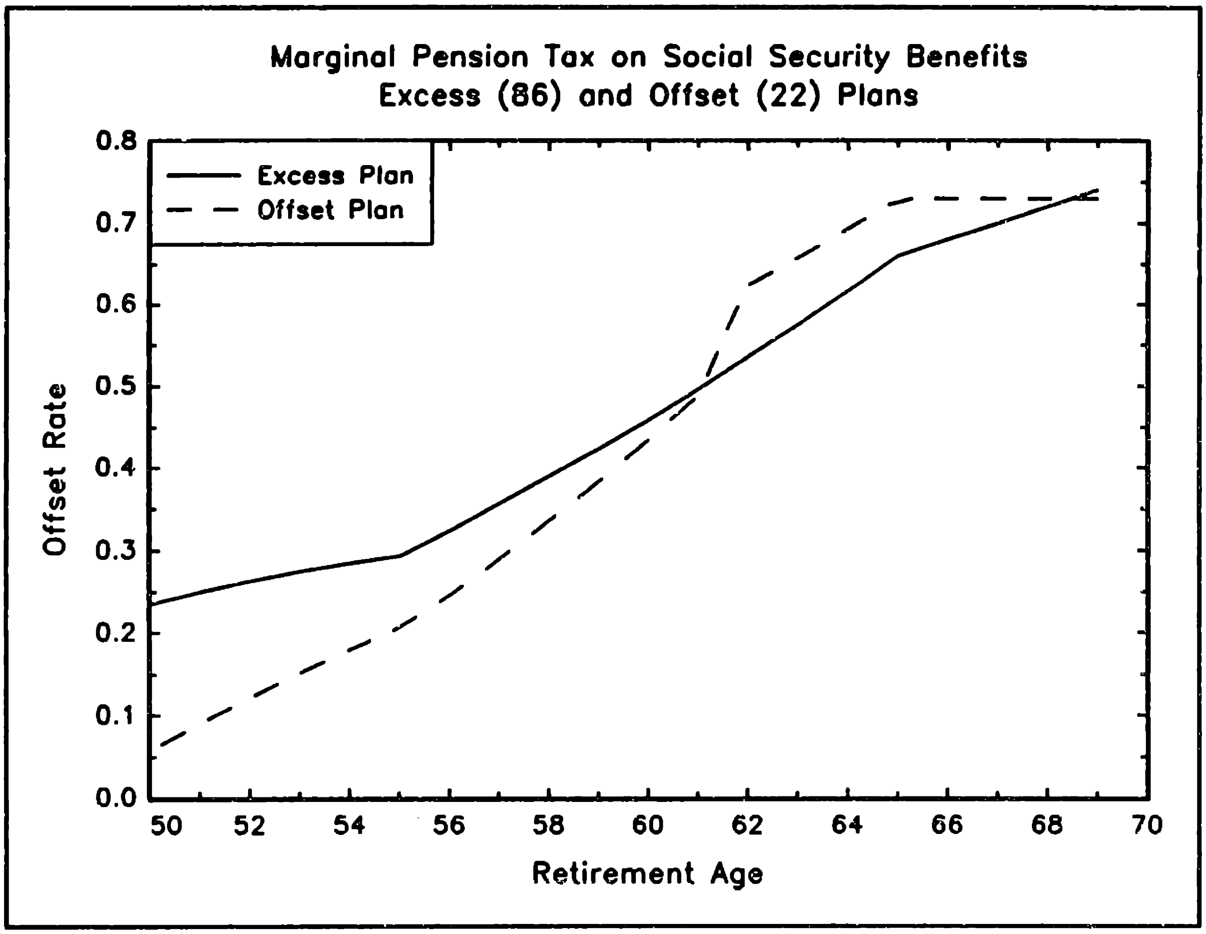

Another way to view the consequences of integrated pension plans is that the pension plans impose a marginal "tax" on the Social Security benefits. Changes in the rules that make Social Security more generous make pensions less generous, so that part of the incremental benefits are taxed away by the pension provider. Figure 9 shows that these marginal tax rates can be quite large, over 75 percent of Social Security benefits for older workers with many years of service. Each schedule graphs the fraction of incremental Social Security benefits that are deducted from the pension benefits when Social Security rules change. Under the Excess plan in Figure 5, a slight increase in the MTE increases Social Security benefits by 15 percent[11] and decreases pension benefits by 0.3% times years of service times the reduction factor (for ER or vested benefits) which itself depends on age. Under the Offset plan in Figures 4 and 8, a slight increase for *any* reason in Social Security benefits available at 65 will be deducted at the rates shown by the dashed line.[12] These rates are increasing functions of age, years of service, and FAP. Thus, both schedules increase with the retirement age and have nonlinearities at the ERA's and NRA's for the plans.

Table 5 shows that a substantial fraction of pension plans are integrated with Social Security.[13] In the full sample of workers in the PPS, roughly 27 percent are

[11]This is technically true only for all remaining years of work which will be counted in the computation of average indexed monthly earnings (AIME). At the MTE, another dollar of AIME increases the primary insurance amount (PIA) on which benefits are based by 15 cents. The schedule in Figure 6 assumes that the MTE is increased for all of the years entering the computation. If the MTE is changed very late in the worker's career, then the change in the MTE will have less of an effect on the Social Security benefits. The "effective" marginal tax rate will be much higher than is graphed here, but it will apply to a smaller incremental change.

[12]A corollary to this point is that generational accounting (see Kotlikoff, 1992) should incorporate pension integration. Changes in the Social Security system designed to promote intergenerational equity will have consequences of changing the value of a firm's pension liability. Part of the burden of the changes will be shifted to the firm's shareholders.

[13]In addition to integration in the form of Excess and Offset plans, many plans have provisions which begin or end benefit payments at the early, normal, and late (62, 65, and 72, respectively, before the 1983 Social Security reforms) retirement ages under Social Security. The purpose is to smooth the flow of total retirement income. When the normal and late retirement ages were changed in the 1983 Social Security amendments, the timing of pension benefits also changed. This type of "integration" is not included in Table 5.

integrated in some way, with the Offset plans covering over one fifth of the workers. Focusing on private sector workers, over 40 percent of the workers in large firms have some integrated pension, with one third of them having an Offset plan.[14] The share for workers in small firms is only 21 percent, but this is likely due to the higher fraction of workers covered only by DC plans, which are almost never integrated in the PPS. Since Table 1 shows that the PPS is less likely to pick up private sector plans than public sector plans and Table 2 demonstrates that the latter are less likely to be integrated, these figures understate the extent of integration in the workforce.

To summarize, the retirement incentives provided by pensions can be quite large, often giving an implied bonus of 100 percent or more of the current wage to employees who work through the Early Retirement Age. They can also be negative near and after the Normal Retirement Age, providing a strong financial incentive to retire. Any analysis of retirement behavior based on a comparison of pre- and post-retirement income must account for these incentives. Moreover, the magnitude and timing of these incentives depend on characteristics of the worker and the pension that may be quite difficult to capture in an *ad hoc* manner. As shown above, strong retirement incentives are more pervasive in DB plans than DC plans. For Flat Rate plans, pension accrual rates are lower at high wages, whereas for Excess plans, pension accrual rates increase with wages. Years of service are perhaps the most important factor affecting the timing

[14] These figures for the share integrated in large firms correspond to those reported in Mitchell (1992), who used data from the Bureau of Labor Statistics' Employee Benefits Survey, for 1982. By 1983, the share of employees in large firms with Excess plans doubled to 20 percent. A leading explanation for this is the increase in OASDI contribution rates enacted in the 1983 legislation. Mitchell also reports that the fraction of pension plan participants covered by any type of integrated DB plan rose from 45 percent in 1980 to 63 percent in 1989.

of large accruals, but even within a given plan, the last important accrual may be a nonmonotonic function of tenure. Finally, the substantial presence of integrated Offset plans makes detailed data necessary to accurately compute both the amount of the pension and the sensitivity of total retirement benefits to Social Security changes.

V. Conclusion

The analysis in Section IV demonstrates that knowledge of an individual's pension formula, rather than a few summary statistics about pension wealth, is essential for understanding the economic opportunity set around retirement. In subsequent work, I use this extra information provided by the PPS to investigate two important economic questions. In Samwick (1993a), I reassess the effect of public and private retirement programs on the timing of retirement in light of the better data on pension accruals found in the PPS. Controlling for the incentive effects of pensions does not affect the overall magnitude of responses to Social Security predicted without detailed pension data, but the pension data is important in accurately predicting the behavior of households younger than the Social Security early retirement age. Pensions are also shown to have statistically significant effects on retirement behavior similar to those of Social Security. In Samwick (1993b), I study the relationship between earnings uncertainty and the risk-compensating elements of pension plan formulas. Essentially, private sector and unionized workers with higher earnings uncertainty are shown to have pensions whose expected present values increase with the amount of wage variability. This finding suggests that, in addition to being vehicles for insurance and saving, pensions are serving to compensate workers for bearing income risk during their working years.

Appendix Using the Pension Provider Survey

The PPS documentation in Curtin (1987) also provides programs that calculate pension entitlements based on the data in the SCF and the PPS. They are sufficiently general so that entitlements can be computed under a variety of economic assumptions for any individual in the SCF or hypothetical. The modifications of these programs for the present analysis were designed to increase their technical and economic applications.

The first modification was purely technical; to change the syntax and file structure so that the programs could be used on any PC. The original programs are written in a version of PASCAL tailored to the University of Michigan MTS mainframe computer. The versions used here have been rewritten in Turbo Pascal, which can compile easily on any PC with expanded memory. The original programs were designed in two stages. The first set read the information on the pension plans from the PPS datafile and created procedures for each plan. The key step was to parse the literal formula codes and generate programming statements from them. The second set of programs read in the information on the plan participants from the SCF and used the code generated by the first set of programs to compute entitlements. It was possible to adapt only this second set of programs to the PC; the pension procedures were generated by the mainframe version and each of the 1011 procedures must be edited individually in order to consider fundamental changes in the structure of pensions. The modifications to Final Average Pay calculations discussed below were made this way.

The second major change was to allow for more general wage growth patterns. The original programs allow for an economy wide growth rate and an individual specific growth rate. The individual's growth rate must remain constant over his entire career, which is an unrealistic assumption. In particular, it does not simultaneously allow for high wage growth during early years on the job and declining (and even negative) wage growth in later years. In order to incorporate variable wage growth over the life-cycle, the coefficients from cross-sectional wage regressions are used. The basic regression is of the form:

$$LN(wage) = \beta_0 + \beta_1 * AGE + \beta_2 * AGE^2 + \beta_3 * EDUC + \beta_4 * EDUC * AGE + \beta_5 * EDUC * AGE^2 + \epsilon$$

where EDUC is the individual's years of education. Separate regressions were run for each combination of {male, female} x {white, nonwhite} x {13 occupational categories}[15] using Current Population Survey (CPS) data on annual wages. Only fulltime workers between the ages of 16 and 64 were used, and estimates were made for each year between 1983 and 1986. The modified programs use sex, race, education, occupation, and age to account for differences in wage growth across individuals.

The third major change concerned the computation of Social Security benefits for use in the Offset and Excess plans. In the original versions, all calculations were based on the average benefit, average covered wage (ACW), and Maximum Taxable Earnings (MTE) as of 1983. Individuals who had current wages above the ACW but below the MTE were assigned the average benefit. Those with less than the ACW received half the average benefit, while those with more than the MTE received 150 percent of the

[15]The occupational classifications are the major 1980 "1-digit" codes used by the CPS.

average benefit. In the new versions, the actual Social Security benefit formulas (and the more realistic wage growth rates) are used so that true variation in the offset rates is revealed in the computed entitlements. Additionally, all changes in the Social Security benefit formulas made in 1983 have been incorporated, so that the effect of those changes on entitlements and retirement can be assessed.

The fourth major change pertained to the calculation of Final Average Pay (FAP) in pensions which had benefits related to earnings. Most such pensions calculate FAP as the average of the highest or highest consecutive years of earnings during some specified period. For example, the highest five years of earnings during the last ten is quite common. When nominal earnings are always increasing with age, which was the case when wage growth was uniform over the lifetime, this formula reduced to the last five years of earnings. The original programs calculated FAP in this simplified way. However, when a more general specification of the wage equation is used, it is possible for nominal earnings at older ages to decline. The simplified FAP calculation would understate FAP (and therefore pension benefits) in such cases. The modified programs compute FAP in the precise way specified in the summary plan description so that this does not occur.

Finally, some portions of the code had to be rewritten so that entitlements of workers as of dates earlier than the survey date could be computed. Even with these changes, however, the programs are very much like the original version in their overall structure, and all substantive modifications could easily be undone. Copies of these modified programs are available from the author upon request.

References

Avery, Robert, Gregory Elliehausen, and Thomas Gustafson (1985). "Pensions and Social Security in Household Portfolios: Evidence form the 1983 Survey of Consumer Finances," *Research Papers in Banking and Financial Economics*. Washington, DC: Board of Governors of the Federal Reserve System.

Bodie, Zvi (1990). "Pensions as Retirement Income Insurance," *Journal of Economic Literature*, Vol. 27, 28-49.

Curtin, Richard T. (1987). *Survey of Consumer Finances Employer Sponsored Pension Benefit Plans: Documentation and Codebooks*. Survey Research Center, University of Michigan.

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

Gustman, Alan L. and Thomas L. Steinmeier (1989a). "An Analysis of Pension Benefit Formulas, Pension Wealth, and Incentives form Pensions," in R. Ehrenberg (ed.) *Research in Labor Economics* 10. Greenwich, CT: JAI Press.

Gustman, Alan L. and Thomas L. Steinmeier (1989b). "Changing the Social Security Rules for Workers Over 65: Proposed Policies and Their Effects," National Bureau of Economic Research Working Paper No. 3087.

Kotlikoff, Laurence J. (1992). *Generational Accounting: Knowing Who Pays, and When, for What We Spend*. New York: Free Press.

Kotlikoff, Laurence J. and David A. Wise (1985). "Labor Compensation and the Structure of Private Pension Plans: Evidence for Contractual Versus Spot Labor Markets," in David A. Wise (ed.) *Pensions, Labor, and Individual Choice*. Chicago: University of Chicago Press.

Kotlikoff, Laurence J. and David A. Wise (1987). "The Incentive Effects of Private Pension Plans," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.) *Issues in Pension Economics*. Chicago: University of Chicago Press.

Kotlikoff, Laurence J. and David A. Wise (1989). "Employee Retirement and a Firm's Pension Plan," in David A. Wise (ed.) *The Economics of Aging*. Chicago: University of Chicago Press.

McDermed, Ann A., Robert L. Clark, and Steven G. Allen (1989). "Pension Wealth, Age-Wealth Profiles, and the Distribution of Net Worth," in Robert E. Lipsey and Helen Stone Tice (eds.) *The Measurement of Saving, Investment, and Wealth.* Chicago: University of Chicago Press.

Merton, Robert C., Zvi Bodie, Alan J. Marcus (1987). "Pension Plan Integration as Insurance Against Social Security Risk," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.) *Issues in Pension Economics.* Chicago: University of Chicago Press.

Mitchell, Olivia S. (1992). "Trends in Pension Benefit Formulas and Retirement Provisions," in John A. Turner and Daniel J. Beller (eds.) *Trends in Pensions 1992.* Washington, DC: United States Department of Labor.

Montgomery, Edward, and Kathryn Shaw (1992). "Pensions and Wage Premia," National Bureau of Economic Research Working Paper No. 3985.

Montgomery, Edward, Kathryn Shaw, and Mary Ellen Benedict (1992). "Pensions and Wages: An Hedonic Price Theory Approach," *International Economic Review*, Vol. 33, No. 1, 111-128.

Samwick, Andrew A. (1993a). "The Joint Effect of Social Security and Pensions on the Timing of Retirement: Some New Evidence," mimeo, Massachusetts Institute of Technology (also Chapter Four in this thesis).

Samwick, Andrew A. (1993b). "Risk Compensation through Employer-Provided Pensions," mimeo, Massachusetts Institute of Technology (also Chapter Five in this thesis).

Watson, Roberta Casper and James R. Feutz (1988). "Integrating Qualified Plan Benefits with Social Security Remains Complex," *Taxation for Lawyers*, Vol. 17, No. 2, 90-93.

Table 1

Pension Provider Survey Take-up Rates, By Employer Type				
Employer Type	Share of Total Workforce	Share Covered by Pension in SCF[1]	Share Covered by Pension in PPS[2]	PPS Take-up Rate[3]
Private Sector, more than 100 employees	0.405	0.681	0.493	0.725
Private Sector, less than 100 employees	0.240	0.215	0.112	0.523
State or Local Government	0.085	0.814	0.748	0.919
Private School	0.016	0.423	0.390	0.922
Public School or College	0.064	0.789	0.761	0.964
Federal Government	0.035	0.818	0.730	0.893
Military	0.008	0.926	0.108	0.116
Self-Employed, (No Pension Data)	0.147			
Full Sample	1.000	0.490	0.372	0.759

Source: Author's calculations using the SCF 1983 and PPS.

1) Percent of workers who claimed to be covered by a pension in SCF household data.

2) Percent of workers who are in (a) and whose plans are represented in the PPS.

3) Ratio of (2) to (1). 1-(3) gives the attrition rate from the SCF to the PPS.

Table 2

Characteristics of the Covered Population in the SCF and Representative Workers						
	Sample Average	Hired at 51 Avg Wage	Hired at 41 Avg Wage	Hired at 31 Avg Wage	Hired at 31 2x Avg Wage	Hired at 31 0.5x Avg Wage
Female	0.4150	0.0000	0.0000	0.0000	0.0000	0.0000
Date of Birth	1941.54	1941.54	1941.54	1941.54	1941.54	1941.54
Date of Hire at current job	1972.26	1992.54	1982.54	1972.54	1972.54	1972.54
Annual hours worked	2051.56	2080.00	2030.00	2080.00	2080.00	2080.00
Annual Wage	22880.61	22880.61	22880.61	22880.61	45761.21	11440.30
Percent Nonwhite	0.1809	0.0000	0.0000	0.0000	0.0000	0.0000
Years of Education	13.44	13.00	13.00	13.00	13.00	13.00
Percent Married	0.7672	0.0000	0.0000	0.0000	0.0000	0.0000
Percent Contributing	0.2034	1.0000	1.0000	1.0000	1.0000	1.0000
Spouse Date of Birth	1941.95	1941.95	1941.95	1941.95	1941.95	1941.95
Percent Contributed	0.0542	0.0500	0.0500	0.0500	0.0500	0.0500

Source: Author's computations from the SCF and PPS

Table 3

Initial Pension Replacement Rates for Representative Workers										
	Hired at 51 Average Wage		Hired at 41 Average Wage		Hired at 31 Average Wage		Hired at 31 2 x Avg Wage		Hired at 31 0.5 x Avg Wage	
Age	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
50			0.013	0.035	0.038	0.091	0.038	0.087	0.038	0.104
51			0.029	0.044	0.055	0.112	0.055	0.108	0.056	0.124
52			0.023	0.057	0.060	0.122	0.060	0.115	0.060	0.133
53	0.002	0.008	0.025	0.062	0.065	0.131	0.064	0.123	0.065	0.142
54	0.003	0.011	0.028	0.068	0.071	0.142	0.070	0.133	0.071	0.154
55	0.006	0.017	0.066	0.109	0.150	0.178	0.145	0.154	0.158	0.205
56	0.014	0.038	0.082	0.127	0.186	0.213	0.179	0.177	0.197	0.265
57	0.017	0.047	0.094	0.143	0.202	0.235	0.194	0.192	0.214	0.293
58	0.022	0.057	0.107	0.162	0.228	0.258	0.218	0.208	0.243	0.318
59	0.027	0.069	0.120	0.180	0.247	0.283	0.236	0.224	0.262	0.344
60	0.042	0.086	0.181	0.237	0.316	0.309	0.303	0.235	0.334	0.376
61	0.087	0.136	0.229	0.274	0.342	0.339	0.328	0.253	0.361	0.417
62	0.118	0.167	0.255	0.303	0.369	0.366	0.354	0.268	0.389	0.448
63	0.134	0.190	0.274	0.332	0.389	0.398	0.371	0.286	0.409	0.485
64	0.149	0.217	0.293	0.362	0.408	0.431	0.389	0.304	0.428	0.524
65	0.188	0.274	0.323	0.393	0.440	0.465	0.420	0.318	0.461	0.566
66	0.204	0.285	0.340	0.403	0.453	0.485	0.433	0.330	0.472	0.575
67	0.216	0.297	0.352	0.416	0.467	0.509	0.446	0.345	0.485	0.587
68	0.230	0.311	0.368	0.430	0.485	0.535	0.463	0.361	0.502	0.603
69	0.244	0.322	0.384	0.443	0.504	0.562	0.480	0.380	0.519	0.618

Source: Author's calculations using the SCF 1983 and PPS.

1) The replacement rate is the ratio of pension benefits that could be received if the worker retired at each age to the wages that could be earned if the worker did not retire.

2) All replacement rates are weighted to reflect the aggregate population covered by pensions in 1983.

Table 4

Age	Hired at 51 Average Wage		Hired at 41 Average Wage		Hired at 31 Average Wage		Hired at 31 2 x Avg Wage		Hired at 31 0.5 x Avg Wage	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
50			0.059	0.108	0.189	0.446	0.188	0.399	0.196	0.496
51			0.335	0.515	0.279	0.581	0.275	0.576	0.287	0.587
52			0.146	0.366	0.193	0.213	0.187	0.187	0.200	0.225
53	0.005	0.232	0.134	0.136	0.201	0.209	0.196	0.187	0.211	0.230
54	0.019	0.164	0.144	0.143	0.222	0.259	0.216	0.239	0.231	0.275
55	0.041	0.143	0.340	0.684	0.622	1.091	0.578	0.896	0.656	1.102
56	0.177	0.473	0.256	0.496	0.407	0.746	0.389	0.635	0.449	1.329
57	0.072	0.149	0.230	0.295	0.282	0.349	0.272	0.263	0.295	0.436
58	0.089	0.214	0.262	0.445	0.367	0.674	0.347	0.630	0.396	0.722
59	0.101	0.314	0.255	0.348	0.299	0.390	0.289	0.307	0.310	0.440
60	0.165	0.367	0.386	0.817	0.455	0.811	0.439	0.676	0.459	0.751
61	0.570	1.044	0.304	0.621	0.316	0.497	0.306	0.421	0.331	0.685
62	0.250	0.702	0.241	0.555	0.244	0.486	0.227	0.394	0.257	0.565
63	0.181	0.325	0.201	0.341	0.200	0.383	0.185	0.271	0.203	0.451
64	0.172	0.336	0.185	0.317	0.178	0.356	0.165	0.219	0.180	0.429
65	0.232	0.608	0.254	0.610	0.267	0.761	0.258	0.719	0.267	0.789
66	0.123	0.346	0.134	0.438	0.091	0.245	0.089	0.181	0.074	0.256
67	0.096	0.191	0.079	0.222	0.077	0.271	0.076	0.201	0.066	0.266
68	0.097	0.150	0.100	0.185	0.102	0.248	0.090	0.172	0.088	0.250
69	0.089	0.139	0.092	0.197	0.096	0.234	0.086	0.181	0.079	0.278

Pension Accruals Relative to Wages for Representative Workers

Source: Author's calculations using the SCF 1983 and PPS.

1) The pension accrual rate is the ratio of the increase in the actuarial present value of pension benefits that results from postponing retirement for one year to the current wage at each age.

2) All accrual rates are weighted to reflect the aggregate population covered by pensions in 1983.

Table 5

Plan Type and Social Security Integration Status, By Employment Characteristics					
	Type of Pension Plans[a]		Type of Social Security Integration[b]		
Employer	Has DC Plans only	Has DB Plans	Has Integrated Plans	Has Offset Plans	Has Excess Plans
Private Sector, more than 100 employees	0.096	0.904	0.416	0.323	0.117
Private Sector, less than 100 employees	0.351	0.649	0.213	0.158	0.056
State or Local Government	0.039	0.961	0.122	0.066	0.065
Private School	0.715	0.285	0.117	0.117	0.000
Public School or College	0.032	0.968	0.064	0.050	0.026
Federal Government or Military	0.000	1.000	0.000	0.000	0.000
Age 50-70 in 1983	0.065	0.935	0.247	0.179	0.080
Covered by Social Security	0.107	0.893	0.312	0.237	0.094
Full Sample	0.095	0.905	0.269	0.204	0.081

Source: Author's calculations using the SCF 1983 and PPS.

1) Percent of covered workers having only Defined Contribution (DC) plans and percent having at least one Defined Benefit (DB) plan (they may also have one or more DC plans).

2) Percent of covered workers having some form of social security integration; percent having an offset plan; and percent having an excess plan. See text for discussion.

Chapter Four

The Joint Effect of Social Security and Pensions on the Timing of Retirement: Some New Evidence

I. Introduction

One of the most persistent trends in the American economy since the end of World War II is the decline in the labor force participation rates of older workers. After remaining essentially constant at 65 percent from 1870 to 1937 (Ransom and Sutch, 1988), the labor force participation rates of men over 60 have fallen continuously. Male labor force participation rates decreased from 62 percent in 1940 to 40 percent in 1985, with much of the decline coming in the 1970's. The rates for older females actually increased over this period from 15 percent to 21 percent, but this is likely due to the large increase in labor force participation rates for women of *all* ages in the postwar period. Nonetheless, female labor force participation rates have declined since the early 1970's from a peak of about 23 percent.[1] Since most persons do work for at least some portion of their adult lives, an important determinant of lower labor force participation is earlier ages of retirement.

A leading explanation for this dramatic decline in labor force participation rates has been the contemporaneous growth of the Social Security program. A legacy of the New Deal, the retirement portion of Social Security is now a 300 billion dollar a year entitlement program that covers just about every worker and retiree in the country. Total contributions to the Old-Age and Survivors Insurance trust fund in 1990 were $267 billion, compared to $14 billion in 1950. In 1989, there were 24 million retired workers who collected average benefits of $6,800. The corresponding figures for 1950 were 4.5

[1]These figures are taken from Lumsdaine and Wise (1990), who document the trend in more detail.

million and $1,050.[2] Chen (1992) reports that the share of aggregate income of households age 65 and over received from Social Security increased from 22 percent in 1958 to 38 percent in 1988.

For over a decade economists have tested the validity of higher Social Security benefits as an explanation for earlier retirement using a variety of estimation methods. Surveys of the literature can be found in Atkinson (1987) and Quinn, Burkhauser, and Myers (1990). Specific examples to be discussed below include Diamond and Hausman (1984), Burtless and Moffitt (1984), Hausman and Wise (1985), Burtless (1986), and Sueyoshi (1989). The general findings of such studies are that although it is possible to estimate statistically significant relationships between the level of Social Security benefits and the likelihood of retirement at various ages, these estimated relationships typically imply a very small economic impact of altering Social Security benefits on the average age of retirement. This suggests that the trend toward lower labor force participation has other explanations.

The standard dataset used in past work on the effect of Social Security on retirement has been the Retirement History Survey (RHS). The RHS is a ten year, biannual panel of several thousand households aged 58 to 63 in 1969. It contains detailed information on labor supply, earnings, Social Security coverage, and health which are necessary to model retirement behavior. Its information on employer-provided pensions, however, is limited to whether or not an individual was eligible for a pension upon retirement; there is not enough information to construct a reliable measure of expected

---

[2] All figures are from the Social Security Administration (1990). All dollar amounts are in constant 1990 dollars.

pension benefits, much less the differences in pension entitlements for alternate years of retirement.

Pension coverage and entitlements also grew rapidly since the early 1940s, however, and could just as readily be the cause of the decline in labor force participation. Over half of the civilian workforce in 1983 was covered by an employer-sponsored pension plan, and this fraction increases past 70 percent when only those workers 25 to 64, working 1000 or more hours per year, with at least one year of service are included (Andrews, 1985).[3] Between 1940 and 1973, the number of retirees receiving pension benefits increased from 200,000 to over 6 million and the number of covered employees rose from 4.1 million to 35 million (Steinberg and Dankner, 1987). Owing to sectoral shifts in employment and sagging economic growth, pension coverage has levelled off at around 50 million workers in very recent years (Bloom and Freeman, 1992). Finally, Chen (1992) reports that the share of aggregate income of households age 65 and over received from pensions (private and government) increased from 14 percent in 1958 to 17 percent in 1988. Moreover, pension plans are often designed to provide strong incentives for retirement at particular ages, as demonstrated in Kotlikoff and Wise (1987, 1989).

Recent work by Stock and Wise (1990a,1990b) has shown that taking account of the incentive effects of pensions on retirement substantially improves the ability of a model to predict the retirement behavior of workers with pensions. In order to analyze

[3]This group of workers is sometimes referred to as the "ERISA workforce," because this is the group that must be covered (in 1983) by a plan in order for it to qualify for tax-favored status under the Employee Retirement Income Security Act of 1974.

pension plans in detail, they select their sample from the personnel records of a single large firm. Their sample is not necessarily representative of the working population nearing retirement age, and their results therefore may not generalize to the population at large. This hampers their ability to assign a fraction of the decline in labor force participation to the growth of pensions. Moreover, using personnel records, they cannot explicitly control for the effects of other information like health status, household composition, and wealth; these factors were shown to be important determinants of retirement probabilities in the prior literature on Social Security and retirement.

This paper is one of the first to use a unique dataset that circumvents the existing tradeoff of detailed pension data versus a nationally representative household survey. The approach here is to link the demographic, employment, and wealth data on households in the *Surveys of Consumer Finances* (SCF) 1983 and 1986 with information on their pension plans in the companion *Pension Provider Survey* (PPS). Briefly, the PPS attempted to interview the plan provider for every worker in the SCF 1983 who claimed to be covered by a pension. The level of detail was such that pension entitlements can be calculated for each worker for any potential year of retirement. Thus, the SCF and PPS together comprise a panel dataset which contains information on the pension entitlements, retirement behavior, and economic characteristics of a representative sample of workers.

The principal findings of the paper are threefold. First, using the PPS to explicitly account for the incentive effects of pensions, the results of previous studies on Social Security and retirement which omitted pensions are upheld. The coefficients on the present value and accrual Social Security benefits are statistically significant in the reduced form models, but simulations of changes in benefit rules reveal modest effects on aggregate retirement behavior. Second, pensions are also found to have statistically significant effects on the probability of retirement, thereby generalizing the results of Stock and Wise (1990a,1990b) to a more representative sample of households. Simulations of extending pension coverage to the one third of the sample without pensions would yield changes in aggregate retirement rates of approximately the same magnitude as those derived from typical Social Security rule changes. Third, the effects of pensions are more heavily concentrated at ages before individuals can begin *collecting* benefits from Social Security. This finding suggests that pensions may affect retirement in part through relaxing liquidity constraints.

The remainder of the paper is organized as follows. In the next section, the data in the SCF and PPS are described. Section III reviews the past literature's treatment of Social Security and pensions in models of retirement. The basic estimation results are presented and discussed in Section IV. In Section V, the effects of changes in Social Security and pension provisions are simulated. Section VI suggests directions for further research and concludes.

II. Data

The Survey of Consumer Finances 1983 (SCF) was designed to be the most comprehensive survey of household wealth ever conducted. In recognition of the growing financial importance of pensions in the economy,[4] the Pension Provider Survey (PPS) was conducted with the households' employers to obtain the details of the pension plans. Because the SCF also contains detailed demographic and economic data and a re-interview survey was performed in 1986, it is possible to use the two datasets to study retirement behavior in the years between 1983 and 1986. This section describes the sample extract used and the construction of the relevant variables for the subsequent analysis.

The full SCF is comprised of 4,303 households in 1983. Of these, 438 are from a sampling frame selected from high income tax returns and 3,824 are from an area-probability sample designed to be nationally representative of the population. In 1986, 2,822 households were re-interviewed, providing a panel of observations whose interim labor market experience can be analyzed.[5] The sample used below includes all respondents and spouses in such households between the ages of 50 and 69 and working full-time (at least 20 hours per week) in 1983. Individuals were excluded if they reported being self-employed or working in the military services as their main occupation. Individuals were also excluded if they reported being covered by a pension but were not present in the PPS.[6]

---

[4]According to the Flow of Funds Accounts (Board of Governors, 1990), pension fund reserves accounted for $1.532 trillion or 13 percent of the $11.640 trillion of household sector net worth in 1983. The corresponding percentage for 1945 was less than 2 percent.

[5]Avery, Elliehausen, and Canner (1984a, 1984b) describe the sample construction of the 1983 survey. Avery, Elliehausen, and Kennickell (1987) discuss the re-interview process and the general quality of the data in the 1986 survey.

[6]Samwick (1993) discusses this censoring of observations in the PPS. Response rates for public sector and educational institutions were all around 90 percent or more. Response rates for firms with 100 or more employees were over 70 percent, while response rates for private firms with fewer than 100

After imposing these restrictions, the sample contains 525 individuals. Table 1 presents means and standard deviations of demographic and economic variables that can be compared with those of a sample from the RHS used by Sueyoshi (1989). The SCF numbers are weighted to maintain comparability in the presence of the oversampling of high income households. The SCF sample contains individuals employed in the public sector and working women; neither of these groups is generally included in the RHS extracts used in past studies. On average, the SCF sample has about two more years of education and $1,500 higher wages. The SCF sample contains many more white collar and service workers than the RHS. Part of this difference may be the result of the high-income sample (even when weighted), but it may also be due to sectoral shifts in employment during the 1970s.

Two variables in particular do not seem comparable across datasets. The health variable in the SCF is self-reported, and only about 1 percent of the households claim to be in poor (as opposed to fair, good, or excellent) health. As such, it is unlikely to pick up important differences in health status across individuals as a determinant of retirement. Moreover, the SCF does not report information on health insurance coverage that would help control for the importance health status on the decision to retire. Also, more than twice as many persons are covered by a pension in the SCF sample than in the RHS, but this discrepancy is a result of the scant attention paid to pensions in the RHS, as the SCF figure agrees with other published sources (see Fields and Mitchell, 1985, for example).

---

employees were slightly over 50 percent. The effect of this sample restriction is therefore to over-represent public sector workers relative to private sector workers.

At the heart of the subsequent analysis of retirement is measuring the incentive effects of Social Security and pensions. For Social Security, the SCF reports employment information on the current job and up to three previous ones for each individual. The re-interview survey in 1986 reports information on the job held in 1986 and the circumstances, if any, under which the 1983 job was terminated For each job, dates of employment, industry, occupation, and final wages are given. Using the Current Population Surveys for 1983-1986, I estimated cross-sectional wage equations as a function of age, sex, race, education, and occupation and used the coefficients to impute wage histories on all jobs.[7] Given the wage histories and the published benefit laws for Social Security, the benefits can be computed for each individual for each year in which he or she could retire. For pensions, the PPS specifies the exact pension formula for each covered individual. Using the information for the current job and the imputation method described above, it is possible to compute the pension benefits for each individual for any year in which he or she could retire.

Table 2 presents summary statistics on Social Security and pension wealth levels. Social Security wealth is the actuarial present value of all Social Security benefits the individual would receive if he or she retired in the current year. Pension wealth is defined analogously.[8] The top panel of the table shows that the average worker has about $47,000 in Social Security wealth and close to $30,000 in pension wealth. The

[7]The appendix to Samwick (1993) describes this estimation and imputation procedure in more detail.

[8]The estimates of Social Security and pension wealth assume an inflation rate of 4 percent and an economy-wide real wage growth rate of 1.5 percent, corresponding to the Social Security Administration's II-B assumptions. For the purpose of discounting, the standard life tables by age and gender from Faber (1982) and a real interest rate of 2 percent are used.

median Social Security wealth is just slightly lower than the mean, but the median pension wealth is zero. This reflects the third of the population not covered by pensions plus some individuals covered by pensions but not yet vested in 1983.[9]

To get a better idea of the distribution of benefits, the bottom panel of Table 2 excludes all observations who are not eligible for both pensions and Social Security. The figures for Social Security are slightly higher, and the mean for pension wealth is now about $40,000 or 75 percent of Social Security wealth. Median pension wealth is only about $18,000; nonetheless, the medians for both types of retirement wealth are considerably higher than that of financial assets, suggesting that most workers will rely primarily on their pensions and Social Security to finance retirement consumption. Additionally, the standard deviation for pension wealth is $65,000 or nearly three times that of Social Security wealth. These statistics show that pension wealth is far more concentrated at the top of its distribution than is Social Security wealth.[10] The right side of Table 2 scales retirement wealths by the wage, and this tends to smooth out the differences in the two distributions.

Another important feature of Social Security and pensions as they affect retirement is their accrual rate each year. The pension accrual rate is defined as the difference in the actuarial present value (APV) of pension benefits conditional on an

[9]An employee is vested in a pension plan when he or she could leave the firm and be entitled to at least some benefits at a later date.

[10]One reason why the average Social Security wealth is higher despite universal pension coverage and some large outliers for pensions in this subsample is the fact Social Security benefits are indexed to inflation whereas fewer than 3 percent of defined benefit pension plans contractually adjust benefits for inflation. Allen, Clark, and McDermed (1992) report that about one quarter of defined benefit pension plans in medium and large firms provided some *ad hoc* increase in benefits between 1983 and 1987. The probability of receiving such an increase is not incorporated in the calculations of pension wealth.

employee's employment history as of age a and the same quantity as of age a-1, suitably increased by the interest rate. Denote this difference by I(a) and its ratio to the current wage W(a) by R(a):

$$I(a) = APV(a) - (1+r)APV(a-1)$$

$$R(a) = \frac{I(a)}{W(a)}$$

I(a) is analogous to a wage in the sense that the worker's lifetime wealth is augmented by this amount by working at age a; it is a form of compensation for being employed at the job in that year. Since all pensions are (weakly) increasing in the level of wages in each year, workers with higher values of W(a) will tend to have higher values of I(a) regardless of the details of the pension formula. Scaling by W(a) focuses on the differences in pension accruals inherent in the pension formulas rather than differences in wages. R(a) thereby allows comparisons to be made between the retirement incentives provided by pensions for workers of different wage levels.[11]

Figure 1 shows a schedule for R(a) that is typical of defined benefit pension plans in the PPS. The two large spikes in the graph occur at the ages when the worker is first eligible for early and normal retirement benefits. Before the early retirement age (ERA) of 55, the worker is entitled only to vested benefits; they cannot be taken immediately and are subject to actuarially fair reduction for every year they are received before the normal retirement age (NRA) of 65. Provided the worker remains with the firm until

[11]Samwick (1993) describes ways in which the R(a) schedule will be affected by the wage level itself. These pension features are economically important when considering the same individual with varying wage levels but not when comparing workers with vastly different wage levels.

the ERA, benefits are available immediately at an actuarially favorable reduction rate for each year they are taken before the NRA. The change in reduction factors determines the size of the spike at the ERA; in this case, the increase in pension wealth is equal to twice the annual wage. Like many plans, this one has a provision whereby normal retirement benefits can be taken before the NRA without any reduction if the worker has sufficient years of service. This accounts for the smaller spike at age 60, which is equal to about 60 percent of the wage that year.

Figure 1

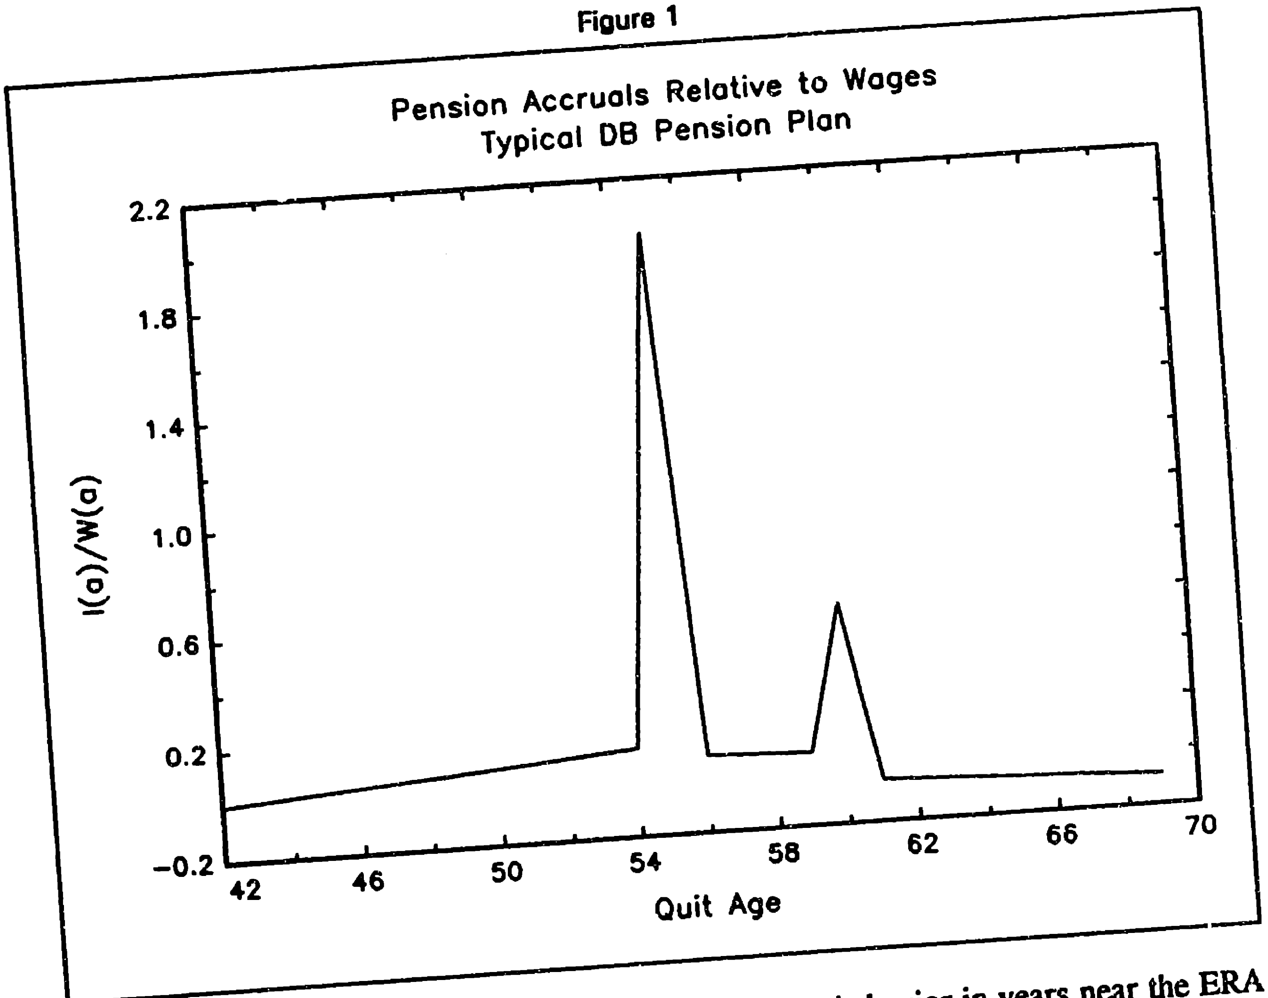

The salient feature of Figure 1 is that retirement behavior in years near the ERA and NRA should be heavily influenced by the opportunity to get large increments to lifetime wealth by working until those ages. It would be unusual, for example, to observe this worker retiring at age 54 given the very high financial rewards to working just one more year. The bottom panel of Table 2 shows that pension accruals are much more dispersed than pension wealth; the coefficient of variation of the former is 1.68 while that of the latter is 6.32. The diversity of such incentives in the PPS is explored in greater detail in Samwick (1993).

The ability to account for the variation in accrual patterns workers actually face is the chief contribution of the PPS to the study of retirement behavior. Table 2 shows that there is far more variation in pension accruals than in Social Security accruals. Failing to account for pension accrual profiles will omit large variation in the incentives for retirement facing workers with pensions.

III. Past Results on Social Security, Pensions, and Retirement

In order to have a benchmark for comparison with later results, this section briefly reviews the conclusions of the existing literature on the impact of Social Security and pensions on retirement. The general conclusions regarding Social Security have been that although a statistically significant effect on retirement behavior can be estimated, the implied result of changing the benefit rules is economically small. More recent work on the effect of pensions on retirement behavior has demonstrated that retirement rates can be extremely sensitive to pension plan features, but the ability of these results to explain general trends in labor force participation is questionable because the samples used were not representative of the entire population.

As discussed in Section I, the most popular dataset in past work on Social Security and retirement is the RHS. Because the RHS contains no information on the magnitude of pension entitlements, all of these studies necessarily concentrated on Social Security in isolation. At most, an indicator variable for pension eligibility was included or individuals known to have pensions were excluded from the sample. Two principal types of model have been estimated using the RHS.

The first type of model utilizes the nonlinear budget set method developed in Burtless and Hausman (1978) and Hausman (1981) to examine the effect of taxes on labor supply. A careful application of this procedure to retirement is found in Burtless (1986). The idea underlying this model is to use the full detail of the Social Security rules to map out the complete, nonlinear opportunity set facing each worker for any potential retirement age. A utility function is then specified in terms of variables like wealth, marital status, health, and income; and unknown parameters to yield a predicted retirement age. Maximum likelihood is used to estimate the parameters that generate the closest match between predicted and actual retirement ages in the sample. Using these estimated parameters, the effects of the unexpected increase in the level of Social Security benefits in the early 1970s on subsequent retirement behavior are examined. Simulations reveal that the long run impact of this policy change was to reduce the average retirement age by about 2 months and to increase the likelihood of retirement at ages 62 and 65 by about 2 percentage points.

The other type of model typically used to estimate retirement probabilities is the hazard model. The hazard model is a reduced-form specification in which economic variables multiply the underlying (or baseline) instantaneous probability of retirement at each age.[12] Hausman and Wise (1985) estimate a hazard model using the RHS that explains retirement behavior in terms of the level and accrual rates of Social Security, health status, earnings, liquid assets, pension eligibility, and demographic variables. Their parameter estimates imply that if Social Security benefit levels had been maintained at the 1969 level, the ensuing reduction in labor force participation rates would have been only two thirds as great.

These findings of statistically significant but economically small effects of Social Security on retirement are the rule among papers written using this data. Also using the nonlinear budget set method, Burtless and Moffitt (1984) found statistically significant effects of Social Security on retirement ages and hours of post-retirement labor supply, but they also calculated that a 20 percent reduction in retirement benefits would lead to an average increase in the retirement age of about 2-2.5 months. Sueyoshi (1989) estimates a competing risks model (essentially a bivariate hazard model) of full and partial retirement. He finds that the large increase in benefits during the 1970s had very little impact on retirement rates, increasing the probability of full early retirement at the expense of partial retirement rather than continued fulltime work. An exception to the findings of small effects is Boskin and Hurd (1978), who estimate a probit model and find that a $1,000 increase in Social Security benefits raises the probability of retirement over the period 1969-1971 by 8 percentage points.

[12]The first application of this model to an economic problem was Lancaster's (1979) study of unemployment duration, and good explanations of the unemployment model can be found in Heckman and Singer (1984), Meyer (1990), and Han and Hausman (1990).

Diamond and Hausman (1984) estimate a hazard model on a different dataset in which they can impute expected pension entitlements for those currently working based on the pension income being received by those who are retired in a year toward the end of the survey. The data is drawn from the National Longitudinal Survey (NLS) of Older Men. The coefficients on Social Security benefits are highly significant and those on pensions are jointly significant. Nonetheless, they still find a large secular trend toward earlier retirement not attributable to economic variables. Their simulations show that delaying the normal retirement age for Social Security by three years would increase the average working life by 0.6 years. Thus, even with a different dataset and imputed pension levels, Social Security is not shown to be the primary factor in explaining the trend toward earlier retirement.

In response to the inability of these efforts to explain the trend toward earlier retirement, Stock and Wise (1990a,1990b) focused their attention on the incentive effects of pension plans on retirement. As discussed in Section II, pensions differ from Social Security in that they often provide large financial incentives to remain with the firm until a particular age and to leave the firm shortly thereafter. Building on the model of pensions in Lazear and Moore (1988), Stock and Wise (1990a) develop an option value model in which the key determinant of whether a worker retires in a given year is the gain in lifetime utility that would result if retirement were delayed until it is predicted to be most advantageous.[13] The model takes full account of the incentives for

[13]This decision-making rule is similar to the dynamic programming rule of Rust (1989), in which the key determinant is whether making the same retirement decision next year yields more utility than retiring now and forsaking that later decision. Lumsdaine, Stock, and Wise (1992) explore the differences between the option value model and related dynamic programming models.

retirement provided by pensions and Social Security. The probability that a worker retires in a given year is the probability that the option value of continued work is negative. The option value itself is a function of structural parameters of an indirect utility function; the estimated parameters are those which maximize the likelihood function generated by the observed retirement behavior in a sample of workers.

Stock and Wise (1990a) estimate the parameters of the option value model on a sample of salesmen for a large *Fortune* 500 firm.[14] Using these estimates, they simulate the effects of two changes in the plan on retirement rates. First, they show that increasing the early retirement age--the age at which benefits can be drawn immediately upon leaving the firm--from 55 to 60 can reduce the fraction of workers with the firm at age 50 who are not employed at age 60 from 65 percent to 42 percent. Second, they show that switching the plan to an actuarially equivalent defined contribution plan can eliminate some, but not all, of the discontinuities in the rate of retirement by age. Additional results in Stock and Wise (1990b) show that changes in the Social Security early retirement age or the reduction factor applied to benefits taken early have negligible effects on retirement rates.

IV. Estimates of Retirement Probabilities

The focus of the empirical work using the SCF and the PPS is on determining which of the conclusions reached in the past literature are robust to estimating a retirement model on the richer dataset. As in all previous work, the features of the model that can be estimated here must be tailored to the limitations of the dataset. The

---

[14]The precise details of this firm's plan can be found in Kotlikoff and Wise (1989).

most important difference between the SCF and the RHS is the number of observations. Table 1 reported that there were roughly one third as many people in the at-risk set for retirement in the SCF (525) as in a typical extract of the RHS (1633). This disadvantage is compounded by the comparatively short panel of three years in the SCF relative to the ten year panel in the RHS. There are two principal consequences of the smaller dataset.

First, an important issue in estimating a retirement model is how to identify the actual transition from the workforce to retirement. An obvious solution would be to use the respondent's self-reported job status, but the post-retirement behavior of people according to this definition can vary considerably. Some people may claim to have retired simply by reducing the amount of hours they work at the same firm.[15] Others may leave a job that they have held for several years and take part-time work elsewhere to supplement their retirement income. In short panels such as the SCF where the respondents are interviewed only twice, there is little choice but to base such classifications on self-reported retirement status. For each individual, the re-interview survey in 1986 solicits information on whether the 1983 job was terminated and if the reason was retirement. Only those who report that they retired or quit are counted as retired.

Second, the dimensions of the panel dictate what econometric specifications are appropriate. All retirement models are estimated essentially by comparing the observed retirement behavior in a sample of workers with the value of a function of their economic

---

[15]Using the RHS, both Rust (1989) and Sueyoshi (1989) have considered the permanence of the first transition out of the labor force from a job that has been held for a long period of time.

characteristics. The simplest examples hold that a worker will retire if and only if a latent variable, say, y^* is greater than zero. In the case of a probit,

$$y^* = X\beta + \epsilon, \qquad \epsilon \sim i.i.d.\ N(0,\sigma^2)$$

where X is a vector of demographic and β is a vector of parameters to be estimated. The probability of retirement is therefore:

$$Prob[\epsilon > -X\beta] = \Phi(\frac{X\beta}{\sigma})$$

The value of β that is optimal is the one that maximizes the log-likelihood function:

$$L(X,\beta) = \sum_{i=1}^{N}\sum_{t=1}^{T_i}(1-\delta_{it})\cdot Log(1-\Phi(X_{it}\beta)) + \delta_{it}\cdot Log(\Phi(X_{it}\beta))$$

where i indexes the individual, t indexes the year of observation, δ_{it} is an indicator variable for whether or not retirement occurs that year. Note that σ has been normalized to unity but could in principle be modeled to incorporate individual effects or serial dependence.

A more complicated model would be a hazard model, which takes advantage of the similarity between labor force attachment and duration analysis to estimate the effects of economic covariates separately from an underlying baseline retirement probability at each age. If the economic variables are changing over time within a given spell, a hazard model is appropriate when there are many observations on each individual. Accurate estimation of a nonparametric baseline hazard also requires many observations of retirement behavior at each age. Given the wide range of ages in the sample and the short panel, a more reasonable approach is to estimate probits and augment the basic specification with various age variables.

Another reason to use the simpler probit specification is that it maintains continuity with the option value models that have been used previously in the pensions literature. The option value differs from the simple probit primarily in that it passes a more general argument to the distribution function $\Phi(\cdot)$. Instead of retiring when the value of $X\beta$ is low, the option value model predicts that people will retire when a function $G(X,\beta,\sigma)$ is low. The innovation of Stock and Wise (1990a) is to formulate $G(X,\beta,\sigma)$ in a structural model to correspond to the option value of continuing work; that is, the loss in expected lifetime utility if one were to retire today instead of at an optimal retirement date in the future.

Table 3 presents the basic retirement probit estimates. The variables of particular interest are those pertaining to Social Security and pension benefits. The present values of these variables are the actuarial present values of the benefits that would be received if the worker stopped earning income covered by Social Security or pensions in the current year. Since this quantity corresponds to a claim on a future stream of payments and leisure during retirement is taken to be a normal good, we expect that higher present values will increase the probability of retirement, other things equal. The accrual rates on Social Security and pensions correspond to the incentives described in Section II, i.e. the value of $R(a)$ including both pensions and Social Security; they reflect the financial gain in retirement wealth that results from working with the same employer during the current year. A high accrual rate indicates that a large increase in lifetime resources will be obtained if retirement is postponed. Thus, we would expect high accrual rates to decrease the probability of retirement.[16]

Models (1) and (2) in Table 3 confirm these intuitions. Both Social Security and pension wealth variables are positive and statistically significant at any conventional level, and both accrual variables are similarly significant and negative. These results are similar to the reduced form estimates in Lumsdaine, Stock, and Wise (1992). Model (3) shows that when all four retirement variables are included simultaneously, the magnitude and significance of the Social Security accrual and pension wealth variables are hardly affected. Social Security wealth and pension accrual retain the correct signs but lose most of their significance. Model (4) restricts the coefficients on the wealth and accrual variables to be the same across Social Security and pensions. A likelihood ratio test cannot reject this restriction at the 5 percent level. In the restricted model, the coefficients on retirement wealth and accrual are between those on the corresponding pension and Social Security variables in the unrestricted model.

The demographic variables in the specification are included more because they may be correlated with the discount rate or the marginal utility of leisure than because their coefficients are inherently interesting. The only such variable that is generally significant is household size; persons with more dependents to support are less likely to retire. Although the coefficient on health is generally positive as would be expected, it

[16]The Stock and Wise (1990a) model generalizes this approach by considering spikes in the accrual rate of pension wealth that may occur in any future year on the job, not just the next one. Using the maximum possible Social Security and pension wealths (i.e. the largest values they could ever attain if the worker stayed with the job indefinitely) instead of the current wealth levels yields quantitatively similar results to those below. A comprehensive treatment of future accruals as in the option value model, however, has thus far been intractable.

is not statistically significant. This result contrasts with past studies which find health to be a prime determinant of retirement behavior, but as discussed in Section II, the measure of health in the SCF is particularly inadequate.

In order to understand what contribution the PPS makes to these estimates, Model (5) emulates the type of specification that could be estimated if only pension eligibility were known. The pension variable is simply an indicator variable for whether the individual could retire and collect a pension. The coefficients on the Social Security variables are virtually identical to those in Model (3), with pension eligibility weakly implying a higher probability of retirement. From this comparison it can be concluded that the magnitudes of the estimates in past studies which lacked detailed pension data were not seriously affected by the omission. They simply failed to capture the independent effect of pensions on retirement behavior.

Figure 2 graphs the actual retirement hazard by age in the sample along with the predicted hazard rates from Models (4) and (5). The large increases at age 62 and 65 are typical of retirement datasets, and their correspondence to the early and normal retirement ages for Social Security benefit receipt seem to suggest a large effect of Social Security on retirement behavior. The predicted hazards, however, do not bear this out, as both fail to capture these large increases in the probability of retirement. Essentially, the present value and accrual of retirement income at ages 61 and 63 are not so different from those at age 62 to justify so large a spike in the model.

Table 4 presents estimates of Models (4) and (5) from Table 3 under alternative specifications. The first panel uses the levels of retirement wealth and accruals, not their

Figure 2

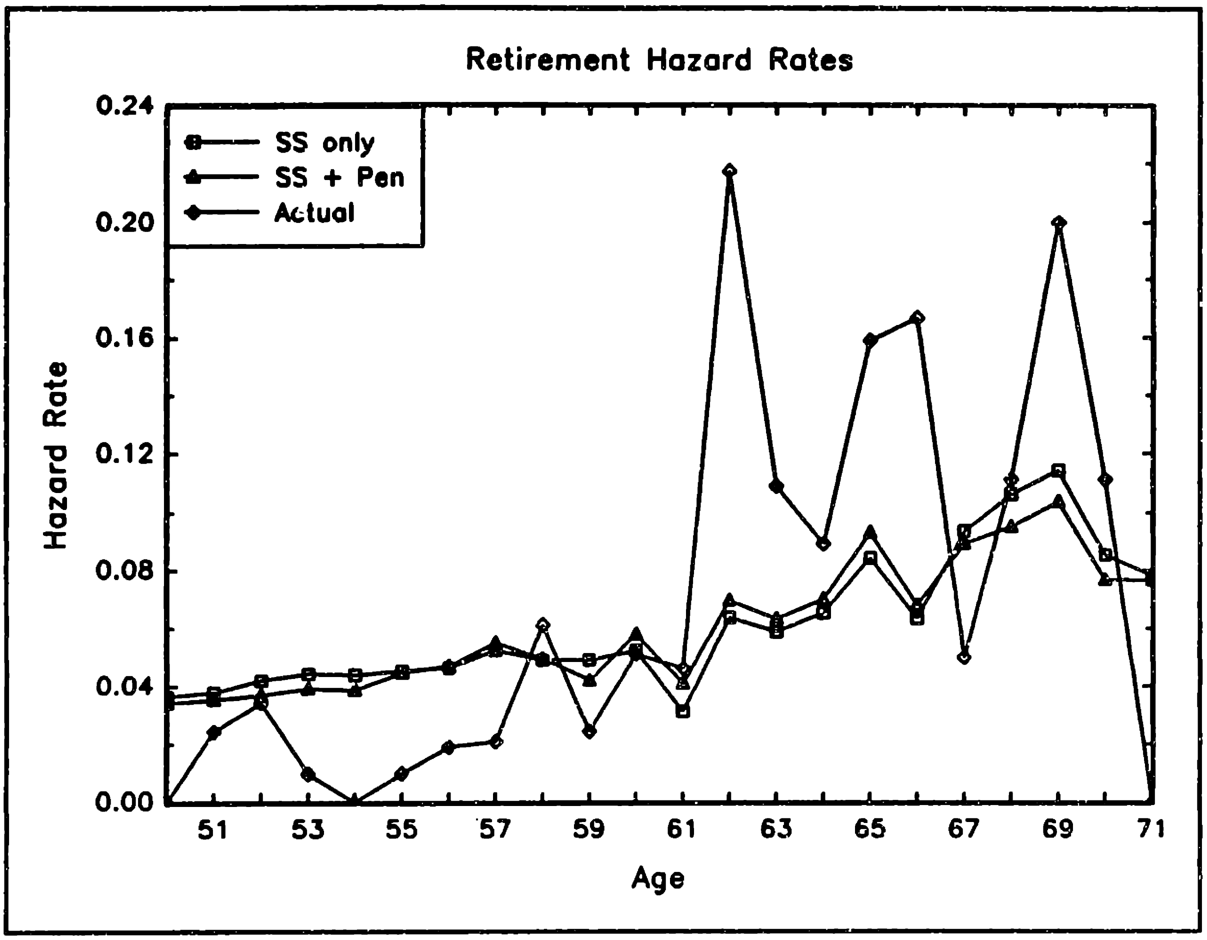

ratios to wages. Because the probit is a reduced-form model, it is not clear *a priori* whether the opportunity cost of retirement is better measured with the wage entering additively or solely in the ratios. In Table 3, earnings are never statistically significant when the ratio is included, suggesting the latter. This panel shows that when wages are not scaling the retirement variables, they are significant and of the appropriate sign. The accrual variable that includes pensions, however, is now positive and significant. Since this does not occur in Model (5) which excludes pension accruals, this anomaly can be attributed to the extremely disperse distribution of pension accruals documented in Table 2.[17]

The middle panel of Table 4 includes age linearly in the specification from Table 3. The probability of being retired naturally increases with age, and it is not surprising to find a strong, positive effect in both models. When age is included, the retirement variables retain the same signs and approximate magnitudes but are generally less significant. To further investigate the age dimension of the retirement model, the third panel estimates a semiparametric hazard model with age at retirement as the measure of spell length. In both models, the retirement present value is positive and significant, but the accruals are positive and insignificant. Recall that in a hazard model, the effect of the covariates (once exponentiated) is to multiply the underlying baseline probability of retirement at each age. In effect, the hazard model attributes the entirety of the large spikes in the accrual in Figure 2 at ages 62 and 65 to factors not explicitly modeled; the covariates can only reflect deviations from the underlying baseline hazard. However, any of the other age-related phenomena that might affect retirement such as a taste for leisure or unobserved declines in health status should be monotonic with age; the empirical hazard function is not. Moreover, Samwick (1993) and Figure 1 show that pension plan formulas do in fact have age and tenure dependent components that can generate nonmonotonicities similar to those in Figure 2. In light of this misspecification,

[17]The specification with retirement variables in levels is extremely sensitive to the presence of outliers. If the individuals with top and bottom 1 percent of nominal pension accruals are excluded from the sample, then the coefficient on retirement wealth accruals is negative and significant. The analogous coefficient in the specification which normalizes by wages also becomes more negative and significant, but the change is comparatively small and does not affect the qualitative results.

it is not surprising that the effect of retirement wealth accruals is blunted in the hazard model.

Table 5 presents three other variations on the estimates in Table 3. The first panel changes the definition of retirement to include any separation from the job held in 1983. The results for Model (4) are qualitatively the same, with slightly smaller and less significant coefficients. The same is true of Model (5), with the difference that the pension eligibility variable is now negative and insignificant. This suggests that although persons with pensions may be more likely to retire, they are much less likely to switch to another job without retiring.

The second panel excludes all women, public sector workers, and individuals under 58 years of age. This makes the sample more like the typical RHS sample used in previous research. The results for Model (4) are again quite similar to those in Table 3, but Model (5) shows that Social Security is relatively more important for this subsample through both the size of the coefficient and, informally, the log-likelihood being better for this model than the one that accounts for pension incentives. One explanation is that by restricting the sample to only those 58 and over, much of the important variation in pensions is excluded. Figure 1 showed that very large pension accruals can happen as early as age 55 for typical pensions. Another explanation is that Social Security is less important for public sector workers because much of their entitlement is based on work at previous jobs and therefore not as greatly affected by the decision to continue working at the present job.

The third panel attempts to explain the finding in Table 3 that financial assets are never significant in the retirement models. By excluding the individuals with extremely high financial assets, these specifications attempt to isolate an effect of other wealth for more typical workers. Although the coefficients are larger in magnitude, they are still quite close to zero and insignificant. Like most demographic variables, financial assets are not found to be important determinants of retirement. This lends credibility to the model of Stock and Wise (1990a), which necessarily omits such variables by using personnel records only.

Though far from definitive, the results of Table 3 do suggest that the earlier studies of Social Security and retirement were not hampered in their estimates by omitting pensions. Both Social Security and pensions are still found to be statistically significant in basic retirement models. The magnitudes of the coefficients on Social Security variables are not affected by specifying pensions as wealth and accruals instead of just pension indicators. When pensions and Social Security variables are combined into total retirement income variables, present values are relatively more important and accruals are relatively less important in predicting retirement behavior.

V. Simulation Results

The results of the previous section demonstrate that both Social Security and pensions have significant effects on the probability of retirement. In this respect, the SCF yields similar results to the RHS. In this section, the prior conclusions that changes in the Social Security program have contributed only slightly to the decrease in labor force participation will be addressed by simulating the effects of past and present changes on predicted retirement behavior. Additionally, the conjecture by Stock and Wise (1990b) that pensions are responsible for these declines can be evaluated by simulating changes in pension coverage. Furthermore, by comparing the simulated results of Models (4) and (5), it is once again possible to assess the effect of the better pension data on the analysis of policy changes.

The first simulation is based on the level of Social Security benefits. Past studies have claimed that over the course of the 1970s there was an unexpected increase of 20-25 percent in Social Security benefits (see Burtless 1986). To undo this change, I assume a 20 percent decrease in the Social Security PIA (for all earnings levels) from its level in 1983 and calculate the new predicted probability of retirement for each observation.[18] These new probabilities yield a new cumulative distribution function for the probability of being retired by each age. For a decrease in benefit levels of this magnitude, both Models (4) and (5) predict an decrease in the cumulative probability of retiring between ages 50 and 70 of 3.5 percentage points. Assuming there was very little retirement before age 50, this represents approximately one third of the total decline over the period (Lumsdaine and Wise, 1990).

Figure 3 shows the important contribution of the PPS to this analysis. For both models, the fraction of this 3.5 percentage point change that is completed by each age is graphed. When pensions are modelled explicitly rather than through a simple indicator

[18]The predictions for Model (4) will also account for the corresponding effects of the change in Social Security on the magnitude of integrated pension plan benefits. In practice, however, these offsets had little effect on the predictions. As described in Samwick (1993), roughly half the private plans are integrated, and they impose around a 50 percent tax on the *change* in Social Security benefits, which itself will be small in comparison with the magnitude of the pension benefits.

Figure 3

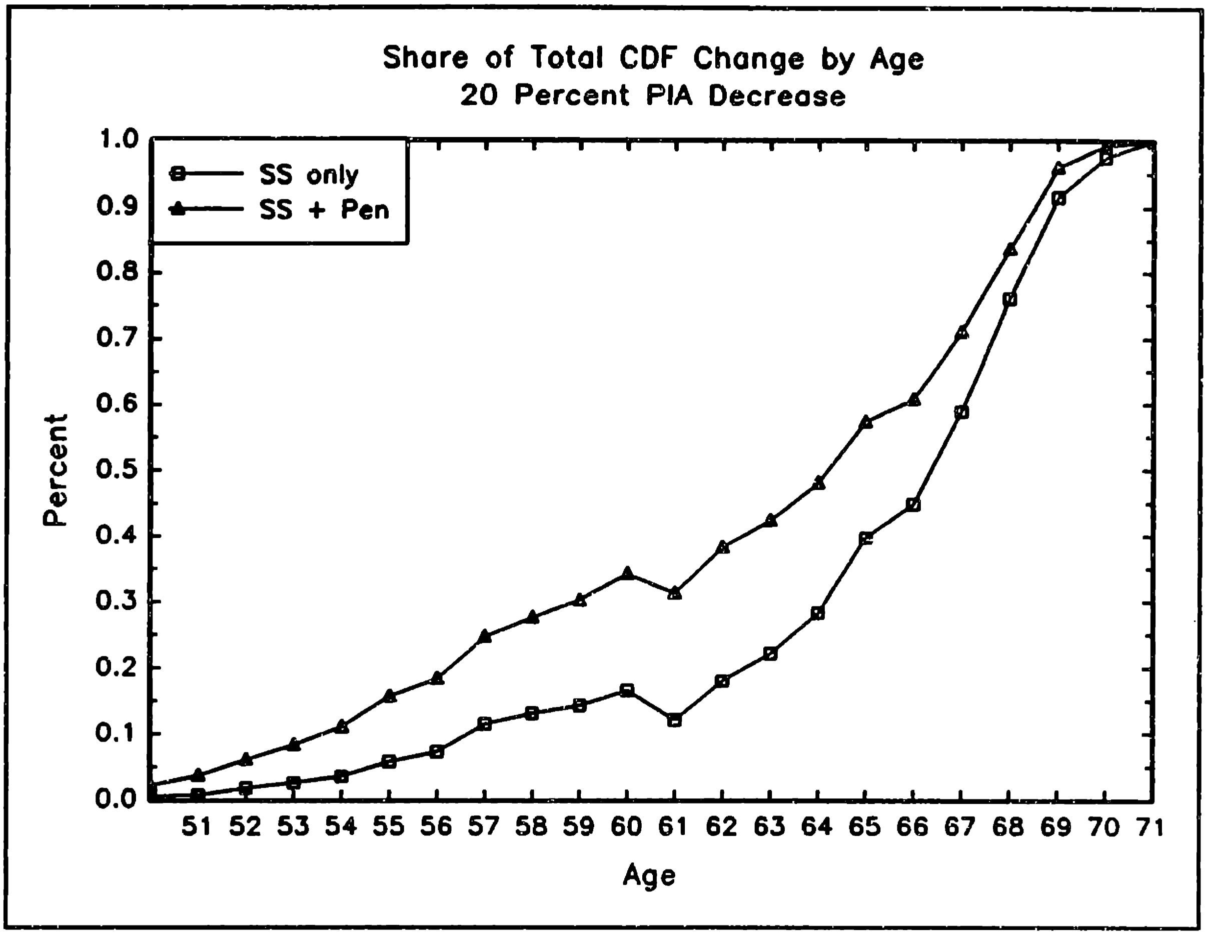

variable, more of the change is predicted to occur at earlier ages. The fraction of the change that is predicted to occur before age 62 and Social Security benefits can actually be collected doubles from 20 to 40 percent. Failing to account for pension incentives increases the predicted effect of Social Security after it can be collected at the expense of the effect in years before it can be collected. This finding also suggests that one way pensions affect retirement is by relaxing liquidity constraints on workers who would like to retire before age 62 but lack sufficient financial assets to support consumption until the benefits actually commence.

Legislation enacted in 1983 raised the normal retirement age for Social Security receipt from 65 to 67, with the increase being phased in over several decades. Applying the same methodology as in the benefit level simulation, both Models (4) and (5) predict a decrease in the cumulative probability of retirement between 50 and 70 of 2.1 percentage points.[19] Thus, roughly 60 percent of the increase in retirement that resulted from the reforms in the 1970s will eventually be undone by the legislation of the early 1980s. Figure 4 shows again that when pension incentives are modelled explicitly, a larger fraction of the change in retirement behavior occurs at earlier ages. In the case of the increase in the NRA, however, both models predict the bulk of the changes between the early and normal retirement ages for Social Security.

Finally, it is possible to reassess the importance of pensions suggested by the analysis of Stock and Wise (1990b) in light of the more representative sample of employees with pensions. One systematic change would be that employers were mandated to provide a pension to all employees. In the current sample, that corresponds to approximately a 50 percent increase in coverage. The actual pension formula corresponds to that graphed in Figure 1 above.

In this case, Models (4) and (5) suggest very different results. In the latter, which constrains the effect of this change to affect the eligibility variable only, the total change in the predicted CDF for retirement between 50 and 70 is 4.3 percentage points.

---

[19] The increase in the NRA also increases the reduction for taking retirement benefits early at age 62 and decreases the delayed retirement credits for postponing retirement. The latter is reflected in Figure 4 by the drop off in the CDF change at later ages. As in the case of the general benefit decrease, changing the NRA on Social Security has ramifications on the level of pension benefits for pension plans that coordinate payments with Social Security benefits, but the impact of this feature is limited in this analysis.

Figure 4

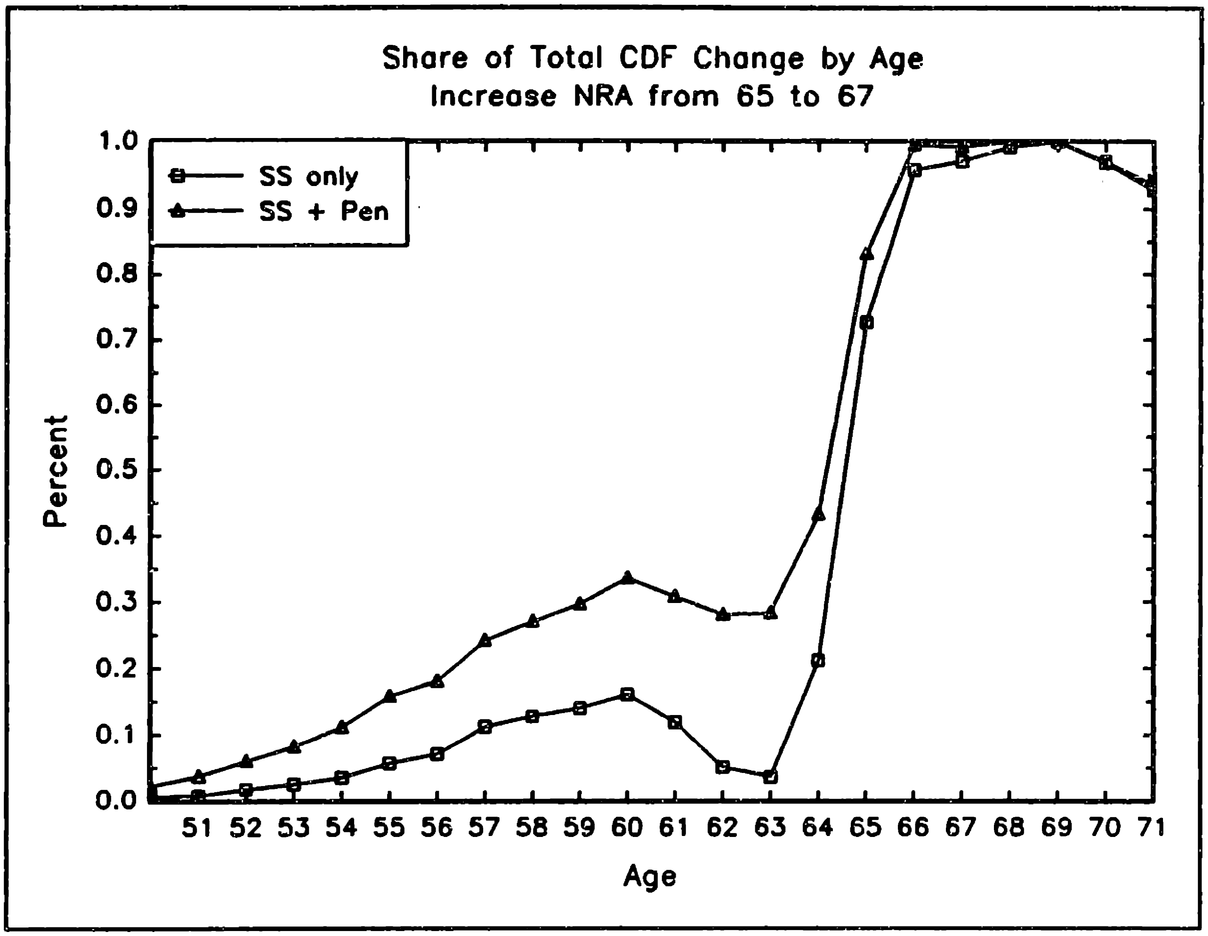

In the former, which incorporates the pensions directly, the change is a somewhat smaller 3.8 percentage points. Both predictions are at least as large as those from the changes in Social Security simulated earlier.

Figure 5 shows that the two models have very different implications for the timing of the changes. Model (5) predicts that the change in retirement probabilities will be proportional to the predicted CDF itself, whereas Model (4) shows that the change will be concentrated in roughly equal parts in the intervals 60-65 and 65-71. This particular result may be due to the pension's normal retirement age falling between 60 and 65, but the general result is that the predicted changes can be extremely sensitive to the detail in the data.

Figure 5

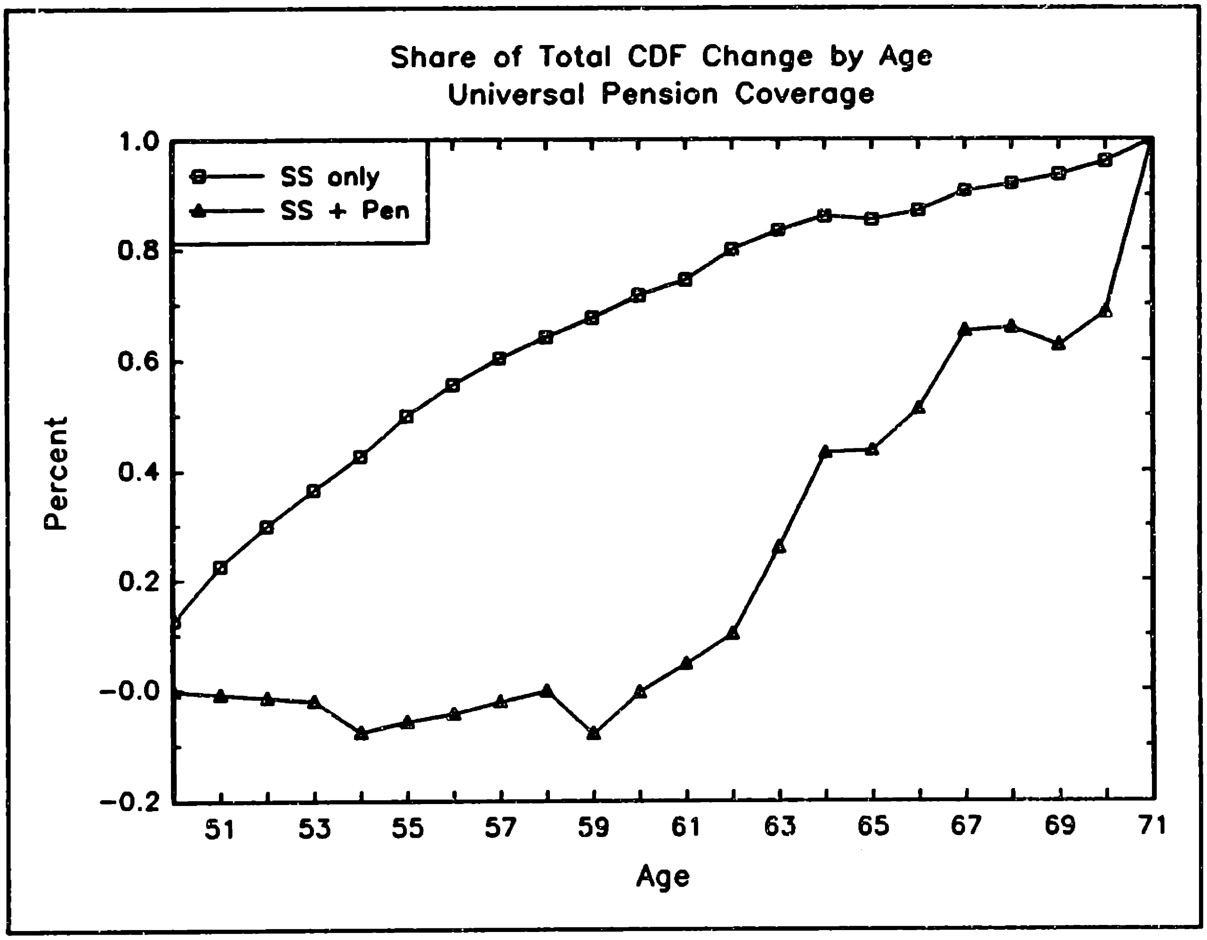

VI. Conclusion

This paper demonstrates that previous findings of statistically significant but economically modest effects of Social Security on retirement are robust to more careful accounting for pension incentives in a richer dataset. Simulations reveal that changes in the Social Security system on the order of those that occurred in the 1970s can explain about one third of the contemporaneous decrease in labor force participation during that period. Other simulations of the effects of the 1983 legislation suggest that about 60 percent of this earlier change will eventually be undone by the higher normal retirement age.

Using the new data, it is also possible to independently consider the effects of increased pension coverage on the probability of retirement. Pensions are also estimated to have statistically significant effects on retirement behavior, which extends the recent work of Stock and Wise (1990b) to a more representative dataset. The simulated effects of increasing pension coverage to those who are not currently eligible for pensions are at least as large as the simulated effects of changing Social Security benefits by 20 percent. Thus, it is possible that increased pension coverage explains yet another third of the decline in labor force participation, although the pension simulation is admittedly quite crude. A better approach would be to actually estimate the structural model of Stock and Wise (1990a) on the new data to get more comparable results. This endeavor is the subject of work in progress.

Finally, the usefulness of datasets like the PPS can now be considered in light of the differences in results presented here that do and do not make use of the pension detail. In most of the estimated models, the inclusion of pensions did not appreciably affect the magnitude of the Social Security coefficients. What was different was the *timing* of all policy simulations. By carefully accounting for the differences in the retirement opportunity set excluding Social Security, the pension data allows for a more accurate analysis of the incidence of policy changes across different age groups. In particular, the finding that about 40 percent of the change in behavior occurs before benefits from Social Security can actually be collected suggests that one role pensions play in affecting retirement is by relaxing financial liquidity constraints for employee's who would otherwise like to retire but not have sufficient assets or cash flow to sustain consumption.

References

Allen, Steven G., Robert L. Clark, and Ann A. McDermed (1992). "Post-Retirement Benefit Increases in the 1980s," in John A. Turner and Daniel J. Beller (eds.) *Trends in Pensions 1992*. Washington, DC: United States Department of Labor.

Andrews, Emily S. (1985). *The Changing Profile of Pensions in America*. Washington, DC: Employee Benefit Research Institute.

Atkinson, Anthony B. (1987). "Income Maintenance and Social Insurance," in Alan J. Auerbach and Martin S. Feldstein (eds.) *Handbook of Public Economics*, Vol. 2. New York: North Holland.

Avery, Robert B., Gregory E. Elliehausen, and Glenn B. Canner (1984a). "Survey of Consumer Finances, 1983," *Federal Reserve Bulletin*, Vol. 70, 679-692.

Avery, Robert B., Gregory E. Elliehausen, and Glenn B. Canner (1984b). "Survey of Consumer Finances, 1983: A Second Look," *Federal Reserve Bulletin*, Vol. 70, 857-868.

Avery, Robert B., Gregory E. Elliehausen, and Arthur B. Kennickell (1987). "Changes in Consumer Installment Debt: Evidence from the 1983 and 1986 Surveys of Consumer Finances," *Federal Reserve Bulletin*, Vol. 73, 761-778.

Bloom, David E. and Richard B. Freeman (1992). "The Fall in Private Pension Coverage in the United States," *American Economic Review*. Vol 82, No. 2, 539-545.

Board of Governors of the Federal Reserve System (1990). *Balance Sheets for the U.S. Economy, 1945-89*. Washington, DC.

Boskin, Michael and Michael D. Hurd (1978). "The Effect of Social Security on Early Retirement," *Journal of Public Economics*, Vol. 10, 361-378.

Burtless, Gary (1986). "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies*, Vol. 53, 781-805.

Burtless, Gary and Jerry A. Hausman (1978). "The Effect of Taxation on Labor Supply: Evaluating the Gary Negative Income Tax Experiment," *Journal of Political Economy*. Vol. 86, 1103-1130.

Burtless, Gary and Robert A. Moffitt (1984). "The Effect of Social Security Benefits on the Labor Supply of the Aged," in Henry J. Aaron and Gary Burtless (eds.) *Retirement and Economic Behavior*. Washington, DC: The Brookings Institution.

Chen, Yung-Ping (1992). "The Role of Private Pensions in the Income of Older Americans," in John A. Turner and Daniel J. Beller (eds.) *Trends in Pensions 1992.* Washington, DC: United States Department of Labor.

Diamond, Peter A. and Jerry A. Hausman (1984). "The Retirement and Unemployment Behavior of Older Men," in Henry J. Aaron and Gary Burtless (eds.) *Retirement and Economic Behavior.* Washington, DC: The Brookings Institution.

Faber, Joseph (1982). *Life Tables for the United States, 1900-2050.* Washington, DC: Social Security Administration.

Fields, Gary S. and Olivia S. Mitchell (1985). *Retirement, Pensions, and Social Security.* Cambridge, MA: MIT Press.

Han, Aaron and Jerry A. Hausman (1990). "Flexible Parametric Estimation of Duration and Competing Risk Models," *Journal of Applied Econometrics*, Vol. 5, 1-28.

Hausman, Jerry A. (1981). "Labor Supply," in Henry J. Aaron and Joseph A. Pechman (eds.) *How Taxes Affect Economic Behavior.* Washington, DC: Brookings Institution, 27-84.

Hausman, Jerry A. and David A. Wise (1985). "Social Security, Health Status, and Retirement," in David A. Wise (ed.) *Pensions, Labor, and Individual Choice.* Chicago: University of Chicago Press.

Heckman, James and Burton Singer (1984). "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, Vol. 52, No. 2, 271-320.

Kotlikoff, Laurence J. and David A. Wise (1985). "Labor Compensation and the Structure of Private Pension Plans: Evidence for Contractual Versus Spot Labor Markets," in David A. Wise (ed.) *Pensions, Labor, and Individual Choice.* Chicago: University of Chicago Press.

Kotlikoff, Laurence J. and David A. Wise (1987). "The Incentive Effects of Private Pension Plans," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.) *Issues in Pension Economics.* Chicago: University of Chicago Press.

Kotlikoff, Laurence J. and David A. Wise (1989). "Employee Retirement and a Firm's Pension Plan," in David A. Wise (ed.) *The Economics of Aging.* Chicago: University of Chicago Press.

Lancaster, Tony (1979). "Econometric Methods for the Duration of Unemployment," *Econometrica*, Vol. 47, No. 4, 939-956.

Lazear, Edward P. and Robert L. Moore (1988). "Pensions and Turnover," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.) *Pensions in the U.S. Economy*. Chicago: University of Chicago Press.

Lumsdaine, Robin L., James H. Stock, and David A. Wise (1992). "Three Models of Retirement: Computational Complexity Versus Predictive Validity," in David A. Wise (ed.) *Topics in the Economics of Aging*. Chicago: University of Chicago Press.

Lumsdaine, Robin L. and David A. Wise (1990). "Aging and Labor Force Participation: A Review of Trends and Explanations," National Bureau of Economic Research Working Paper No. 3240.

Meyer, Bruce D. (1990). "Unemployment Insurance and Unemployment Spells," *Econometrica*, Vol. 58, No. 4, 757-782.

Mitchell, Olivia S. (1992). "Trends in Pension Benefit Formulas and Retirement Provisions," in John A. Turner and Daniel J. Beller (eds.) *Trends in Pensions 1992*. Washington, DC: United States Department of Labor.

Quinn, Joseph, Richard Burkhauser, and Daniel Myers (1990). *Passing the Torch: The Influence of Economic Incentives on Work and Retirement*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Ransom, Roger L. and Richard Sutch (1988). "The Decline of Retirement and the Rise of Efficiency Wages: U.S. Retirement Patterns, 1870-1940," in R. Ricardo-Campbell and E. Lazear (eds.) *Issues in Contemporary Retirement*. Stanford, CA: Hoover Institution, 3-37.

Rust, John (1989). "A Dynamic Programming Model of Retirement Behavior," in David A. Wise (ed.) *The Economics of Aging*. Chicago: University of Chicago Press.

Samwick, Andrew A. (1993). "Retirement Incentives in the 1983 *Pension Provider Survey*," mimeo, Massachusetts Institute of Technology (also Chapter Three in this thesis).

Social Security Administration (1990). *Social Security Bulletin Annual Statistical Supplement*. Washington, DC: Department of Health and Human Services.

Steinberg, Richard M. and Harold Dankner (1987). *Pensions: A Financial Reporting and Compliance Guide*. New York: John Wiley & Sons, Inc.

Stock, James H. and David A. Wise (1990a). "Pensions, the Option Value of Work, and Retirement," *Econometrica*, Vol. 58, No. 5, 1151-1180.

Stock, James H. and David A. Wise (1990b). "The Pension Inducement to Retire: An Option Value Analysis," in David A. Wise (ed.) *Issues in the Economics of Aging*. Chicago: University of Chicago Press.

Sueyoshi, Glenn T. (1989). "Social Security and the Determinants of Full and Partial Retirement: A Competing Risks Analysis," National Bureau of Economic Research Working Paper No. 3113.

Table 1

Comparison of Household Data in SCF and RHS				
	Survey of Consumer Finances[a]		Retirement History Survey[b]	
Variable	Mean	Std. Deviation	Mean	Std.Deviation
Education (years)	12.583	3.080	10.805	3.447
Poor health	0.011	0.106	0.319	0.466
Married	0.760	0.427	0.890	0.313
Non-white	0.118	0.323	0.097	0.296
Household size	2.638	1.284	2.532	1.240
Clerical, Service & Sales	0.354	0.478	0.219	0.413
Craftsman	0.163	0.369	0.247	0.431
Laborer	0.127	0.333	0.292	0.455
Manager	0.174	0.379	0.129	0.335
Professional & Technical	0.189	0.392	0.067	0.251
Earnings (x $1000)	22.028	25.571	20.654	14.314
Pension Eligibility	0.654	0.476	0.301	0.459
Number of Observations	525		1633	

Notes:
1) Author's calculations from SCF 1983. Relevant sample is all fulltime workers (excluding military and self-employed) ages 50 to 69 in 1983 who are re-interviewed in SCF 1986 and not omitted from the PPS if they report being covered by a pension in the SCF. Data are weighted.
2) SCF sample is 46 percent female and 30 percent employed by the public sector. The RHS sample contains only males working in the private sector.
3) Taken from Sueyoshi (1989), Table 3. Earnings are adjusted to 1983 dollars.

Table 2

Pension and Social Security Wealth and Accruals in 1983	Levels			Ratio to Wages		
	Mean	Median	Std Dev	Mean	Median	Std Dev
Full Sample						
Social Security Wealth	47,018	46,437	25,551	3.243	2.577	4.110
Social Security Accrual	-805	-544	2,111	-0.071	-0.026	0.246
Pension Wealth	29,710	0	63,330	1.138	0	1.812
Pension Accrual	956	0	9,767	0.042	0	0.299
Earnings	22,028	17,440	25,571			
Financial Assets	58,132	8,986	355,771			
Housing Net Equity	58,931	46,506	95,136			
Eligible for both Social Security and Pensions						
Social Security Wealth	52,535	51,134	23,616	2.835	2.582	1.715
Social Security Accrual	-851	-703	2,110	-0.048	-0.035	0.128
Pension Wealth	39,152	17,877	65,823	1.526	0.947	1.761
Pension Accrual	1,873	55	11,830	0.075	0.003	0.359
Earnings	23,707	19,378	26,098			
Financial Assets	62,038	11,045	416,629			
Housing Net Equity	55,755	47,475	86,661			

Notes:
1) Author's calculations using SCF and PPS.
2) Sample size is 525 for full sample, 305 for sample restricted by eligibility.
3) See text for explanation of variables.

Table 3

Retirement Probits -- Basic Results					
Variable	(1)	(2)	(3)	(4)	(5)
Constant	-1.3462 (0.3217)	-1.3245 (0.3212)	-1.3996 (0.3530)	-1.4144 (0.3291)	-1.3225 (0.3301)
Education	-0.0295 (0.0199)	-0.0134 (0.0201)	-0.0202 (0.0206)	-0.0139 (0.0201)	-0.0207 (0.0201)
Poor Health	0.1837 (0.5055)	0.1127 (0.5089)	0.1676 (0.5068)	0.1297 (0.5081)	0.2086 (0.5113)
Married	0.2120 (0.1715)	0.0883 (0.1735)	0.0974 (0.1747)	0.0974 (0.1731)	0.0966 (0.1716)
Female	0.1440 (0.1307)	0.1697 (0.1328)	0.2055 (0.1345)	0.1749 (0.1329)	0.1466 (0.1305)
Nonwhite	0.0096 (0.2119)	0.0300 (0.2162)	0.0458 (0.2161)	0.0266 (0.2164)	0.0042 (0.2130)
Household Size	-0.1305 (0.0685)	-0.1269 (0.0691)	-0.1284 (0.0702)	-0.1271 (0.0690)	-0.1307 (0.0688)
Financial Assets	0.00003 (0.00004)	0.00003 (0.00003)	0.00003 (0.00004)	0.00003 (0.00004)	0.00003 (0.00004)
Housing Net Equity	-0.0005 (0.0006)	-0.0007 (0.0006)	-0.0006 (0.0006)	-0.0007 (0.0006)	-0.0006 (0.0006)
Earnings	0.00009 (0.0009)	-0.0001 (0.0009)	0.00007 (0.0009)	0.00008 (0.0009)	-0.0004 (0.0009)
Social Security Present Value	0.0435 (0.0119)	-----	0.0051 (0.0194)	-----	0.0033 (0.0195)
Social Security Accrual	-----	-0.8067 (0.1876)	-0.8211 (0.2740)	-----	-0.8198 (0.2719)
Pension Present Value	0.1004 (0.0281)	-----	0.0800 (0.0325)	-----	-----
Pension Accrual	-----	-0.6388 (0.2300)	-0.3246 (0.2289)	-----	-----
SS plus Pension Present Value	-----	-----	-----	0.0217 (0.0154)	-----
SS plus Pension Accrual	-----	-----	-----	-0.5808 (0.1828)	-----
Pension Eligibility	-----	-----	-----	-----	0.1426 (0.1288)
Log-Likelihood	266.458	262.718	259.877	261.959	266.759

Notes:
1) Number of workers: 525 Number of observations is 1401.
2) All dollar amounts are in thousands.
3) Retirement variables are ratios of present value or accrual to the wage level.

Table 4

Retirement Models -- Variations of Basic Specification						
	Use Levels, not Ratios		Enter Age Linearly		Hazard Model, not Probit	
Variable	(4)	(5)	(4)	(5)	(4)	(5)
Constant	-1.0541 (0.3030)	-1.3487 (0.3449)	-5.9243 (0.9347)	-6.2950 (0.9469)	-----	-----
Education	-0.0358 (0.0195)	-0.0315 (0.0201)	-0.0155 (0.0213)	-0.0147 (0.0210)	0.0047 (0.0410)	-0.0064 (0.0414)
Poor Health	0.1466 (0.5087)	0.1810 (0.5167)	0.2336 (0.5095)	0.3089 (0.5132)	0.8729 (1.0241)	0.9047 (1.0324)
Married	0.1234 (0.1699)	0.0546 (0.1743)	0.0082 (0.1808)	0.0232 (0.1811)	-0.2658 (0.3600)	-0.1561 (0.3654)
Female	0.1189 (0.1305)	0.2217 (0.1472)	0.2490 (0.1408)	0.2247 (0.1390)	0.4342 (0.2859)	0.3769 (0.2791)
Nonwhite	0.0037 (0.2081)	0.0054 (0.2138)	-0.0031 (0.2286)	-0.0113 (0.2257)	0.0719 (0.4323)	0.1214 (0.4353)
Household Size	-0.1299 (0.0671)	-0.1039 (0.0683)	-0.0599 (0.0681)	-0.0639 (0.0680)	-0.0410 (0.1249)	-0.0739 (0.1261)
Financial Assets	0.000008 (0.00004)	0.00002 (0.00004)	0.00001 (0.00004)	0.00002 (0.00004)	0.00005 (0.00008)	0.00005 (0.00008)
Housing Net Equity	-0.0006 (0.0006)	-0.0009 (0.0007)	-0.0009 (0.0007)	-0.0008 (0.0007)	-0.0038 (0.0018)	-0.0037 (0.0018)
Earnings	-0.0035 (0.0016)	-0.0014 (0.0011)	-0.0003 (0.0010)	-0.0007 (0.0010)	-0.0009 (0.0023)	-0.0015 (0.0023)
Social Security Present Value	-----	0.0023 (0.0030)	-----	0.0178 (0.0181)	-----	0.0549 (0.0277)
Social Security Accrual	-----	-0.0865 (0.0228)	-----	-0.2286 (0.1338)	-----	0.4196 (0.3797)
SS and Pension Present Value	0.0021 (0.0007)	-----	0.0256 (0.0147)	-----	0.0507 (0.0250)	-----
SS and Pension Accrual	0.0037 (0.0019)	-----	-0.2252 (0.1636)	-----	0.1472 (0.2875)	-----
Pension Eligibility	-----	0.1269 (0.1279)	-----	0.1824 (0.1338)	-----	0.4612 (0.2734)
Age	-----	-----	0.0742 (0.0142)	0.0808 (0.0142)	-----	-----
Log-Likelihood	271.441	264.534	247.498	249.381	334.580	333.851

Notes:
1) Total number of observations is 1401.
2) All dollar amounts in thousands.
3) See text for further description.

Table 5

Retirement Models -- Variations of Basic Specification						
	Use General Turnover, not Retirement		RHS sample Restrictions		Exclude Financial Assets over $50,000	
Variable	(4)	(5)	(4)	(5)	(4)	(5)
Constant	-0.7879 (0.2575)	-0.6844 (0.2574)	-1.7207 (0.5534)	-1.8970 (0.5867)	-1.2956 (0.3965)	-1.1763 (0.3981)
Education	-0.0146 (0.0157)	-0.0164 (0.0157)	-0.0108 (0.0322)	-0.0268 (0.0337)	-0.0012 (0.0264)	-0.0092 (0.0259)
Poor Health	0.3037 (0.3858)	0.3033 (0.3852)	0.5694 (0.5841)	0.9374 (0.6054)	0.1998 (0.5102)	0.2978 (0.5159)
Married	0.1176 (0.1339)	0.1075 (0.1133)	0.1944 (0.3639)	0.3130 (0.3720)	-0.0044 (0.2043)	0.0430 (0.2029)
Female	0.0972 (0.1024)	0.0607 (0.1011)	-----	-----	0.1242 (0.1809)	0.0289 (0.1784)
Nonwhite	-0.0161 (0.1689)	-0.0398 (0.1681)	-0.1057 (0.3834)	-0.1234 (0.3790)	0.0820 (0.2363)	0.0597 (0.2288)
Household Size	-0.1550 (0.0528)	-0.1527 (0.0522)	-0.0909 (0.0986)	-0.1252 (0.1044)	-0.1247 (0.0847)	-0.1436 (0.0854)
Financial Assets	0.00002 (0.00003)	0.00003 (0.00003)	0.000003 (0.00005)	0.000004 (0.00005)	-0.0037 (0.0064)	-0.0039 (0.0063)
Housing Net Equity	-0.0012 (0.0006)	-0.0011 (0.0007)	-0.0010 (0.0010)	-0.0008 (0.0009)	-0.00006 (0.0032)	0.0004 (0.0031)
Earnings	-0.000006 (0.0007)	-0.0001 (0.0007)	0.0007 (0.0010)	0.0004 (0.0010)	-0.0090 (0.0095)	-0.0120 (0.0107)
Social Security Present Value	-----	0.0201 (0.0142)	-----	-0.0048 (0.0250)	-----	0.003 (0.0214)
Social Security Accrual	-----	-0.3119 (0.2070)	-----	-1.0980 (0.3748)	-----	-0.7598 (0.2801)
SS and Pension Present Value	0.0154 (0.0131)	-----	0.0295 (0.0165)	-----	0.0022 (0.0195)	-----
SS and Pension Accrual	-0.4752 (0.1469)	-----	-0.3710 (0.2139)	-----	-0.8950 (0.2465)	-----
Pension Eligibility	-----	-0.0815 (0.0966)	-----	0.6328 (0.2417)	-----	0.1809 (0.1684)
Log-Likelihood	467.828	474.772	103.686	99.580	172.168	178.829

Notes:
1) Total number of observations in the three panels are 1401, 696, 888.
2) All dollar amounts in thousands.
3) See text for further description.

Chapter Five

Wage Risk Compensation Through Employer-Provided Pensions

I. Introduction

In a series of articles, Eaton and Rosen (1980a, 1980b, 1980c) and Varian (1980) propose that redistributive taxation can improve welfare by providing insurance against labor income risk. A key assumption made is that "the market fails to provide insurance against the vagaries of wage rates."[1] The justification for this assumption is that moral hazard problems associated with insuring the returns to human capital overwhelm private efforts to provide insurance. In more recent work, Kaplow (1991a, 1991b) has pointed out that moral hazard problems which limit private insurance are not remedied by having the government provide the insurance, as the tax authority is no more capable of observing effort (and thereby negating moral hazard) than is a private insurer. Therefore, any model which generates a result that some redistributive taxation is optimal also proves that some degree of private insurance should exist despite the moral hazard problem. Kaplow goes on to show that given the existence of private insurance arrangements, it is inefficient for the government to offer insurance through the tax system (if moral hazard is the only concern) because it distorts the price of the private insurance.

The issue of whether the government has any role to play in relieving labor income risk rests crucially on the answer to two questions not addressed in past studies. First, do workers actually face much uncertainty in their labor income over their working careers? Empirical estimates of the uncertainty in labor income are rare, yet there is no moral hazard problem in the absence of income risk. Second, are there any private

[1]Quoted in Eaton and Rosen (1980a, p. 366). Varian (1980) also cites moral hazard as one reason private insurance markets do not exist.

institutions which shift the burden of income uncertainty between firms and workers actually in place in the economy? If so, then Kaplow's criticism of the earlier studies is valid, and any government arrangements for insuring risk that propose to increase welfare must take account of the existing private insurance mechanisms.

This paper establishes that the answer to both questions is affirmative. Using the methods developed in Topel and Ward (1992) and Carroll and Samwick (1992) and data from a large firm and the *Panel Study of Income Dynamics* (PSID), the stochastic component of within-job wages is estimated to be quite large. Thus, there is a reason to expect some private arrangements to spread this risk to exist, even if no formal market for human capital exists. Using the *Survey of Consumer Finances* (SCF) and the companion *Pension Provider Survey* (PPS), it is then demonstrated that employer-provided pension plans, which cover approximately half of the workforce, often contain features which serve both to compensate employees for bearing labor income uncertainty during their working careers and to filter it out of their post-retirement income. Furthermore, it is precisely those workers with high income uncertainty that have pensions which have such features. This finding is consistent with the use of the pension formula as a device for counterbalancing the utility decreasing consequences of stochastic wage processes.

It should be noted at the onset that the premise here is not that the primary reason that pensions are offered is to compensate workers for bearing labor income risk. Rather, the origin of pension plans is well described by Ransom, Sutch, and Williamson (1992) to be an outgrowth of two economic phenomena of the early twentieth century.

First, following a scandal in the insurance industry in 1906, many states outlawed the most popular form of individually-purchased pension. This created a void in the provision of retirement annuities to be filled. Second, the growth of internal labor markets led to a gradual lengthening of the implicit employment contract between firms and workers. The institutions of mandatory retirement and steeper optimal age-earnings profiles put forth in Lazear (1979) subsequently emerged. In this context, pensions become necessary to provide the appropriate incentives for work effort for older workers.[2] The subsequent growth of pensions can be linked to the tax incentives provided in the Revenue Act of 1942 and the Employee Retirement Income Security Act of 1974 (ERISA).

Within this general framework, however, the firm has considerable discretion in structuring the pension plan formula. The most important decision is to determine the relative size of pension benefits for workers with different job histories and wage profiles in a straightforward manner. Firms have historically chosen to make the relationship between lifetime wages and the pension benefits nonlinear, and the question being addressed is whether the risk-compensating aspects of these relationships are systematically related to the amount of wage uncertainty to which the workers are exposed.

---

[2]In particular, if there is a lag between "shirking" and detection, then some amount of lifetime income must be withheld until after the worker leaves the job to account for this lag. Since pensions are paid only after retirement, they serve this function where a steep age-earnings profile does not. Additionally, as Lazear (1983) discusses, the early retirement features of many pension plans are akin to severance pay, which is a feature of the optimal wage contract that helps to control turnover for older workers.

The organization of the paper is as follows. In the next section, a simple model of pension valuation in the presence of uncertainty in labor income is developed. Simulations are used to illustrate the potential for pensions to improve welfare in the presence of stochastic wages. Institutional features of typical pension plans that compensate for income risk are identified in Section III. Section IV documents the degree of uncertainty in within-job wages and the extent to which it differs across individuals in the population. In Section V, it is shown that the degree to which a worker's pension possesses these desirable insurance properties is significantly related to the amount of income uncertainty in his wage process. That is, the distribution of pension plans in the population is consistent with their use as mechanisms to compensate for wage uncertainty. There is a brief concluding section which also discusses the implications of these results for Social Security policy, precautionary saving, and the valuation of pensions more generally.

II. The Effects of Pensions on Welfare with Wage Uncertainty

In order to illustrate the effect of pensions on expected utility when wages are stochastic, I numerically solve a three period life-cycle model with a very simple structure. The employee earns a certain wage in the first period, an uncertain wage in the second period, and a retirement pension in the third period. The employee's optimization is to choose a value for saving in each of the first two periods. By varying the relationship between the pension benefit and the random wage paid in the second period in ways that correspond to typical pension formulas, it is possible to show that a risk-averse employee prefers some pension formulas to others. In particular, when the uncertainty in the wage is not insurable and the pension must be related to actual earnings paid, such an employee will prefer a pension formula that is convex in wages to one that is not.

Figure 1

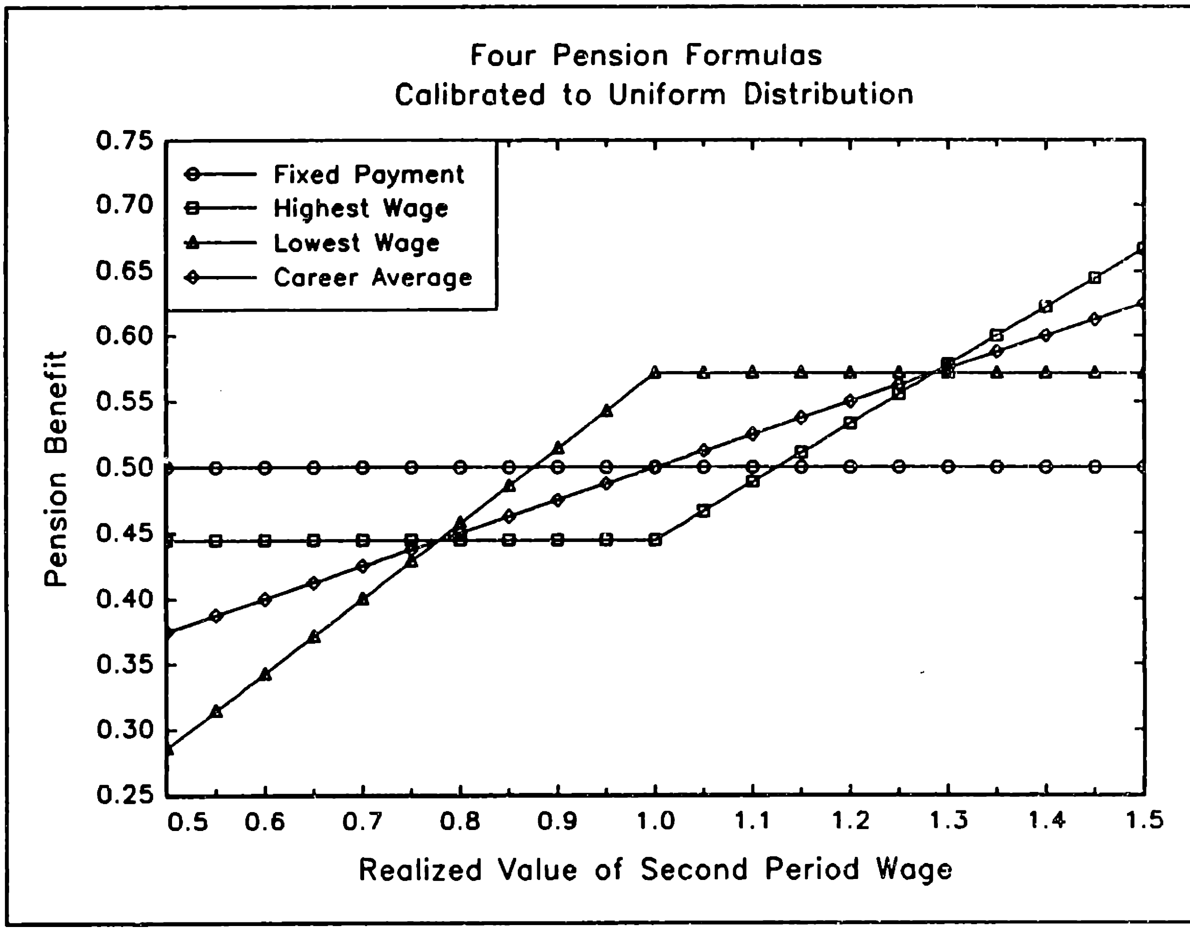

I consider four types of pensions of the form $p = p(w_1,w_2,\alpha)$, where w_t is the wage paid in period *t* and by assumption w_2 is a random variable in the first period. The four plans are graphed in Figure 1.[3] The actual pension formulas take the form given in the following chart:

[3]All pensions in Figure 1 have values of w_1 and E(p) normalized to 1 and 0.5, respectively, and w_2 distributed uniformly between 0.5 and 1.5.

Pension Type	Pension Formula
Fixed Payment	$p = \alpha_1$
Highest Wage	$p = \alpha_2 * \max(w_1, w_2)$
Lowest Wage	$p = \alpha_3 * \min(w_1, w_2)$
Career Average	$p = \alpha_4 * (w_1 + w_2)$

The α's are chosen (conditional on the ex-ante distribution of w_2) so that all formulas yield benefits of equal expected value. Thus, any differences in welfare or behavior under different pension regimes will be due to the pensions' treatment of the uncertainty in the second period wage.

Two of the pension formulas are linear in the second period wage. In the Fixed Payment (or Flat Rate) pension, the benefit amount does not depend on the actual value of w_2 that is paid and is therefore constant for all possible values of w_2.[4] The Career Average pension will increase linearly with the actual value of w_2 because the benefit amount is just a weighted average of w_1 and w_2. The other two formulas contain a nonlinearity in the benefit schedule at the point where $w_1 = w_2$. In the case of the Highest Wage pension, the benefit formula is convex in the value of w_2, whereas the benefit formula for the Lowest Wage pension is concave in the value of w_2.

[4]This statement is not necessarily true if the pension formula is renegotiated over the employee's tenure at the firm. To the extent that the firm's profitability is reflected in the adjustments made to the formula over time, then firm- or industry-specific shocks to earnings will also affect the value of the pensions. In this case, the Fixed Payment pension will begin to resemble the Career Average pension. Since all of the other pension formulas are subject to this renegotiation, it is still accurate to claim that Fixed Payment pensions are valuable in this model by removing the individual-specific component of income variation from the pension calculation.

The insurance consequences of these nonlinearities are informative. Relative to the Career Average pension, the Highest Wage pension trades off lower pension values when w_2 is close to its mean for higher values when w_2 is far from its mean. By contrast, the Lowest Wage pension trades off higher pension values when w_2 is close to its mean for lower values when w_2 is far from its mean relative to the Career Average pension. For an employee who is risk-averse, the tradeoff implicit in the Highest Wage pension is to be preferred to that of the Lowest Wage pension. In this model, a bad state of the world is identified by a low value of w_2. Thus, the Highest Wage pension compensates relatively more when income has been low (and the marginal utility of consumption is consequently higher). Similarly, the Fixed Payment pension corresponds to a fully insured pension; the pension amount is the same regardless of the value of w_2.

The importance of this tradeoff is depicted in Figure 2, which graphs expected utility as a function of the degree of uncertainty in the distribution of w_2. Formally, the model being solved in each case is:

$$
\begin{aligned}
&\max_{\{c_1,c_2\}} \sum_{t=1}^{3} \beta^{t-1} E u(c_t) \\
&\text{s.t.} \\
A_{t+1} &= (1+r)A_t + w_t - c_t \\
A_1 &= A_4 = 0 \\
w_2 &= \begin{cases} w_1(1+\gamma), & \pi \\ w_1(1-\gamma), & 1-\pi \end{cases} \\
w_3 &= p(w_1, w_2, \alpha) \\
u(c_t) &= \frac{c_t^{1-\rho}}{1-\rho}
\end{aligned} \tag{1}
$$

where c_t and w_t are consumption and income during period t and A_t is non-pension wealth at the beginning of period t. The first two conditions define the accumulation of non-pension savings assuming no bequests are received or given. The next condition specifies the distribution of w_2 to be binomial with mean w_1 and variance parameterized by γ. The next condition specifies that income in the third period is comprised of the retirement pension, and the last condition assumes Constant Relative Risk Aversion utility with coefficient ρ. Figure 2 plots the maximized value of Equation 1 obtained by solving the dynamic programming problem against the value of γ.[5]

For all values of γ, the Fixed Payment pension affords the greatest expected utility. Among the earnings-related pensions, the Highest Wage pension yields the greatest expected utility, followed by the Career Average and Lowest Wage pensions. That is, for earnings-related pensions in this model, welfare is increasing with the

[5]Because the model is homothetic in wages, w_1 is normalized to unity, giving the interpretation of percents of first period wage to all nominal amounts below. Other parameters are given reasonable values: $r = 0.05$, $\beta = 0.95$, $\rho = 3$, $\pi = 0.50$, and $\alpha_1 = 0.50$. The qualitative results are unaffected by other sensible choices of these parameters; they are naturally weakened if a nonstochastic component of earnings growth which is large relative to the stochastic part (γ) is present.

Figure 2

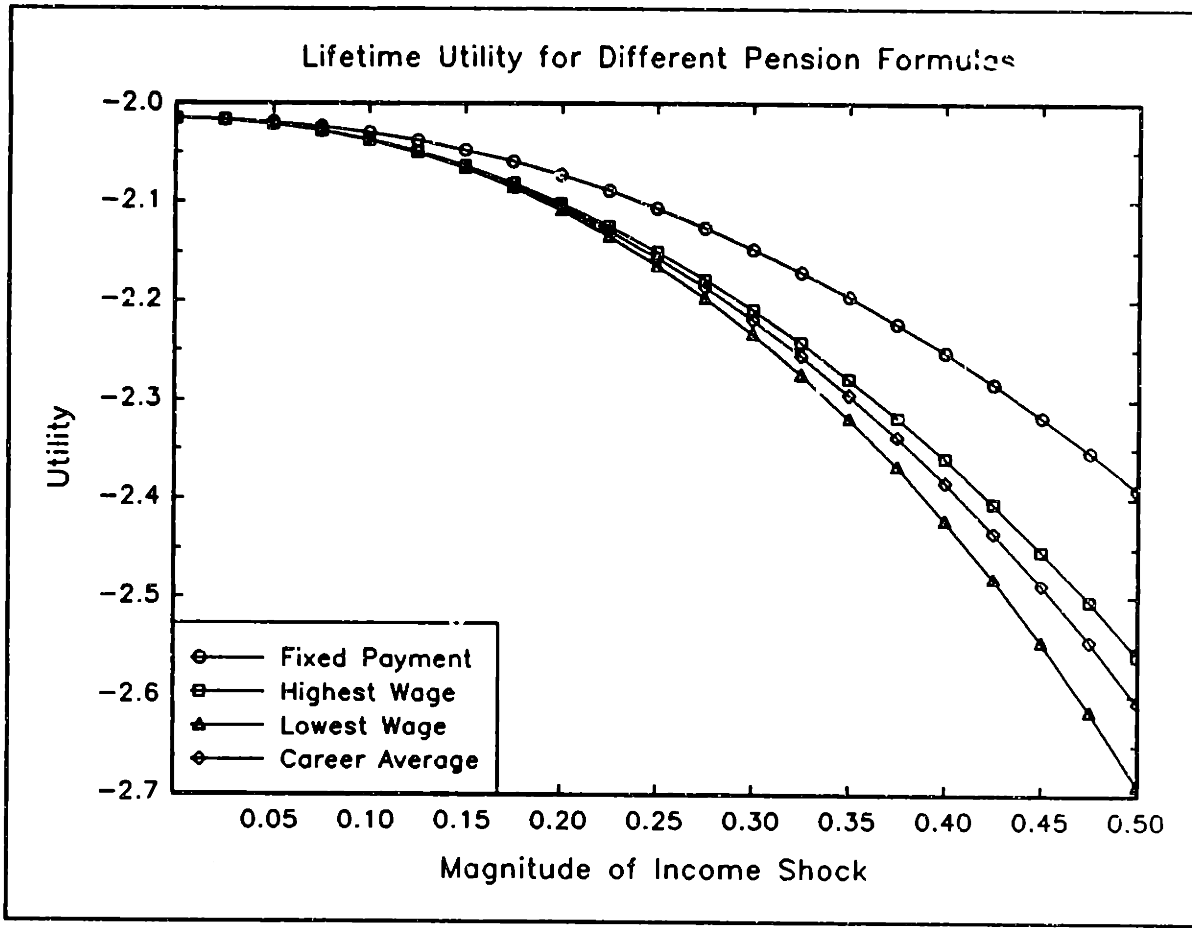

convexity of the pension formula.[6]

In order to measure how much better off a worker would be under each formula, the equivalent variation of each earnings-related pension relative to the Fixed Payment plan is graphed in Figure 3. The equivalent variation is the amount of income a worker would need under an earnings-related pension to be as well off as under the Fixed

[6]This result does not generalize transparently. It will be the case that the convex pension outperforms an equally concave pension as they are defined here, i.e. the pension is (weakly) increasing in the realized value of w_2 and the kink is located at the same value in each formula. However, it is possible to make the upward sloping portion of the schedule in Figure 1 so steep that the linear pension formula is preferred to the convex pension. In the simulation model described here, such reversals did not occur until the slope was increased from 1 to almost 3. The actual pension features that give rise to convex formulas to be discussed in Section III are typically less than 1, so the empirical importance of these pathological cases is likely to be trivial.

Figure 3

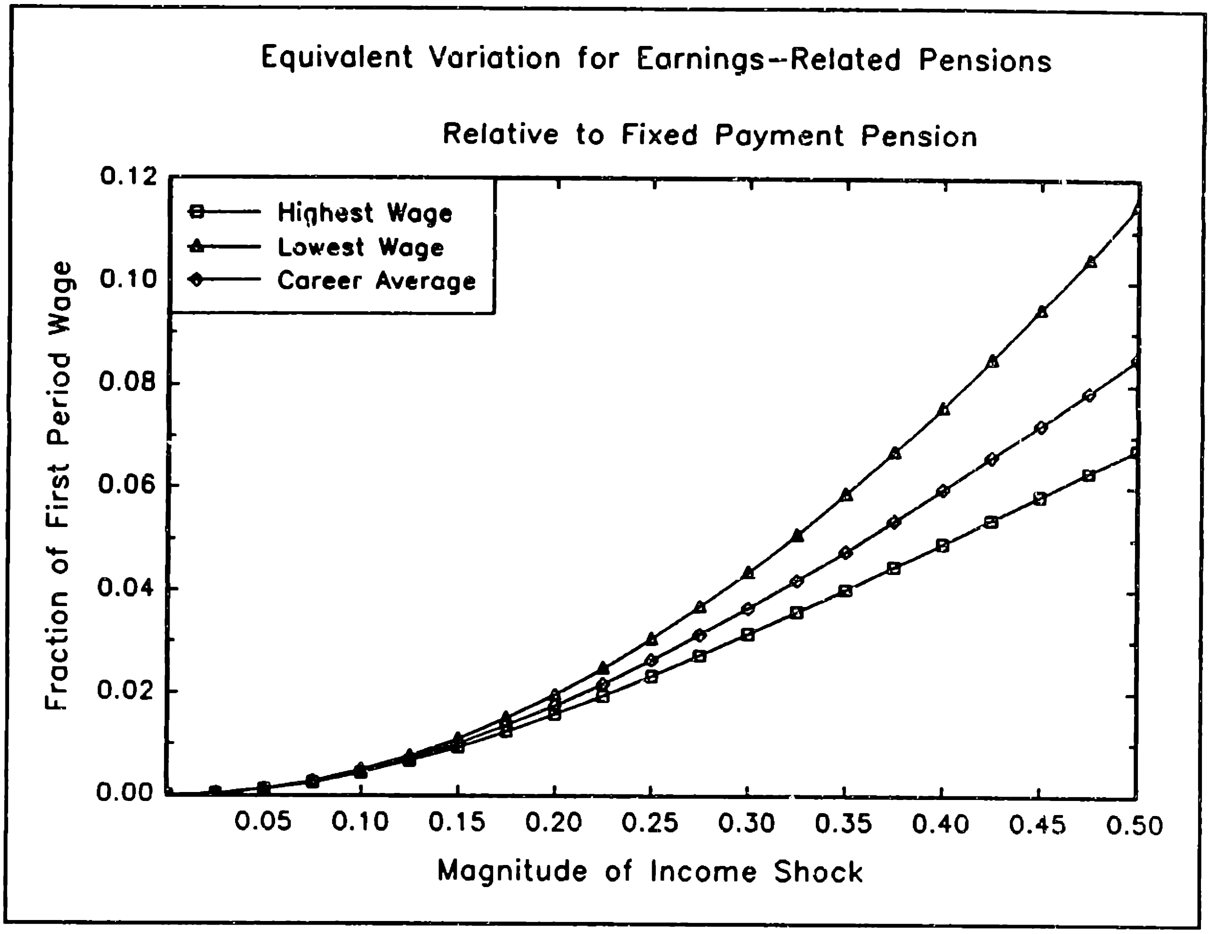

Payment pension of equal expected value. Figure 3 shows that when γ is 0.50, a worker would be indifferent between the Fixed Payment pension, the Highest Wage pension with $A_1 = 0.07$, the Career Average pension with $A_1 = 0.085$, and the Lowest Wage pension with $A_1 = 0.115$.

The choice of pension formula will also affect economic behavior. Since the utility function is CRRA, the individual has a precautionary saving motive.[7] Figure 4 graphs the amount of first-period saving (the only saving decision made under uncertainty

[7]The welfare results above hold for any utility function that exhibits risk aversion. In order for the optimal decision variable (savings in the first period) to be affected by the income uncertainty, the utility function must also exhibit prudence, which is a convex first derivative. Kimball (1990) develops this theory and gives several examples.

Figure 4

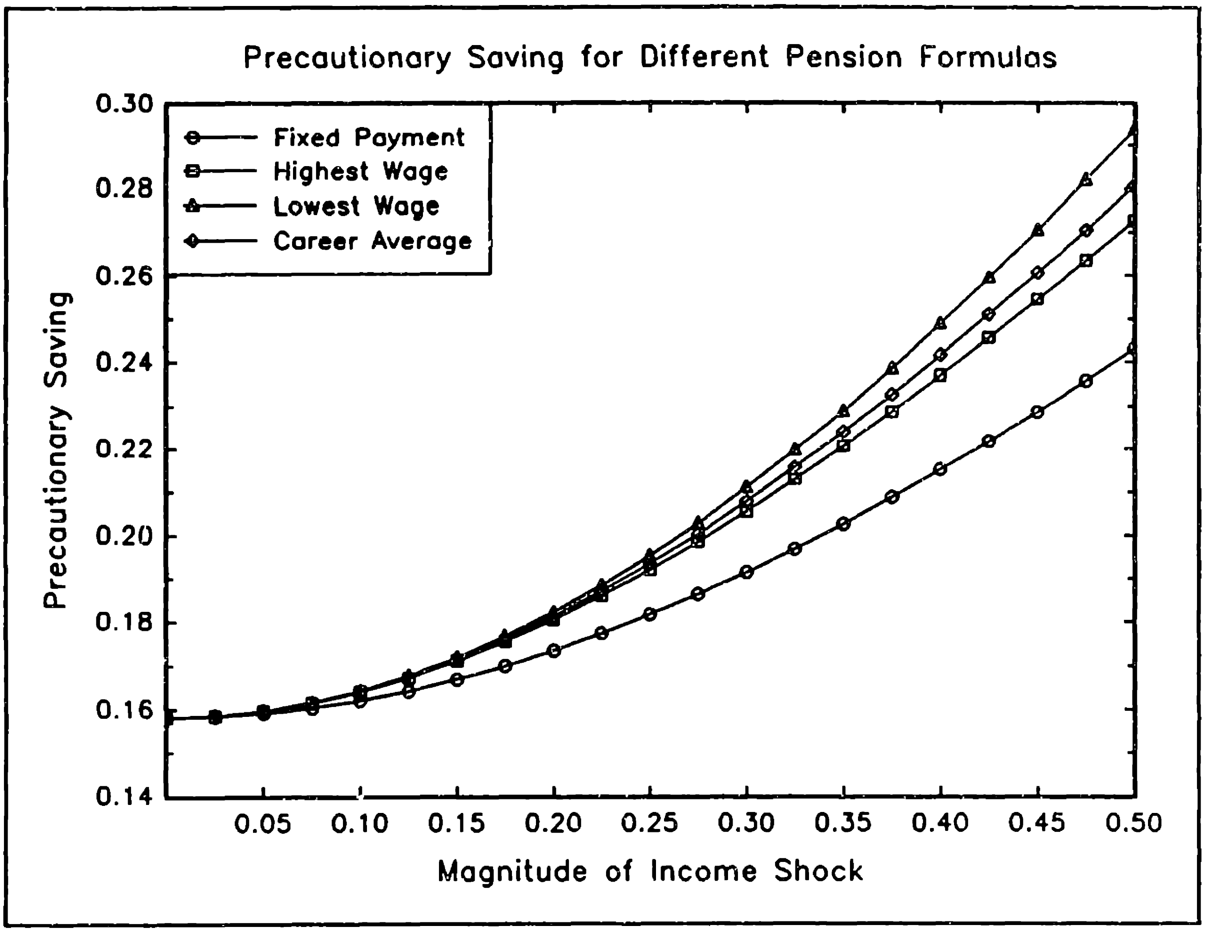

in this model) against the value of γ. The amount of precautionary saving is inversely related to the convexity in the earnings-related pension plan (subject to the same caveats listed in Footnote 6). The graph shows that without any uncertainty, approximately 16 percent of first period income would be saved for life-cycle reasons alone. When uncertainty is as high as that of $\gamma = 0.50$ and the pension is unrelated to earnings (Fixed Payment), an additional 8 percent is saved. At this level of uncertainty, having an earnings-related pension of any form will result in approximately 4 percent more income being saved. Earnings-related pensions therefore increase precautionary saving by approximately 50 percent in this model.[8]

Two main results emerge from these simulations. First, the pension that is unrelated to actual earnings received is superior to any earnings-related pension formula of equal expected value. This is because less of future income in the first period is uncertain. Desirable incentive effects of total compensation on effort may prevent companies from severing the tie between productivity and any portion of compensation, however, making widespread use of Fixed Payment plans infeasible. Second, of those pensions which are related to the actual wages received, the formula with the convex kink (Highest Wage) yields higher lifetime utility than the linear schedule (Career Average) and the concave kink (Lowest Wage).

III. Risk Compensating Features of Pension Formulas

The simulation results in Section II established that two types of pensions will be particularly valued when wages are stochastic: those with fixed payments regardless of the individual worker's wage level and those with formulas that are convex in actual earnings paid. The former are easy to identify in the PPS, as only Defined Benefit plans with no measure of Final Average Pay in the benefit formula can qualify. Identifying the convexity in pension plan formulas is a more formidable task, as the curvature of the relationship between pensions and wages will depend on the interaction of many of the pension plan's details.

---

[8]Carroll and Samwick (1993) find that 25 to 40 percent of wealth holdings are attributable to earnings uncertainty. Assuming life-cycle saving in the second period is approximately 16 percent on average, this model yields a figure of 20 percent without an earnings-related pension and 30 percent with an earnings-related pension under the highest uncertainty level specified.

The following lemma demonstrates that one metric of the convexity of pension formulas is the extent to which the expected value of pension benefits increases when the uncertainty in the wage process increases.

LEMMA:
Given two pension plans, p(w) and q(w), and an arbitrary distribution of wages, F(w,r), where increases in r reflect mean-preserving spreads[9] to the distribution of wages,

$$p''(w) > q''(w) \quad \forall\, w \in [\underline{w},\overline{w}] \rightarrow \frac{dE[p(w)]}{dr} > \frac{dE[q(w)]}{dr}, \quad \forall\, w \in [\underline{w},\overline{w}]$$

Stated another way, the lemma establishes that any pension plan which is more convex than another at every wage level will have a relatively larger increase in the expected value of the pension benefit for a given increase in the uncertainty in the wage process.[10] Moreover, the effect of wage uncertainty on the expected value of the pension gives a very natural interpretation to the insurance role being played. When wage uncertainty increases, so does the expected value of the pension; the pension automatically compensates the worker for bearing more wage uncertainty.

---

[9]This terminology is from Rothschild and Stiglitz (1970). Given a family of distributions for w, F(w,r), increases in the index r will be mean-preserving spreads if the following two conditions are satisfied:

$$\int_{\underline{w}}^{\overline{w}} F_r(w,r)dw = 0$$

$$\int_{\underline{w}}^{w} F_r(z,r)dz \geq 0 \quad for\ \underline{w} \leq w \leq \overline{w}$$

[10]The case for using the change in the expected value of the pension benefit with a higher level of wage uncertainty would be strengthened if the causality could be demonstrated in the reverse direction as well. The lemma currently shows only that if one pension has an expected value that is more sensitive to the uncertainty in the wage process than another, it cannot be the case that the first is everywhere less convex than the second. Given the nature of the pension plan features (described below) that actually generate the convexity of pension benefit schedules, however, the sensitivity of the expected pension benefit to wage uncertainty is likely to be a good predictor of the convexity of the pension formula.

Given this metric for indexing the convexity of the pension formula in the wage, the next task is to examine whether actual pension formulas do vary in the ways that distinguish the four pensions in Figure 1. The three features of pension plan formulas that contribute most directly to compensation for wage risk are the definition of Final Average Pay, integration with provisions of the Social Security benefit formula, and multiple benefit schedules.

In order to systematically compare the impact of these features, it is useful to adopt a more general notation than that of Section II. The time-path of wages over a worker's career will be denoted by $\mathbf{w} = \{w_1,\ldots,w_T\}'$ and the pension $p(\mathbf{w})$ will be characterized by the α_t's in the equation:

(2)
$$p(w,Z) = \alpha_0(Z) + \sum_{t=1}^{T} \alpha_t(w,Z)w_t$$
$$w \sim \left[\mu, \Sigma\right]$$

where t indexes the years in which the wages were earned and **Z** contains nonwage information relevant to the pension calculation (e.g. date of birth, date of hire, retirement date). The key issue is whether $E[p(\mathbf{w})]$ increases with (the elements of) Σ, the variance of **w**. The elements of Σ loosely correspond to the parameter r that indexes the uncertainty in wage distributions; references to a "larger Σ" are meant to imply that uncertainty in **w** is higher.

To demonstrate, two benchmark pension formulas can be identified. A "Flat Rate" pension specifies the pension benefit as a function of age at retirement and years of service but *not* pre-retirement earnings. Thus, $\alpha_t(\mathbf{w},\mathbf{Z}) = 0, \forall\, t > 0$, corresponding to the Fixed Payment pension discussed above. A "Percent of Earnings" (POE) pension specifies the pension benefit as a linear function of pre-retirement income. Final Average Pay (FAP) is a weighted average of wages each period, and the benefit itself is a linear function of FAP. That is, $\partial\alpha_t/\partial w_t = 0, \forall\, t > 0$ as in the "Career Average" pension described above.

The first way a pension can have E[p(w)] vary with Σ is to have FAP defined as a nonlinear function of pre-retirement earnings. For example, most earnings-related DB plans define FAP as the best *x* years out of a specified period of *y* years of earnings. Different realizations of the stochastic component of wages can alter the actual years that enter the computation of FAP. This provision is a more general version of the Highest Wage pension above and allows the values of $\{\alpha_t\}$ to depend on realized $\{w_t\}$ in a way that weights good income draws more heavily. A larger Σ will result in larger positive shocks, and this provision mitigates the otherwise offsetting effects of the larger negative shocks by substituting different years in FAP when they occur. Thus, E[p(w)] increases with Σ.

It is important to note that the inflation rat will have meaningful effects on this provision because most pension formulas do not index **w** to the growth in the price level when selecting the highest wage years.[11] If inflation is extremely high, then FAP is essentially a function of the wage at retirement and there is no scope for switching years in the computation. If inflation is low and real wage growth is flat, then the best years are simply those with the largest positive shocks, regardless of how big they are. Thus,

[11]This contrasts with the Social Security benefit formula, which indexes most elements of w for not only increases in the price level but economy-wide real wage growth as well.

this feature is operative only when expected nominal wage growth is of roughly the same magnitude as the standard deviation of the shock to income.[12]

The second and more important avenue for E[p(w)] to increase with Σ is through the integration of pension and Social Security benefits in either Excess or Offset plans.[13] An Excess plan allows the marginal replacement rate on FAP to be higher for the component of FAP in excess of the Social Security Maximum Taxable Earnings (MTE) than it is for the component below the MTE. If these two rates are r_2 and r_1, respectively, then expected pension benefits can be expressed as

$$E[p(w)] = r_1 E[FAP(w)] + (r_2 - r_1) * E[\max(0, FAP(w) - MTE)] \tag{3}$$

The second term on the right side of equation (3) is analogous to the max(·) term in the Highest Wage pension formula. It will increase with the variance of FAP(w), as a higher variance implies a greater probability of having this term strictly positive. This analysis pertains to any convex kink in the marginal replacement rate schedule, which by law may not occur above the MTE and in practice is not initiated below it. This would seem to imply that integration only affects workers with relatively high earnings (the MTE in 1992 was \$55,500), but for some older plans the kink was fixed at the nominal value of the prevailing MTE (e.g. \$10,800 in 1973) rather than the MTE in the

---

[12]Nominal wage growth has averaged approximately 5 percent since the 1950s, whereas average values for the permanent shock to earnings estimated below are between 10 and 15 percent. I am indebted to Larry Katz for proposing a link between inflation and the interpretation of this feature.

[13]The direct effect of integration is to increase the average replacement rate with pre-retirement earnings. ERISA and subsequent legislation permit this in part to mitigate the progressivity of the Social Security benefit schedule. Samwick (1993) discusses Social Security integration in the context of pension accrual rates. Merton, Bodie, and Marcus (1987) discuss the role of integration in insuring against unforeseen changes in Social Security rules.

year benefits are first paid. Thus, it is possible to find such kinks remaining at wage levels below the current MTE.

An Offset plan subtracts a fraction of the employee's primary social security benefit from the earnings related component of the pension formula. Letting r_2 be the replacement rate on FAP and r_1 be the offset rate on Social Security benefits (SSB), the offset plan has a formula similar to

$$E[p(w)] = E[\max(0, r_2 FAP(w) - r_1 SSB(w))] \tag{4}$$

If FAP(w) and SSB(w) were perfectly correlated, then there would be no effect of Σ on E[p(w)] because the probability of $r_2 \bullet FAP(w) - r_1 \bullet SSB(w)$ being negative would be invariant to Σ. However, less than perfect correlation will exist between FAP(w) and SSB(w) for at least two reasons. First, Social Security benefits are based on 35 years of wages, whereas offset plans are generally based on 5 that are within 10 years of the retirement date, and the relevant wage history for Social Security will include years spent on other jobs and hence be unaffected by Σ for the current job. Second, the Social Security benefit formula is extremely progressive (i.e. the replacement rate on pre-retirement income declines with the level of income), making it a concave function of earnings, whereas FAP is typically convex for reasons given above. Thus, large positive shocks will disproportionately affect FAP relative to SSB, and the effect will be more pronounced at higher wage levels.[14]

[14]Nonetheless, one might wish to consider the effects of uncertainty on total retirement income. To do this, the term SSB(w) would be added to the right-hand side of equation (10), making it resemble the Excess plan in equation (4). Expected total retirement income would therefore still be sensitive to Σ in the manner described here.

The 1980s witnessed a sustained increase in the fraction of plans that are integrated with Social Security. Mitchell (1992, Table 9.11) reports that the fraction of private defined benefit pensions in medium and large firms which integrated with Social Security increased from 45 percent in 1980 to 63 percent in 1989. Over this period, two thirds of the integrated plan take the Offset form. Despite the Tax Reform Act of 1986's limits on the amount of the "spread" permissible across the MTE (r_2-r_1 in equation (3)) in Excess plans and the magnitude of the offset rate (r_1 in equation (4)) in Offset plans, integrated plans and their consequences for risk compensation continue to be important.

The third way for E[p(w)] to vary with Σ is the direct imposition of floors or ceilings on the level of benefits or the magnitude of FAP in the pension formula. In many pensions, there are several normal retirement benefit schedules, with the final benefit paid defined as the maximum or minimum of those computed under the several formulas. The most common example is a pension that gives the maximum of a Flat Rate formula and an earnings-related formula. As shown above, the convexity of the max($\cdot$) function and the concavity of the min($\cdot$) function in their arguments imply that E[p(w)] will be increasing or decreasing with Σ, respectively. It may also be the case that, even when there is only one formula, the FAP itself has a floor or a ceiling, which generate the same results as a max($\cdot$) or min($\cdot$) applied to two different formulas. In practice, such a pension plan provision will dominate the effects of any other nonlinearities, and about 2 percent of the plans in the PPS contain explicit limits on FAP.

Table 1 presents information on these characteristics of pension plans in the *Pension Provider Survey* (PPS). Included are the 924 plans under which someone under

70 and currently working in the SCF expected to receive benefits.[15] All figures are weighted by the sample weights of the households covered. The top panel gives a breakdown of the plans by the number of different normal retirement schedules they contain. The more schedules, the more likely it is that the plan takes a max($\cdot$) or a min($\cdot$) over them to make the benefits a nonlinear function of wages. A few of the plans contain different benefit schedules based on years before and after ERISA, but most adjust for the new requirements by creating a new FAP definition rather than a new benefit schedule.[16] Over half of the plans are DB plans with 2 or more normal retirement schedules, and over one quarter have 3 or more.

The middle panel of the table reports that about 25 percent of the plans are integrated with Social Security, with most of them taking the offset form. This figure is substantial but still less than that reported in Mitchell (1992) for 1983. One reason is that government workers are included in this sample but not Mitchell's, and, as of 1983, federal government workers and many state government workers were not covered by Social Security. Thus, there would be no reason for their plans to be integrated. As documented in Samwick (1993), government plans are not only present but overrepresented in the PPS because the attrition rate for small private firms between the SCF and the PPS is relatively high.

---

[15]This set of plans is identical to that of Samwick (1993).

[16]ERISA was the first major piece of pension legislation to establish vesting and funding standards. Firms with plans not meeting these standards would not qualify for the tax advantages of the pension, so most changed their plans (if necessary) to comply.

The bottom panel of Table 1 shows the percentage of plans in the PPS having various types of FAP formulas. Recall that only FAPs which allow for the best years over a particular period induce a dependence of E[p(w)] on Σ. The last four lines show that at least half of all plans have such a feature. Less than half of this set of plans formally require that the years come in the final year or years of the job, but most do focus only on blocks of consecutive years rather than individual years. Nearly 10 percent are Flat Rate plans for which Σ has no effect on any of the pension moments. About 14 percent have at least one FAP which is based on the career average pay.[17] A DC plan is similar to a Career Average plan in this respect, as contributions proportional to wages are made in each year a person is covered by a DC plan. Thus, roughly 30 percent of plans have a benefit schedule which uses all years of earnings.

Table 1 shows that there is considerable opportunity for pensions to compensate employees for larger Σ. Plans also differ in the particular method through which this is accomplished. Integration with Social Security, minimum and maximum benefit formulas, and the definition of FAP are some of the plan features that can be used. Additionally, about 10 percent of the pensions are unrelated to earnings altogether, corresponding to the Fixed Payment pension in Section II, which would generally dominate other earnings-related pensions of the same expected value.

[17]The sum of the percentages in the bottom panel exceeds 100 because some plans have more than one FAP associated with them. Most of the overlap involves plans having both a career average plan and one of the FAPs based on the highest years in some period. What is of interest here is the fraction of all plans with a particular feature, regardless of the other features it also has.

IV. Estimating Within-job Wage Processes

The modeling of risk-compensating features of pensions thus far presupposes that workers do face wage uncertainty during their tenure with a given employer. In order to test the implications of that model, therefore, it is necessary to obtain an estimate of the uncertainty in annual within-job wages for each employee. Unfortunately, the job history and wage information in the *Survey of Consumer Finances* (SCF) is not adequate to compute such a measure because the longitudinal dimension of the data is limited to at most two observations on wages at an interval of three years. Econometric estimates of wage processes using panel data such as Abowd and Card (1989), MaCurdy (1982), and Lillard and Willis (1978) have found evidence of a highly persistent--if not permanent--shock to earnings levels. Such a shock cannot be identified without a panel of earnings or unpalatably stringent distributional assumptions.

To make up for this shortcoming in the dataset that is linked to the pension formulas, a long panel of earnings from the *Panel Study of Income Dynamics* (PSID) will be substituted. Using this data, it is possible to estimate the magnitude of the shocks to wages for each worker. The distribution of income shocks will then be parameterized in terms of variables present in both surveys, and the predicted values based on the data in the SCF will be used in subsequent analysis. The main results of this analysis are that the uncertainty in within-job wages is substantial and systematically different across groups in the population.

Any reasonable method of measuring wage uncertainty in microeconomic data will rely on the differences between the actual wages received and what a consistent, ex-ante wage forecast predicted for them. Ideally, the forecast used would be the worker's own expectation of future wages, but in practice household surveys do not solicit information on the expected distribution of wages. Instead, a wage process based on observable characteristics of the worker must be estimated and the residuals used to measure uncertainty. Because the information sets of the worker and the econometrician will differ with the former generally including the latter, measures of uncertainty based on such residuals will likely overestimate the true uncertainty faced by the worker.

Two steps are taken in the following analysis to safeguard against this caveat. First, two estimation strategies that are robust against different forms of misspecification in the wage process are used. Both frameworks yield similar parameter estimates. Second, before presenting the results for the PSID, both methods will be implemented on a sample of workers derived from the personnel records of a single large firm. The parameter estimates are also similar across the two types of datasets. Taken together, the generality of the estimation methods and the comparability of results across datasets suggest that the effect of this data shortcoming is not too important for subsequent analysis.

The firm from which the personnel records are drawn is a *Fortune* 500 company with a defined benefit pension plan.[18] The sample used here consists of male workers aged 39 to 55 in 1969 who had positive earnings in each year between 1969 and 1980, inclusive. Drawing the sample from the same firm ensures that in addition to working in the same occupation and having the same gender, all the workers operate in the same

---

[18]This sample and the details of its pension plan are discussed in Kotlikoff and Wise (1989).

industry and are likely to have similar educational backgrounds. Measured uncertainty will therefore not be strongly influenced by any of these factors. The sample restriction that the workers must be employed over the twelve year period makes the estimate of uncertainty a conservative one, as any extreme shocks to earnings that cause a worker to leave the firm force that job spell to be excluded from the sample. Additionally, personnel records reported directly by the firm are less likely than earnings reported in other surveys to suffer from measurement error, which would tend to inflate the estimated wage uncertainty. The finding of sizable wage uncertainty in this conservatively chosen and homogeneous sample lends credibility to similar findings in the PSID sample.

The extract from the PSID contains data on heads of households from 1976-1987. The virtue of using the PSID is that, given the basic results established using the single firm, it allows variation in wage uncertainty across characteristics like occupation, industry, and education to be explored. Although there is considerable noise in the PSID related to earnings data, especially in identifying changes in employers, the mean estimates are quite close to those found using the single firm.

A Model of Wage Uncertainty

A standard model for the wage process at a particular job is given by:

$$w_{jt} = X_{jt}\beta + \theta_j + \epsilon_{jt} \qquad (5)$$

where w_{jt} is the natural logarithm of the real wage, X_{jt} contains economic variables describing job j at time t, θ_j is a fixed effect specific to job j which allows for heterogeneity across jobs in the earnings capacity for an individual worker, and ϵ_{jt} is a time-varying random component of measured earnings.[19] For example, if θ_j were equal to zero and ϵ_{jt} were equal to 0.1, then the actual wages received on job *j* in year *t* would be 10 percent higher than their expected value based on the information in X_{jt}.

The stochastic component of the wage process that is of interest for studying the risk-compensating effects of pensions is ϵ_{jt} rather than θ_j. High variance in the match-specific component of wages suggests that turnover in the early years of a job is to be encouraged so that bad matches are weeded out and only good matches remain. Thus, the workers bear the burden for this uncertainty precisely by severing their ties with a firm, thereby rendering pensions--which are explicitly long term contracts--ineffective as insurance mechanisms. In this case in particular, a higher average level of wages in the early years of a job is a more effective method of compensating for the risk due to θ_j.[20]

By contrast, the movements in ϵ_{jt} correspond to productivity shocks against which workers would like to be insured and firms, with their greater access to capital markets and diversification, are more capable of bearing. To fix ideas, ϵ_{jt} can be further decomposed as:

[19] Note that a person appears once in the sample for each job he holds over the sample period, and each such occurrence is denoted by a different value of *j*. Thus, the components of X_{jt} may include descriptive information on the person, the job, or both.

[20] Models of job-matching are discussed in Jovanovic (1979) and Topel (1986).

$$\epsilon_{jt} = \gamma_{jt} + \eta_{jt} \tag{6}$$
$$\gamma_{jt} = \rho\,\gamma_{jt-1} + \nu_{jt}$$

where ν_{jt} is an autoregressive shock to earnings and η_{jt} represents a transitory shock or measurement error in the earnings data. Their joint distribution is given by:

$$\left\{\eta_{jt}, \nu_{jt}\right\} \sim \textit{i.i.d.}\ (0,\Sigma) \tag{7}$$
$$\Sigma = \begin{bmatrix} \sigma_\eta^{\,2} & 0 \\ 0 & \sigma_\nu^{\,2} \end{bmatrix}$$

Equations (6) and (7) define a set of parameters $\Gamma = (\rho,\sigma_\nu,\sigma_\eta)$ that characterize the stochastic component of within-job wage evolution of interest. Equation (7) imposes the identifying assumption that $\{\eta_{jt}\}$ and $\{\nu_{jt}\}$ are mutually uncorrelated at all leads and lags.

The autoregressive shock ν_{jt} is the result of economic events that will affect the worker's wage earning capacity over several years and can be introduced at a variety of levels. Note that if $\rho = 1$, then the effect of this shock will never dissipate; when this random walk is imposed, ν_{jt} will be referred to as a permanent shock. At the individual level, a persistent shock to earnings may come in the form of a disability that limits the tasks the worker can perform. Within a firm, "tournaments" that determine promotions among similar workers will contain a random component. That is, uncertainty about career advancement within a firm is wage uncertainty in this context. Unforeseen changes in product demand can affect the productivity and, consequently, the wages of all workers in a given industry. Similarly, technological advances and changes in factor demand can permanently alter the returns to the skills of a given occupation in all industries.

The transitory shock η_{jt} is the result of short term variations in earnings. An example of such a shock is layoffs that last only a year or so. Illness or emergency that temporarily limits the hours available to work will also be transitory shocks to earnings. Fluctuations in product demand due to the business cycle may result in more or less overtime being available during a given year. For workers whose wages are negotiated at intervals of a few years, transitory shocks may still arise from unexpected inflation over the contract period or temporary give-backs to the firm. Measurement error in the dataset will also contribute to the measure of the transitory shock to earnings.

This decomposition of ϵ_{jt} into persistent and transitory shocks is useful for two reasons. First, as described in Section III, pension formulas embody long term contracts between the firm and the worker. Most take an average over three or more years to compute the final average pay, thereby mitigating the effect of purely transitory shocks on the level of the pension benefit. While the magnitude of transitory shocks has been shown to be economically important in other respects, pension benefits in general are not sensitive to them.[21] Second, there have been few systematic attempts to quantify how much uncertainty different groups in the working population have in their income processes. For example, it is commonplace in econometric work to simply assert, as in Skinner (1988), that farmers and self-employed workers have more income uncertainty than other groups. Empirical estimates in Carroll and Samwick (1992) show, however,

[21]Carroll and Samwick (1992) show that transitory shocks to total *noncapital* income have statistically significant effects on the household wealth levels. The explanation for this difference is that transitory shocks to noncapital income are measured to be much larger than those to wages at a single employer. A secondary explanation is that saving behavior is better characterized than labor supply by high discount rates, thereby magnifying the importance of transitory, relative to permanent, income fluctuations.

that this perception is largely due to much higher transitory rather than permanent shocks to income for these groups. Thus, considering the two stochastic processes separately helps to clarify which groups in the population are facing more uncertainty that could be compensated for in a long term contract such as a pension.

Wage Uncertainty at a Single Firm

The first step in forming an estimate of Γ is to remove the match-specific component of the error θ_j in equation (6) from the specification. One way to accomplish this is to difference the equation at any lag k:

$$(8) \qquad \begin{aligned} w_{jt} - w_{jt-k} &= \left(X_{jt} - X_{jt-k}\right)\beta + \left(\epsilon_{jt} - \epsilon_{jt-k}\right) \\ \Delta w_{jt}{}^{k} &= \Delta X_{jt}{}^{k}\beta + \Delta \epsilon_{jt}{}^{k} \end{aligned}$$

where the superscript k denotes the length of the lag. Equation (8) can be consistently estimated by OLS under the appropriate exogeneity condition for the regressors:

$$(9) \qquad E\left(X_{jt} - X_{jt-k}\right)\cdot\left(\epsilon_{jt} - \epsilon_{jt-k}\right) = 0$$

Since the parameters of the model pertain to the variance of the shocks to earnings, a flexible strategy for estimating Γ is to use the Generalized Method of Moments (GMM) on second degree functions of the residuals from equations like (8).

One set of moment conditions are those of Topel and Ward (1992), who examine wage growth and job mobility for young workers. They use the covariagram of $\Delta\epsilon_{jt}{}^{1}$:

$$C_0 \equiv E\left[(\epsilon_{jt} - \epsilon_{jt-1})\cdot(\epsilon_{jt} - \epsilon_{jt-1})\right] = \frac{2\sigma_v^2}{(1+\rho)} + 2\sigma_\eta^2$$

(10) $$C_1 \equiv E\left[(\epsilon_{jt} - \epsilon_{jt-1})\cdot(\epsilon_{jt-1} - \epsilon_{jt-2})\right] = -\frac{\sigma_v^2(1-\rho)}{(1+\rho)} - \sigma_\eta^2$$

$$C_k \equiv E\left[(\epsilon_{jt} - \epsilon_{jt-1})\cdot(\epsilon_{jt-k} - \epsilon_{jt-k-1})\right] = -\frac{\sigma_v^2\rho^{k-1}(1-\rho)}{(1+\rho)}, \quad \forall k \geq 2$$

The final expressions in the moment conditions for this model are denoted by $F_C(\Gamma)$. Given consistent estimates of equation (8), the asymptotic distribution of the parameters is given by:

(11) $$\hat{C} - F_C(\Gamma) \sim N(0, \Omega_C)$$

and the GMM estimate, Γ_C^*, minimizes the quadratic form:

(12) $$S_C = \left[\hat{C} - F_C(\Gamma)\right]' \hat{\Omega}_C^{-1} \left[\hat{C} - F_C(\Gamma)\right]$$

with respect to Γ, where the elements of $\hat{C} = (\hat{C}_0, \hat{C}_1, \ldots, \hat{C}_k)'$ are the sample analogues to the theoretical covariances depicted on the left side of equation (10).[22]

The next step in the estimation of Γ is therefore to estimate equation (8) with k = 1. Because the only nonwage data included in the personnel records are age, tenure, and calendar year, the column-space of X_{jt} in equation (6) includes Age, Age^2, Tenure,

[22]The ^ over the covariance matrix in the equation indicates that a two-step estimator is used, with the first step using the identity matrix for Ω and the second step using the residuals from the first step to construct an asymptotically efficient estimate of Ω.

Tenure^2, and indicator variables for each year except the first.[23] Equation (5) is estimated in Table 2. The marginal effect of Age and Tenure on the wage are positive through 42 and 23, respectively, after which they are negative. Equation (8) with k = 1 is estimated in Table 3. As in Table 2, both Age and Tenure are significant predictors of the wage level.

Finally, the residuals from equation (8) are used to construct $\hat{C}$. Table 4 presents the GMM estimates for Γ using $\hat{C} = (\hat{C}_0,\hat{C}_1,\hat{C}_2,\hat{C}_3,\hat{C}_4)'$.[24] The first row of parameter estimates pertains to the full sample of 4001 individuals used in the estimation of β. The estimate for ρ is very close to unity, which would imply that ν_{jt} is a random walk or permanent shock to income, although this hypothesis can be statistically rejected. When such a restriction is imposed on the data, the estimates for σ_ν and σ_η are hardly affected, but the χ^2-test for misspecification from Newey (1985) significantly rejects the null hypothesis of no misspecification.

The estimates for σ_ν are highly significant but unusually large. The next set of estimates shows that the results depend on whether individuals who experience an extremely large negative shock to income are included in the sample. If individuals who had the ratio of actual to predicted (from equation (5)) income in one or more years fall to below 20 percent of the average of this ratio over the sample period are excluded, the estimates for σ_ν fall from 0.24 to 0.15, a figure which is more consistent with Topel and

[23]Including higher order terms for Age and Tenure has minimal impact on the estimates of Γ and no effect on the qualitative results for either the single firm or PSID dataset.

[24]Γ_C^* was estimated using the generalized method of moments algorithm of Hansen, Heaton, and Ogaki documented in Ogaki (1992). Financial support for that research was given by the National Science Foundation.

Ward's (1992) estimate of 0.13. Carroll (1991) argues that these "near-zero" income events are better modeled as a separate transitory income process.[25] The bottom two rows of the table show that increasing the near-zero threshold from 0.2 to 0.4 has comparatively little effect on the parameter estimates.

Using the income threshold of 0.2 as a baseline, then, the GMM estimates for the single firm show that the standard deviation of the permanent shock to earnings is approximately 15 percent in this sample. If the shocks were lognormally distributed, this would imply that two thirds of the workers would have an income realization between 0.85 and 1.15 times their one-year ahead income forecast each year and 95 percent would have a realization between 0.70 and 1.30 times this predicted income. The estimated standard deviation of the transitory shock (which exists for one year only and compounds the permanent shock that period) is 5.53 percent, which is much lower than Topel and Ward's 11 percent. One explanation is that personnel records are less likely to contain measurement error than even a carefully measured panel dataset. Another is that this particular sample nearing retirement itself induces less variation in hours worked over the year than the younger workers in that study.

Carroll and Samwick (1992) present an estimation method that imposes stronger exogeneity restrictions on the model in equation (5) but allows for more robustness to

[25]Since the focus here is on permanent rather than transitory shocks within a given firm, the approach taken will be to further select the sample so that such observations do not affect the estimates of Γ rather than formally modeling the "near-zero" process.

misspecification in equation (7). They use the residuals from equation (5) directly to construct moments based on the theoretical variances of $\Delta\epsilon_{jt}^{k}$:

$$(13) \qquad W_k \equiv E\left[\left(\epsilon_{jt} - \epsilon_{jt-k}\right)^2\right] = 2\cdot\sigma_\eta^{\,2} + \frac{2}{1+\rho}\cdot\left(\sum_{l=1}^{k}\rho^{l-1}\right)\cdot\sigma_v^{\,2}, \quad \forall\, k \geq 1$$

The condition for consistent estimates of β in equation (5) is now:

$$(14) \qquad E\left(X_{jt}\cdot\left(\epsilon_{jt} + \theta_j\right)\right) = 0$$

It is clear from equation (13), however, that all the parameters of Γ can be estimated using any values of k, whereas the method based on the covariagram of $\Delta\epsilon_{jt}^{1}$ required that C_0 and C_1 be included for σ_η to be estimated. By using only values of k greater than some level q, the estimates of equation (13) will be robust against serial correlation in the transitory shock up to order MA(q).[26] In particular, it can accommodate the MA(2) serial correlation documented by both MaCurdy (1982) and Abowd and Card (1989) in their studies of labor income.

Under the assumption that w_{jt} follows a random walk ($\rho = 1$), equation (13) reduces to a very convenient formula:

$$(15) \qquad W_k = 2\cdot\sigma_\eta^{\,2} + k\cdot\sigma_v^{\,2}, \quad \forall\, k \geq 1$$

which can be estimated by using OLS of a sample analogue of W_k on a vector of 2's and a vector of lag lengths; the coefficients of this regression are σ_η^2 and σ_ν^2, respectively. Table 5 presents the results of this regression for a variety of near-zero thresholds and MA(q) robustness assumptions. In general, this methodology yields a similar pattern of

[26] That is, $\mathrm{Cov}(\eta_{jt}, \eta_{jt-k}) = 0$ and, therefore, $\mathrm{Var}(\eta_{jt}-\eta_{jt-k}) = 2*\sigma_\eta^2$ for $k > q$, as in equation (13).

coefficients by near-zero threshold level. The values for σ_ν are lower and those for σ_η are higher than their counterparts in Table 4.

Allowing for robustness against MA(2) serial correlation slightly increases the estimate of σ_ν and reduces the estimate of σ_η. Both effects are less dramatic when the near-zero threshold is larger. In fact, when no observations are removed due to near-zero events and the procedure is made robust to MA(2) serial correlation, the coefficient estimated for σ_η^2 is negative. This is a feature of any estimation strategy for a variance which cannot constrain the estimate to be nonnegative. The estimate is still consistent, and for the purposes of comparing the relative uncertainty of different groups within the population, the negative estimate is simply an indication of a very small transitory shock to wages.

Wage Uncertainty in a Cross-Section

The two methods applied to the data from the single firm show that even in a conservatively chosen sample, wage uncertainty is estimated to be large. On average, workers in this firm experience a permanent shock to earnings with a standard deviation of 10 to 15 percent per year, compounded by a transitory shock of about 5 to 7 percent per year. These figures compare favorably to the average annual growth rate of 3.5 percent over the sample period. This firm, however, need not be representative of the population at large or the workforce covered by pensions. To verify wage uncertainty in the general population, the same two methods can be applied to the PSID. Since the

PSID is a more representative sample of the population, the estimates of uncertainty generated from it will apply broadly to other cross-sectional surveys such as the *SCF*.[27]

The sample is comprised of all persons who were heads of households between the ages of 25 and 55 at the beginning of the sample period and held at least one job over 1976-1987. Persons who were not heads of households are excluded because data collection for them is less consistent over the sample period. Because most pensions deal explicitly with wages on one particular job at a time, each year of wages had to be matched up with a particular employer. Additionally, because the focus is on employer-provided pensions, any jobs for which the respondent was self-employed are excluded. However, it should be noted that it is not possible to determine pension coverage in the PSID for all years, so estimates of the wage process cannot be made conditional on pension coverage.[28]

Since the PSID contains no employer identification number or direct question (in every wave) as to whether the respondent is working for a new employer this year, employer changes had to be identified using responses to the question: "How many years' experience do you have altogether with your present employer?" Presumably, a decrease in this variable from one wave to the next indicates a change in employers, but

[27]The PSID contains a representative national survey and a supplemental *Survey of Economic Opportunity* comprised exclusively of low-income households. All estimates below include the low-income households, but none of the qualitative results change when they are omitted.

[28]In some sectors such as construction and transportation, multi-employer pension plans are common. Multi-employer plans are portable across some firms within the same industry. In such cases, it might be better to treat every job covered by the same pension (regardless of employer) as a single job. This would be more relevant if the subsequent analysis focused on the uncertainty of tenure length rather than wages. Additionally, Internal Revenue Service Form 5500 data presented in the Appendix to Turner and Beller (1992) show that multi-employer plans account for only 13 percent of all pension plan participants and 10 percent of all pension assets and contributions.

implementing this approach is not so straightforward. For example, suppose the number of months decreases from 132 to 108 in one year. Does this represent recall error, a newly found confusion over "present employer" and "present job," or an actual transition from a job in which the respondent has 11 years of experience to a previously held job which he left after 8 years of experience?

Brown and Light (1992) investigate such problems thoroughly and conclude that econometric results will not be greatly affected by ambiguous responses like these as long as some method of internal consistency is imposed on the data.[29] That is, tenure must be anchored at some value and incremented by one for each year on the job. The first step in the approach followed here was to identify earnings histories that appeared to be the same job over the whole period, making allowances for observations such as the one above to be counted as the same job. Next, a group of observations who did not, over the twelve year period, have at least four years of consistent employment and tenure responses was discarded. The tenure variable was not reported for 1978 (if over age 45), 1979, and 1980. Of the remaining earnings histories, ambiguities about job starts in those years were resolved by using changes in occupation and industry to identify the year in which the job began if it is not clear from the tenure variables in 1977 and 1981.

---

[29]Studies of the effect of seniority on wages by Altonji and Shakotko (1987), Abraham and Farber (1987), and Topel (1991) also discuss methods of synthesizing a tenure series from the variables reported in the PSID for various years.

The initial response for tenure is taken to be valid, and the values for other years are constrained to increase by one per year.[30]

This procedure identifies 2,171 jobs encompassing 15,554 years of wages. The measure of wages used is the log of wage and salary income. X_{jt} in equation (5) augments the specification for the single firm by interacting {1, age, age^2} with indicator variables for (six) education levels, (seven) occupation categories, and (seven) industry groups. Demographic variables for gender, race, marital status, union status, census region, and public sector employment are also included. The R^2 is 0.5163 for this specification. Although it is quite likely the case that unobserved variables are correlated with some elements of X_{jt} (tenure, for example), the idea here is to try to find the forecast of wages that most closely approximates that of the worker. The inherent tradeoff is between the bias in the coefficients in the regression itself and overestimating the true amount of uncertainty. Since the coefficients themselves are not of primary interest, emphasis is placed on guarding against the latter.[31] As in the previous sample, the estimate of equation (8) simply estimates equation (5) in first differences and has an R^2 of 0.0128.

Table 6 presents the estimates of Γ_C^* for the PSID sample under the same restrictions as in Table 4. The results are far less sensitive to the near-zero threshold imposed. The point estimates of ρ are lower than those in Table 4, but the null

---

[30]A complete description of the algorithm is available on request. A non-scientific examination of the resulting tenure sequences suggested that about 5 percent of the individual job histories still have some ambiguities.

[31]Excluding all variables that may be related to the job reduces the R^2 to 0.28 but does not substantially affect the results in Tables 6-8.

hypothesis of a random walk cannot be rejected at the 5 percent significance level. The imposition of this restriction leads to insignificant changes in the estimates for σ_ν and σ_η, and the χ^2-test for misspecification cannot reject at conventional significance levels. Compared to the single firm data, the permanent shock to earnings is slightly smaller (0.1338 instead of 0.1525) while the transitory shock considerably larger in the PSID data. Some of the latter effect is undoubtedly due to greater measurement error in the household survey data than in the personnel records.

The results of the variance decomposition applied to the PSID sample are shown in Table 7. The pattern of coefficients by near-zero threshold is similar to that in Table 5, but both types of shock are estimated to be larger in this data than in the sample of salesmen. The estimates of σ_ν are very similar to the GMM estimates in Table 6 whenever some positive near-zero threshold is imposed. The consequences of making the procedure robust to MA(2) serial correlation are more apparent in the PSID. Compared to the GMM estimates for σ_η, the variance decomposition yields a higher value when the procedure is not robust to MA(2) and a lower value when it is. For the single firm, the variance decomposition always gave a higher estimate for σ_η than the GMM estimate. Overall, the two methods suggest a value of about 13.5 percent for the permanent shock and 12 percent for the transitory shock.

The close correspondence between the results in Tables 6 and 7 and those in Tables 4 and 5 suggests that within-job wage uncertainty is reliably estimated in the PSID. Both methods give estimates for permanent uncertainty close to those found using better data in the single firm. As expected, transitory uncertainty is found to be larger, but this can be attributed in part to the different sample selection criteria and to the greater degree of measurement error in a cross-sectional survey.

Since each method of estimating the magnitude of shocks to income is based on aggregating the estimates at the individual level, it is also possible to estimate the parameters on subsamples of the PSID. The top panel of Table 8 presents the mean estimates of (σ_ν,σ_η) for workers of selected educational attainment using the variance decomposition method. The magnitude of the permanent shock declines with education, with the major difference coming between those with a high school diploma and those without one. One perquisite of the diploma is that it gains access to jobs with more secure wage paths. Those with the added education of a college degree do not have smaller permanent shocks to income; rather, the contribution of the degree to lower uncertainty is manifested entirely in smaller transitory shocks to income.

The mean estimates of the uncertainty parameters by occupation are given in the middle panel of Table 8. All estimates of the permanent shock are statistically different from zero. Managers have the largest estimate and craftsmen have the smallest. Other occupations have estimates at roughly the sample average values. The magnitudes of the transitory shocks do show more variation across occupation, with the entirely white collar occupation of professionals and managers having the least and service workers having the most.

Lastly, uncertainty estimates by industry are given in the bottom panel of the table. The values for manufacturing industries, which comprise one third of the sample, are quite close to the full sample estimates in Table 7. Workers in wholesale and retail trade have much larger permanent shocks and much smaller transitory shocks by comparison. The situation is exactly the reverse for workers in finance, insurance, and real estate, who have relatively low estimates of permanent uncertainty and relatively high estimates of transitory uncertainty. Public administrators have lower estimates for both types of uncertainty than do workers in manufacturing industries.

The estimates in Tables 8 show that for all but a few subsamples, permanent uncertainty is statistically different from zero and generally quite large. There are also pronounced differences between at least some of each type of group, with industrial groups showing the most variation. It is natural to observe within-job movements in wages to be more a function of industry than occupation, as the strongest underlying economic movements that generate uncertainty probably operate in the markets for goods rather than those for factors of production.

V. Empirical Evidence on Risk Sharing

Given the presence of risk-compensating pension plan features described in Section III and the considerable wage uncertainty estimated in Section IV, the next step is to determine whether they are systematically related. If pension formulas are in fact designed to more efficiently share the burden income risk between firms and workers, then we would expect the pension plans of workers with higher income risk to compensate them for it directly with convex pension formulas or insure retirement income against it with Flat Rate pensions. Conversely, we would expect pension plans which compensate for or mitigate earnings uncertainty to support a higher level of uncertainty.

The approach to testing this hypothesis begins with simulating pension benefits under a uniform level of uncertainty for each pension in the dataset. From these simulations, the sample analogues to moments of each pension's distribution can be constructed. The two moments of interest are the expected value and coefficient of variation of the pension benefits. The former, when compared to the pension benefits when there is no wage uncertainty, will describe the convexity of the pension plan formula. The latter will determine how close the pension formula is to the Flat Rate pension that completely insures retirement income against pre-retirement wage uncertainty. Each worker in the SCF can then be matched up with his or her pension moments from the PPS and within-job wage uncertainty from the PSID. If the sample correlation of these two variables is such that persons with higher income uncertainty also have pensions with higher expected benefits relative to the certainty case and lower coefficients of variation of pension benefits, then the hypothesis is not rejected.

There are 1,485 individuals in the SCF that are covered by one or more of the 924 pensions in the PPS summarized in Table 1. The mean characteristics of this sample relevant for pension calculations are given in Table 9. The 34 million workers represented in this sample comprise roughly 75 percent of the workforce covered by pensions in 1983. The remainder are excluded because the entity responsible for providing the pension (usually the employer) could not be contacted by the survey.[32]

The annual shocks to earnings to be used in the calculation of the pension benefit distributions were drawn from lognormal distributions with standard deviations of {0.05,

[32]The workers from the SCF and the pensions from the PPS are the same as those used in Samwick (1993).

0.10, 0.15, 0.20} in the underlying normal distribution.[33] This set bounds the values found in the empirical estimates in Section IV for the PSID. A uniform value of uncertainty is used for each worker rather than the actual (predicted) uncertainty to avoid spurious correlation between the pension moments and the level of uncertainty. For example, if actual uncertainty were used and everyone had the same pension plan formula, then the expected pension value ratio would be highly--if not perfectly--correlated with the level of uncertainty even though no meaningful differences in risk compensation existed. Because transitory shocks are far less important in determining the level of pension benefits, the transitory shocks were not explicitly modeled.

For each plan and value of σ_η, the entitlements at age 65 for an individual with the mean (or modal, if an indicator variable) value of the characteristics in Table 9 were calculated under 100 different earnings paths.[34] As with the level of uncertainty, a representative individual is used in these calculations, rather than the actual individual, so that personal decisions about how much to contribute to the pension, how many hours to work, and when to join the firm do not induce a spurious correlation between the pension plan moments and the amount of uncertainty.

For each earnings path, the amount of average annual earnings and the initial pension benefit were computed. From the set of earnings paths, the sample moments of the

[33]The means of the underlying normal distributions were adjusted to account for the effect of the standard deviation of the underlying normal distribution on the mean of the actual distribution. That is, the mean value of the multiplicative shock to income is 1.0, regardless of the parameter for the underlying standard deviation. This allows an increase in the standard deviation of the underlying normal distribution to represent a mean-preserving spread to the distribution of wages.

[34]Without loss of generality, it is assumed that all pensions are paid in the form of joint life annuities.

distribution of pension benefits were computed for each pension. As a benchmark, the predicted value of the pension and average wages were also computed under an assumption of no uncertainty.

Table 10 contains descriptive statistics for these pension moments for each possible value of σ_η. The top panel summarizes the ratio of E[p(w)] under the given value of σ_η to p(w) when $\sigma_\eta = 0$. All of the summary statistics are increasing with the level of σ_η. The mean value of 1.0576 when $\sigma_\eta = 0.15$ indicates that expected pension benefits are an average of 5.76 percent higher when the variance of the permanent shock to earnings is near its estimated value in the PSID rather than zero. It is also noteworthy that the mean ratio for $\sigma_\eta = 0.05$ is actually lower than 1, indicating that for low amounts of uncertainty, pensions punish riskier earnings path on average. This suggests that pension plan formulas may be locally concave but globally convex around the mean level of wages in the sample.

There is also considerable variation in this ratio. The minimum ratio observed in the sample is 0.05 and the maximum is 5.59, indicating the presence of extremely concave and convex pension formulas, respectively. Figure 5 graphs the cumulative distribution function for this ratio with the top and bottom two percent of the distribution omitted (from the graph, not the CDF itself). The large increase at unity is the result of Flat Rate, Percent of Earnings, and Defined Contribution plans that are by definition invariant in expected value to wage uncertainty. Only 10 percent of the plans are actually concave in the underlying wage (the ratio is strictly less than 1), whereas nearly two thirds are convex (the ratio is strictly greater than 1). Finally, over half the plans

Figure 5

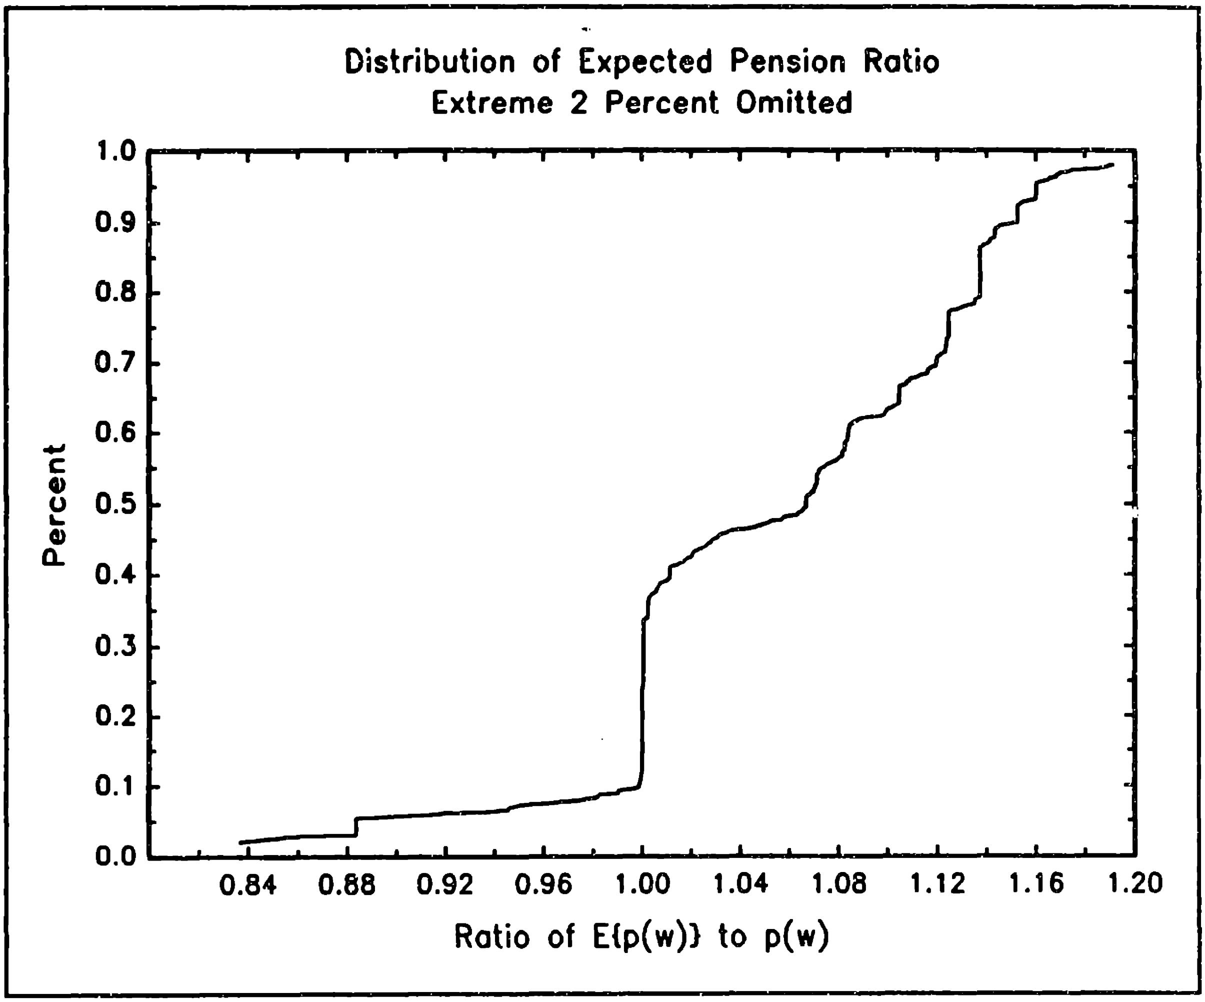

have ratios in excess of the mean value of 1.0576.

The bottom panel of Table 10 summarizes the coefficient of variation of p(w) for each value of σ_η. The mean value rises with the level of uncertainty. A value of zero generally indicates a Flat Rate plan which is independent of actual earnings received.[35] The intercept of the cumulative distribution function in Figure 6 indicates that approximately 12 percent of the plans are Flat Rate. The distribution is very flat over most of its range, indicating a fairly dispersed distribution of the coefficient of variation.

[35]It may also indicate a plan which takes the minimum or maximum of an earnings-related and a flat rate formula in which the flat rate formula is locally determining benefits.

Figure 6

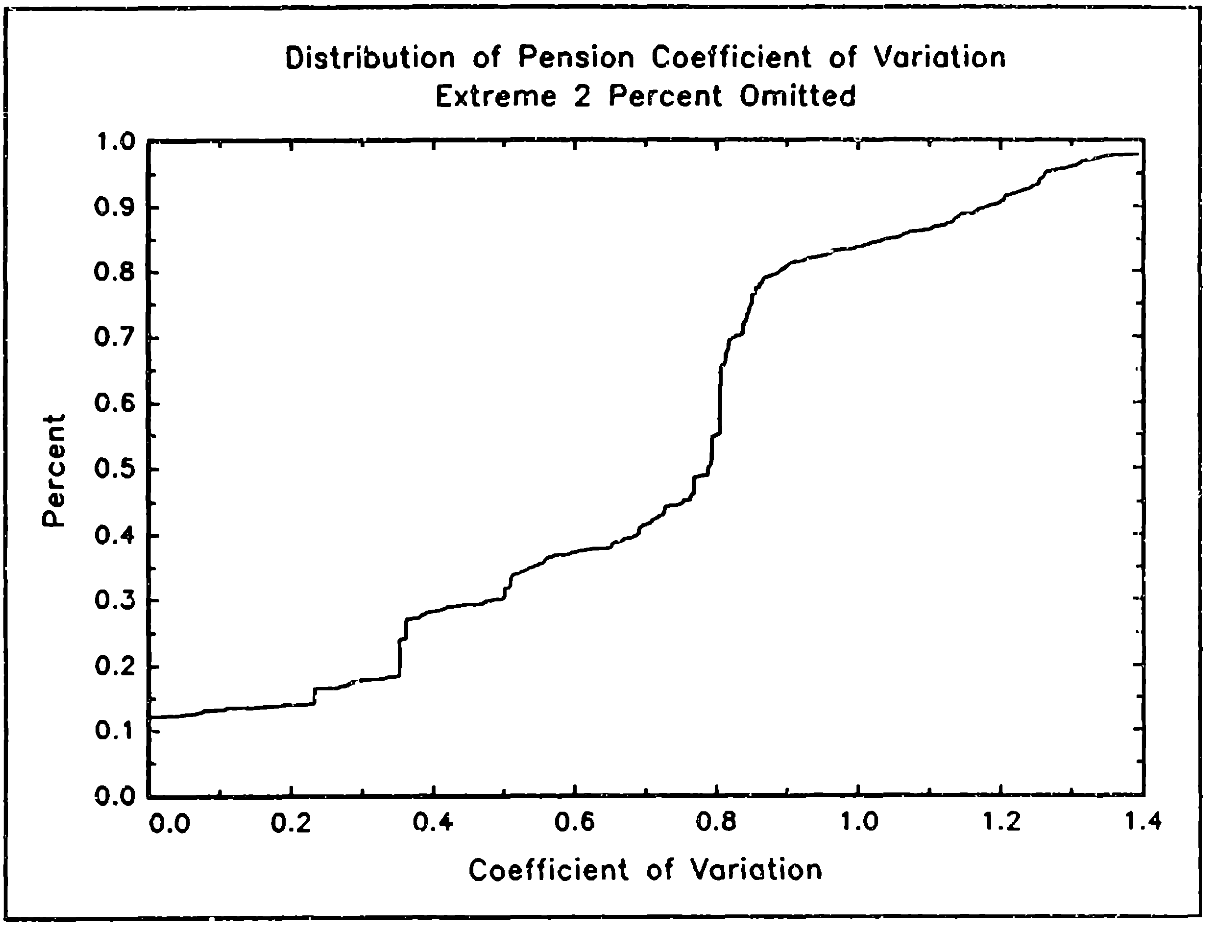

One way to test the hypothesis that pensions are serving to compensate for the uncertainty in wages or insulate retirement income from it is to compute the sample correlation of the relevant pension moment and the magnitude of the income uncertainty. Equivalently, the following regressions could be estimated:

$$\text{(16)} \qquad \begin{aligned} \textit{Pension Ratio} &= \beta_0 + \beta_1 \sigma_\eta^2 + \epsilon_1 \\ \textit{Pension C.V.} &= \gamma_0 + \gamma_1 \sigma_\eta^2 + \epsilon_2 \end{aligned}$$

If pensions are being used to compensate workers for income uncertainty, then workers with higher income uncertainty should have pensions that increase relatively more in expected value in the presence of uncertainty. Similarly, if pension formulas are serving to insulate retirement income from uncertain wages, then workers with higher income uncertainty should have pensions with lower coefficients of variation for a given level of uncertainty. That is, β_1 should be positive and γ_1 should be negative. Of course, employers may be using other methods, such as the average level of wages themselves, to compensate for higher uncertainty even if these coefficients are insignificantly different from zero. The hypothesis being tested here is only that *some* amount of the uncertainty documented in Section IV is being offset by pensions with better risk characteristics.

In order to estimate equation (16), each worker must be assigned a value for the magnitude of the permanent shock to income. Because the estimates in Section IV are based on data in the PSID rather than the SCF, the individual estimates for the variance of the permanent shock to earnings are projected onto a set of variables including Age, Age^2, Tenure, Tenure^2, and indicator variables for occupation, industry, education, region, sex, race, marital status, public sector, and union status. The R^2 of this regression was 0.055 and the standard error of the estimate was 0.173. Since all of these variables are also present in the SCF, the predicted values of $\sigma_\eta{}^2$ can be constructed for each worker.

Table 11 presents the estimates of equation (16) using the pension moments calculated for $\sigma_\eta = 0.15$ and $\sigma_\eta = 0.10$.[36] Separate results are presented for private and public sector workers, union and non-union workers, and the four groups comprised

[36]The estimates in Table 11 are corrected for the attenuation of the coefficients toward zero in the presence of measurement error which results from using the predicted value for uncertainty rather than the actual value. The correction is generally an upward revision of no more than 1 percent of the OLS estimate.

of the interaction of these two categories. For each subgroup and each value of σ_η, the estimate of β_1 is always positive, indicating that the simple correlation of wage uncertainty and pension plan convexity is positive. Thus, in each sector, those workers who face the greatest income uncertainty have pensions which compensate them relatively more in the presence of uncertainty. This basic result demonstrates the presence of wage risk compensation through private pensions.

In most cases, however, the parameter estimates are not significantly different from zero. One explanation of this finding is that the construction of both the pension moments and the measure of uncertainty removed much of their individual variation in order to safeguard against endogeneity in the relationship. Regarding the ratio of expected pension benefits to pension benefits without uncertainty, using uniform values of wages and σ_η may be more than what is absolutely necessary to purge the relation of spurious correlations described above. The consequence may be that the dependent variable describes the compensation for risk at a wage level or value of σ_η which will never apply to the worker in question. We would naturally expect such a relationship to be less apparent than the true relationship. Similarly, the R^2 of 0.055 on the "first stage" regression of uncertainty on personal characteristics in the PSID indicates that much of the variation in the actual uncertainty faced by workers is not present in the predicted values in the SCF. Although this is the desirable outcome of instrumenting, a better set of exogenous variables to use for the first stage regression might improve the statistical significance in the estimation of equation (16).

An alternate explanation is that the risk compensation through pensions occurs only within some of the labor markets in the sample. For the estimates with $\sigma_\eta = 0.15$, β_1 achieves its highest magnitude and significance for the private sector, unionized workers, and union workers in only the private sector. This contrasts with the magnitude of the estimate for β_0, which is higher in the public and non-union sectors. This pattern suggests that although the amount of risk-compensation per se is higher for non-union workers and public sector workers, it is not necessarily higher within those sectors based on the amount of uncertainty faced by the workers. The estimates with $\sigma_\eta = 0.10$ are similar in nature, with a weaker result regarding the effect of union status.

The estimates for γ_1, the correlation between wage uncertainty and the coefficient of variation in pension benefits, are not as robust. Although the estimate is never significantly greater than zero for either value of σ_η, it is now negative *only* when the sample is restricted to private sector workers, union workers, or both. As the estimates for γ_0 show, the coefficient of variation of pension benefits is on average lower in the private and unionized sectors. Undoubtedly this is due to a more widespread use of flat rate formulas in these sectors. The results for γ_1 demonstrate that within these sectors, workers covered by less variable pensions tend to be the ones facing the highest income uncertainty. In these sectors, pensions are also offsetting the harmful effects of wage uncertainty by offering less variable pension benefits to those with more variable wages.

VI. Conclusion

This paper has documented the presence of considerable within-job earnings uncertainty and the effects of its interaction with nonlinear pension benefit formulas. It shows that workers with higher wage uncertainty have pensions that disproportionately compensate them for it. The correlation between wage uncertainty and compensation for it through the pensions is most significant for workers within the private and unionized sectors of the economy. Within those sectors, there is also evidence that workers with higher wage uncertainty have pension formulas that, *ceteris paribus*, generate less variable pension benefit distributions.

These findings partially refute the assertion made by Eaton and Rosen (1980a, 1980b, 1980c) and Varian (1980) that there are no private arrangements for insuring the returns to human capital. Although no formal markets exist, some private institutions do. Therefore, their argument for progressive income taxation as a form of insurance is weakened. What is of greater importance in light of the results presented here is to ascertain the consequences of simultaneous government insurance of risk through the tax and transfer system and private compensation for risk through pension plans.

These findings also imply that optimal regulation of pension plans and optimal design of Social Security policy must take account of differences in pension coverage and risk insurance through pensions. The proliferation of pensions that are integrated with Social Security suggests that firms are already aware of at least some of this interaction. In particular, the welfare improving effects of the Flat Rate pension's low variance of pension benefits and the nonlinearity in integrated plan formulas which acts like a call on wages with an exercise price determined by the Social Security benefit or MTE might be incorporated into Social Security for the benefit of the half of the workforce that has no pension to compensate it for earnings uncertainty.

The results also have implications for studying precautionary saving behavior. The model in Section II showed that simple POE or DC plans imply that pre- and post-retirement income will be highly correlated and that pension benefits will inherit the uncertainty of the wage process, so that saving through the pension plan is an imperfect substitute for precautionary saving the individual would otherwise have to do. Flat Rate plans do not compensate for earnings risk, but they also guarantee a minimum level of post-retirement income independent of wages actually received over the career, which make them a better substitute for precautionary saving. Thus, empirical estimates of the amount of precautionary saving in Carroll and Samwick (1993) may find another source of identification in the type of pension plan the workers have.

Finally, the explicit accounting for the effect of uncertainty on the value of pension entitlements is not a feature of studies that examine the distribution of pension entitlements in the population (McDermed, Clark, and Allen 1989) or those that seek to use the pension to estimate compensating differentials in wage contracts (Montgomery, Shaw, and Benedict 1992). The part of the pension that resembles a call option should be evaluated under the amount of uncertainty the workers face and included in estimates of the net worth embodied in pensions. The stylized estimates presented here suggest that the risk-adjusted present value of pension benefits is approximately 6 percent higher. Given the wide variation in uncertainty in the population, however, the consequences could be far worse if differing pension present values are used to achieve identification.

Appendix

Proof of the lemma:

Let F(w,r) denote an arbitrary distribution of w. Increases in r will be mean-preserving spreads to the distribution of w if conditions (A.1) and (A.2) hold.

$$\text{(A.1)} \qquad \int_{\underline{w}}^{\overline{w}} F_r(w,r)\,dw = 0$$

$$\text{(A.2)} \qquad \int_{\underline{w}}^{w} F_r(z,r)\,dz \geq 0 \quad \textit{for } \underline{w} \leq w \leq \overline{w}$$

The expected value of the pension is given by the equation:

$$\text{(A.3)} \qquad E[p(w)] = \int_{\underline{w}}^{\overline{w}} p(w) F_w(w,r)\,dw$$

The derivative of the expected value with respect to r is:

$$\text{(A.4)} \qquad \frac{dE[p(w)]}{dr} = \int_{\underline{w}}^{\overline{w}} p(w) F_{wr}(w,r)\,dw$$

Integrating the right side by parts once yields:

$$\text{(A.5)} \qquad \frac{dE[p(w)]}{dr} = F_r(w,r)p(w)\Big|_{\underline{w}}^{\overline{w}} - \int_{\underline{w}}^{\overline{w}} p'(w) F_r(w,r)\,dw$$

The first term on the right side of the equation is necessarily zero because $F_r(w,r) = 0$ at both the upper and lower limits of a CDF for w. Integrating the second term on the right side of the equation by parts gives:

$$\text{(A.6)} \qquad \frac{dE[p(w)]}{dr} = -p'(w)\cdot\left[\int_{\underline{w}}^{w} F_r(z,r)\,dz\right]\Bigg|_{\underline{w}}^{\overline{w}} + \int_{\underline{w}}^{\overline{w}} p''(w)\cdot\left[\int_{\underline{w}}^{w} F_r(z,r)\,dz\right] dw$$

The first term on the right side of the equation is zero by condition (A.1). This leaves the equality:

$$\text{(A.7)} \qquad \frac{dE[p(w)]}{dr} = \int_{\underline{w}}^{\overline{w}} p''(w)\cdot\left[\int_{\underline{w}}^{w} F_r(z,r)dz\right]dw$$

Next subtract off the same relation for a different pension, q(w):

$$\text{(A.8)} \qquad \frac{dE[p(w)]}{dr} - \frac{dE[q(w)]}{dr} = \int_{\underline{w}}^{\overline{w}} (p''(w)-q''(w))\cdot\left[\int_{\underline{w}}^{w} F_r(z,r)dz\right]dw$$

The term in brackets on the right hand side is always nonnegative under (A.2). If p''(w) > q''(w) for all values of w, then the right side is nonnegative as well. This proves the result.

References

Abraham, Katharine G. and Henry S. Farber (1987). "Job Duration, Seniority, and Earnings," *American Economic Review*, Vol. 77, No. 3, 278-297.

Altonji, Joseph G. and Robert A. Shakotko (1987). "Do Wages Rise with Job Seniority?" *Review of Economic Studies*, Vol. 54, 437-459.

Abowd, John M. and David Card (1989). "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, Vol. 57, No. 2, 411-445.

Brown, James N. and Audrey Light (1992). "Interpreting Panel Data on Job Tenure," *Journal of Labor Economics*, Vol. 10, No. 3, 219-257.

Carroll, Christopher D. (1991). "Buffer Stock Saving and the Permanent Income Hypothesis," Board of Governors of the Federal Reserve System, Economic Activity Section Working Paper No. 114.

Carroll, Christopher D. and Andrew A. Samwick (1992). "The Nature and Magnitude of Precautionary Wealth," Board of Governors of the Federal Reserve System, Economic Activity Section Working Paper No. 124 (slightly revised as Chapter One of this thesis).

Carroll, Christopher D. and Andrew A. Samwick (1993). "How Important Is Precautionary Saving?" mimeo, Massachusetts Institute of Technology (also Chapter Two of this thesis).

Eaton, Jonathan and Harvey S. Rosen (1980a). "Labor Supply, Uncertainty, and Efficient Taxation," *Journal of Public Economics*, Vol. 14,365-374.

Eaton, Jonathan and Harvey S. Rosen (1980b). "Taxation, Human Capital, and Uncertainty," *American Economic Review*, Vol. 70, 705-715.

Eaton, Jonathan and Harvey S. Rosen (1980c). "Optimal Redistributive Taxation and Uncertainty," *Quarterly Journal of Economics*, Vol. 95, 357-364.

Jovanovic, Boyan (1979). "Job Matching and the Theory of Turnover," *Journal of Political Economy*, Vol. 87, 972-990.

Kaplow, Louis (1991a). "Incentives and Government Relief for Risk," *Journal of Risk and Uncertainty*, Vol. 4, 167-175.

Kaplow, Louis (1991b). "A Note on Taxation as Social Insurance for Uncertain Labor Income," National Bureau of Economic Research Working Paper No. 3708.

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

Kotlikoff, Laurence J. and David A. Wise (1989). "Employee Retirement and a Firm's Pension Plan," in David A. Wise (ed.) *The Economics of Aging*. Chicago: University of Chicago Press.

Lazear, Edward P. (1979). "Why Is There Mandatory Retirement?" *Journal of Political Economy*, Vol. 87, No. 6, 1261-1284.

Lazear, Edward P. (1983). "Pensions as Severance Pay," in Zvi Bodie and John B. Shoven (eds.) *Financial Aspects of the United States Pension System*. Chicago: University of Chicago Press.

Lillard, Lee A. and Robert J. Willis (1978). "Dynamic Aspects of Earning Mobility," *Econometrica*, Vol. 46, No. 5, 985-1012.

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

McDermed, Ann A., Robert L. Clark, and Steven G. Allen (1989). "Pension Wealth, Age-Wealth Profiles, and the Distribution of Net Worth," in Robert E. Lipsey and Helen Stone Tice (eds.) *The Measurement of Saving, Investment, and Wealth*. Chicago: University of Chicago Press.

Merton, Robert C., Zvi Bodie, and Alan J. Marcus (1987). "Pension Plan Integration as Insurance Against Social Security Risk," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.) *Issues in Pension Economics*. Chicago: University of Chicago Press.

Mitchell, Olivia S. (1992). "Trends in Pension Benefit Formulas and Retirement Provisions," in John A. Turner and Daniel J. Beller (eds.) *Trends in Pensions 1992* Washington, DC: United States Government Printing Office.

Montgomery, Edward, Kathryn Shaw, and Mary Ellen Benedict (1992). "Pensions and Wages: An Hedonic Price Theory Approach," *International Economic Review*, Vol. 33, No. 1, 111-128.

Newey, Whitney K. (1985). "Generalized Method of Moments Specification Testing," *Journal of Econometrics*, Vol. 29, 229-256.

Ogaki, Masao (1992). "GMM: A User's Guide," mimeo, University of Rochester.

Ransom, Roger L., Richard Sutch, and Samuel H. Williamson (1992). "Inventing Pensions: The Origins of the Company-Provided Pension in the United States, 1900-1940," mimeo, University of California, Berkeley.

Rothschild, Michael and Stiglitz (1970). "Increasing Risk I: A Definition," *Journal of Economic Theory*, Vol. 2, 225-243.

Samwick, Andrew A. (1993). "Retirement Incentives in the 1983 Pension Provider Survey," mimeo, Massachusetts Institute of Technology (also Chapter Three of this thesis).

Skinner, Jonathan (1988). "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics*, Vol. 22, 237-255.

Topel, Robert H. (1986). "Job Mobility, Search, and Earnings Growth: A Reinterpretation of Human Capital Earnings Functions," in *Research in Labor Economics*, Vol. 8, Part A (Greenwich, CT: JAI Press, Inc.), 199-233.

Topel, Robert H. (1991). "Specific Capital, Mobility, and Wages: Wages Rise with Seniority," *Journal of Political Economy*, Vol. 99, No. 1, 145-176.

Topel, Robert H. and Michael P. Ward (1992). "Job Mobility and the Careers of Young Men," *Quarterly Journal of Economics*, Vol. 107, No. 2, 439-479.

Turner, John A. and Daniel J. Beller (eds.) (1992). *Trends in Pensions 1992*. Washington, DC: United States Government Printing Office.

Varian, Hal R. (1980). "Redistributive Taxation and Social Insurance," *Journal of Public Economics*, Vol. 14, 49-68.

Table 1

Characteristics of Plans in the PPS		
Number of Normal Retirement Formulas	Percent	Cumulative Percent
DC Plan (1)	16.76	16.76
1	29.50	46.26
2	28.21	74.47
3	11.16	85.63
4	4.78	90.41
5 or more	9.59	100.00
Social Security Integration		
Offset Plan	18.06	18.06
Excess Plan	6.82	24.88
Not Integrated DB Plan	58.36	83.24
DC Plan	16.76	100.00
FAP Formula		
Flat Rate Plan	9.19	
Career Average	13.92	
DC Plan	16.76	
Final Average	12.03	
Final Average Highest	5.65	
Final Average Highest Consecutive	18.14	
Average Highest	16.70	
Average Highest Consecutive	20.99	

Source: Author's Tabulations from SCF 1983 and PPS
1) All figures are weighted by covered workers in the SCF.
2) FAP formula shares total more than 100 due to multiple FAPs in many plans.
3) See text for definitions of categories.

Table 2

Estimate of Equation (6) for Single Firm, 1969-1980				
Variable	Coefficient	Standard Error	T-Statistic	Sample Mean
Constant	7.6008	0.1640	46.3369	1.00
Age	0.0927	0.0065	14.2758	51.00
Tenure	0.0280	0.0013	22.1132	16.75
Age^2	-0.0011	0.00006	17.0787	2628.66
$Tenure^2$	-0.0006	0.00004	17.3808	327.03
R^2: 0.111 Number of Persons: 4001 Number of Observations: 48012				

Notes:
1) Dependent variable is the natural log of annual wages.
2) Regression also includes 11 year effects (not shown).

Table 3

Estimate of Equation (9) with k = 1 for Single Firm, 1970-1980				
Variable	Coefficient	Standard Error	T-Statistic	Sample Mean
Constant	0.2425	0.0156	15.5523	1.00
ΔAge^2	-0.0025	0.00016	15.7848	102.00
$\Delta Tenure^2$	-0.0007	0.00011	6.2257	33.50
R^2: 0.1392 Number of Persons: 4001 Number of Observations: 44011				

Notes:
1) Dependent variable is the one-period increase in natural log of annual wages.
2) Regression also includes 10 lagged year effects (not shown).

Table 4

Method of Moments Estimators for $\Gamma = (\rho, \sigma_p, \sigma_\eta)$, Single Firm				
Sample Restrictions	AR(1) term ρ	Perm SD σ_p	Trans SD σ_η	Number of Individuals Included, Excluded
t = 0.0				
ρ estimated	0.9856 (0.0046)	0.2471 (0.0081)	0.0332 (0.0113)	4001, 0
ρ = 1, imposed	——	0.2442 (0.0081)	0.0376 (0.0099)	4001, 0
t = 0.2				
ρ estimated	0.9533 (0.0077)	0.1587 (0.0020)	0.0553 (0.0022)	3816, 185
ρ = 1, imposed	——	0.1525 (0.0018)	0.0613 (0.0017)	3816, 185
t = 0.4				
ρ estimated	0.9461 (0.0088)	0.1407 (0.0015)	0.0567 (0.0020)	3666, 335
ρ = 1, imposed	——	0.1345 (0.0011)	0.0628 (0.0015)	3666, 335

Notes:
1) Standard errors are in parentheses.
2) t refers to the threshold such that if an individual had an annual income less than t*(predicted income that year) he is excluded from the sample.

Table 5

Variance Decomposition for Single Firm				
Sample Restrictions	Permanent Shock σ_ν	Transitory Shock σ_η	Number Included	Number Excluded
t = 0.0				
q = 0	0.1780 (0.00012)	0.0850 (0.00013)	4001	0
q = 2	0.1954 (0.00010)	——	4001	0
t = 0.2				
q = 0	0.1190 (0.00004)	0.0899 (0.00005)	3816	185
q = 2	0.1228 (0.00005)	0.0699 (0.00012)	3816	185
t = 0.4				
q = 0	0.1070 (0.00003)	0.0889 (0.00005)	3666	335
q = 2	0.1077 (0.00004)	0.0857 (0.00009)	3666	335

Notes:
1) Standard errors are in parentheses.
2) t refers to the threshold such that if an individual had an annual income less than t*(predicted income that year) he is excluded from the sample.
3) q refers to the degree of MA(q) serial correlation in the transitory shock to which the estimates of Γ are robust.

Table 6

Method of Moments Estimators for $\Gamma = (\rho, \sigma_\nu, \sigma_\eta)$, PSID				
Sample Restrictions	AR(1) term ρ	Perm SD σ_ν	Trans SD σ_η	Number of Individuals Included, Excluded
t = 0.0				
ρ estimated	0.9362 (0.0504)	0.1568 (0.0113)	0.1536 (0.0098)	1254, 0
ρ = 1, imposed	——	0.1475 (0.0089)	0.1599 (0.0079)	1254, 0
t = 0.2				
ρ estimated	0.9354 (0.0517)	0.1419 (0.0085)	0.1318 (0.0084)	1232, 22
ρ = 1, imposed	——	0.1338 (0.0055)	0.1381 (0.0059)	1232, 22
t = 0.4				
ρ estimated	0.9518 (0.0495)	0.1263 (0.0072)	0.1186 (0.0071)	1185, 79
ρ = 1, imposed	——	0.1213 (0.0051)	0.1226 (0.0054)	1185, 79

Notes:
1) Standard errors are in parentheses.
2) t refers to the threshold such that if an individual had an annual income less than t*(predicted income that year) he is excluded from the sample.

Table 7

Variance Decomposition for PSID – Full Sample				
Sample Restrictions	Permanent Shock σ_ν	Transitory Shock σ_η	Number Included	Number Excluded
t = 0.0				
q = 0	0.1783 (0.0004)	0.1763 (0.0003)	1829	0
q = 2	0.1988 (0.0011)	0.0418 (0.0068)	1405	424
t = 0.2				
q = 0	0.1490 (0.0004)	0.1570 (0.0003)	1796	33
q = 2	0.1376 (0.0004)	0.1167 (0.0007)	1375	454
t = 0.4				
q = 0	0.1204 (0.0004)	0.1473 (0.0003)	1727	102
q = 2	0.1207 (0.0003)	0.1037 (0.0006)	1316	513

Notes:

1) Standard errors are in parentheses.
2) t refers to the threshold such that if an individual had an annual income less than t*(predicted income that year) he is excluded from the sample.
3) q refers to the degree of MA(q) serial correlation in the transitory shock to which the estimates of Γ are robust.

Table 8

Variance Decomposition for PSID – Education, Occupation, and Industry			
Category	Permanent Shock σ_ν	Transitory Shock σ_η	Fraction of Sample
Education			
0-8 years	0.1717 (0.0238)	0.1103 (0.0605)	0.1112
High School Diploma	0.1216 (0.0190)	0.1432 (0.0263)	0.3483
College Degree	0.1215 (0.0293)	0.0622 (0.0934)	0.1469
Occupation			
Professional and Technical	0.1581 (0.0194)	——[1]	0.1964
Managers	0.1333 (0.0315)	0.0775 (0.0885)	0.1047
Craftsmen	0.1005 (0.0301)	0.1607 (0.0307)	0.2029
Operatives and Laborers	0.1347 (0.0198)	0.1520 (0.0288)	0.2582
Service Workers	0.1374 (0.0347)	0.1801 (0.0432)	0.0815
Industry			
Manufacturing	0.1384 (0.0170)	0.1247 (0.0308)	0.3346
Wholesale and Retail Trade	0.1783 (0.0234)	0.0662 (0.1027)	0.1069
Finance, Insurance, and Real Estate	0.0732 (0.1017)	0.1686 (0.0720)	0.0335
Public Administration	0.1281 (0.0311)	0.0609 (0.1066)	0.1171

Notes:
1) The point estimate of σ_η^2 for professionals is negative but insignificant.
2) Standard errors are in parentheses.
3) Number of Observations: 1375
4) Near-zero threshold is: 0.2, Robust to MA(2) in transitory shock

Table 9

Characteristics of the Covered Population in the SCF	
Percent Female	41.50
Date of Birth	1941.54
Date of Hire at current job	1972.26
Hours worked per year	2051.56
Annual Wage	22880.61
Percent Nonwhite	18.09
Years of Education	13.44
Percent Married	76.72
Percent Making Voluntary Contributions	20.34
Spouse Date of Birth (if married)	1941.95
Percent of Wages Contributed (if positive)	5.42
Weighted number of observations (millions)	34.08

Source: Author's computations from the SCF and PPS

Table 10

Sample Pension Moments				
	Mean	Std. Deviation	Minimum	Maximum
Ratio of E[p(w)] to p(w) under certainty				
$\sigma_\eta = 0.05$	0.9869	0.1062	0.0521	3.9909
$\sigma_\eta = 0.10$	1.0142	0.1323	0.0530	5.1482
$\sigma_\eta = 0.15$	1.0576	0.1578	0.0544	5.5907
$\sigma_\eta = 0.20$	1.0874	0.1839	0.0561	5.5943
Coefficient of Variation for p(w)				
$\sigma_\eta = 0.05$	0.1955	0.1338	0.0000	2.3634
$\sigma_\eta = 0.10$	0.4115	0.2396	0.0000	1.8953
$\sigma_\eta = 0.15$	0.6706	0.3866	0.0000	2.3860
$\sigma_\eta = 0.20$	0.8549	0.4905	0.0000	3.3861

Source: Author's computations from the SCF and PPS

1) All pension moments are weighted to reflect the covered population in 1983.

2) See text for further explanation.

Table 11

Estimates of Equation (16)--Sample Correlations of Pension Moments and Uncertainty Levels									
	Standard Deviation of Permanent Shock to Income is 0.15				Standard Deviation of Permanent Shock to Income is 0.10				
Sample	β_0	β_1	γ_0	γ_1	β_0	β_1	γ_0	γ_1	Number of Observations
All workers	1.0579 (0.0056)	0.1080 (0.0833)	0.6459 (0.0134)	0.0245 (0.1989)	1.0153 (0.0047)	0.0686 (0.0689)	0.4004 (0.0083)	0.0199 (0.1227)	1237
Private	1.0181 (0.0042)	0.1278 (0.0622)	0.5858 (0.0212)	-0.0525 (0.3160)	0.9911 (0.0030)	0.0679 (0.0439)	0.3582 (0.0129)	-0.0204 (0.1918)	729
Public	1.1141 (0.0116)	0.0993 (0.1732)	0.7313 (0.0105)	0.1601 (0.1565)	1.0500 (0.0102)	0.0834 (0.1533)	0.4607 (0.0067)	0.1006 (0.1001)	508
Union	1.0546 (0.0046)	0.0982 (0.0734)	0.5574 (0.0217)	-0.5770 (0.3423)	1.0207 (0.0023)	0.0492 (0.0361)	0.3429 (0.0135)	-0.3246 (0.2130)	503
Non-union	1.0602 (0.0089)	0.1123 (0.1288)	0.7117 (0.0161)	0.2839 (0.2325)	1.0116 (0.0076)	0.0839 (0.1098)	0.4421 (0.0099)	0.1673 (0.1429)	734
Private, Union	0.9982 (0.0051)	0.1533 (0.0818)	0.3923 (0.0330)	-1.3803 (0.5290)	0.9937 (0.0027)	0.0881 (0.0425)	0.2308 (0.0200)	-0.7800 (0.3208)	257
Public, Union	1.1136 (0.0046)	0.0389 (0.0742)	0.7307 (0.0149)	0.2386 (0.2416)	1.0489 (0.0023)	0.0074 (0.0368)	0.4605 (0.0092)	0.1361 (0.1481)	246
Private, Non-union	1.0293 (0.0058)	0.1051 (0.0844)	0.7003 (0.0239)	0.3967 (0.3503)	0.9895 (0.0043)	0.0627 (0.0629)	0.4315 (0.0146)	0.2091 (0.2116)	472
Public, Non-union	1.1150 (0.0219)	0.1443 (0.3172)	0.7310 (0.0149)	0.1109 (0.2146)	1.0514 (0.0195)	0.1380 (0.2830)	0.4605 (0.0098)	0.0795 (0.1418)	262

Source: Author's estimates from the SCF and PPS.