Discussion Paper 6

Institute for Empirical Macroeconomics
Federal Reserve Bank of Minneapolis
250 Marquette Avenue
Minneapolis, Minnesota 55480

January 1989

BEHAVIOR OF MALE WORKERS AT THE END OF THE LIFE-CYCLE: AN EMPIRICAL ANALYSIS OF STATES AND CONTROLS

John Rust

Department of Economics
University of Wisconsin

ABSTRACT

This paper estimates the expectations of older male workers in the form of a 130 million element Markov transition probability matrix specifying the joint stochastic process for workers' income, health, martial and employment status, conditioned on workers' decisions about labor force participation and collection of Social Security benefits. The estimated transition matrix will be used in subsequent work to estimate the unknown parameters of workers' utility functions under the assumption that their behavior is governed by the solution to a dynamic programming model. The paper also discusses some of the problems involved in constructing good measures of workers' states and decisions.

I am extremely grateful for financial support from the National Institute on Aging and National Science Foundation grant SES-8721199, the latter which provided the computer resources that made this work possible. I am equally thankful to the Institute for Empirical Macroeconomics at the Federal Reserve Bank of Minneapolis which financed six months of uninterrupted research time that made this work possible. Finally, I'd like to thank B. J. Lee for his able research assistance, and to Lawrence Christiano, Finn Kydland, Edward Prescott, David Runkle, Christopher Sims, and Dick Todd for helpful comments on my work during seminars presented at the Federal Reserve.

This material is based on work supported by the National Science Foundation under Grant No. SES-8722451. The Government has certain rights to this material.

1. Introduction

This is the second installment in a series of three papers studying the behavior of men at the end of the life-cycle. The first paper, "A Dynamic Programming Model of Retirement Behavior" (Rust, 1988a), constructed a theoretical model based on the hypothesis that workers maximize expected discounted lifetime utility. The model treats observed behavior as a realization of a *controlled stochastic process* $\{x_t,d_t\}$ derived from the solution to a stochastic dynamic programming problem (DP). Estimation of the DP model requires observations of the worker's *state* x_t and *control* d_t, and a specfication of the markov transition probability density $\pi(x_{t+1}|x_t,d_t)$ representing a stochastic "law of motion" that embodies workers' beliefs of uncertain future events.

This paper uses the Retirement History Survey (RHS) to construct state and control variables $\{x_{ti},d_{ti}\}$, $t=1,\ldots,T_i$, $i=1,\ldots,I$, for a sample of I=8131 male respondents interviewed biannually from 1969 to 1979. I discuss some of the conceptual problems involved in constructing measurements of $\{x_t,d_t\}$ so that the resulting discrete-time, discrete-state DP model makes the best possible approximation to the underlying continuous-time, continuous-state decision process. I present my solutions to the measurement problems and conduct an extensive comparative data analysis to assess the overall quality of the resulting variables. Finally, I present estimates of workers' expectations, in the form of an estimated transition probability matrix $\hat{\pi}$.

All this work is building up to the third paper of the series, which will use the constructed state and control variables and the estimated transition probability matrix as inputs to a "nested fixed point" maximum likelihood algorithm (Rust 1988b) to estimate the unknown parameters of workers' utility functions. The success of the final stage depends critically on accurate measurements of $\{x_t,d_t\}$ and correct specification of workers' beliefs π.

The paper is organized as follows. Sections 2 and 3 summarize the principal findings. Section 2 describes the state and control variables constructed from the RHS dataset and presents the main conclusions of the data analysis. Section 3 specifies the functional form of workers' beliefs, and summarizes the main empirical findings. The remaining sections present details on the construction of $\{x_t, d_t\}$ and the numerical estimates of $\hat{\pi}$ that comprise the evidence for the conclusions drawn in sections 2 and 3.

2. State and Control Variables: Main Findings

Following the notation of Rust (1988a) the DP model requires a vector of state variables, $x_t \equiv (w_t, y_t, aw_t, h_t, a_t, e_t, ms_t)$, defined by:

w_t : accumulated net financial and tangible nonfinancial wealth
y_t : total income from earnings and assets
aw_t: Social Security "average monthly wage"
h_t : health status of worker (good health/poor health/disabled/dead)
a_t : age of worker
e_t : employment status (full-time/part-time/not employed)
ms_t: marital status (married/single)

and control variables, $d_t \equiv (s_t, c_t)$ defined by:

s_t : employment search decision (full-time/part-time/exit labor force)
c_t : planned consumption expenditures.

In the last twenty years, several panel datasets have accumulated sufficiently detailed data to permit construction of the required variables: the Panel Survey on Income Dynamics (PSID), the National Longitudinal Survey (NLS), and the Retirement History Survey (RHS). The RHS is the largest and most comprehensive of the three, explicitly designed by the Social Security Administration (SSA) to study the dynamics of retirement behavior. A special feature of the RHS is the availability of matching records from Census and SSA that permit direct validation of response error in several key variables. The Social Security Earnings Record (SSER) contains each covered worker's wage earnings (up to the statutory maximum taxable earnings) and quarters of coverage from 1939 to 1974. The Social Security Master Beneficiary Record (SSMBR) contains actual payments of Social Security Old Age, Survivors, Disability and Death benefits (OASDI) to each respondent, spouse and dependent from 1969 to 1978. The combination of finely detailed data, large sample size, long duration, plus the existence of linked Census and SSA records, make the RHS the data set of choice for estimating the DP model.

Having said this is not to deny the sober truth that even with the linked RHS records, there is a limit to how accurately one can measure the "true" states and decisions of individuals. Besides the obvious problems of missing data, response and coding error, estimation of the DP model presents three additional problems: 1) choice of time discretization, 2) choice of state discretization, and 3) construction of observable indicators of latent state and control variables.

Although the individual's actual decision process is best modelled in continuous-time, the data are collected and the DP model is formulated in discrete-time. In theory, use of discrete-time models is not a limitation since it has been shown that under very general conditions one can formulate a discrete-time DP model that approximates an underlying continuous-time DP model arbitrarily closely as the time interval goes to zero (van Dijk, 1984). One can account for absence of data on (x_t, d_t) between survey dates by forming a marginal likelihood function that integrates out the missing observations. In practice, however, computational and data limitations forced me to use fairly coarse two year time intervals. The computational limitations arise from the numerical integrations required to form the marginal likelihood function and the "curse of dimemsionality" inherent in DP models with fine time grids. The data limitations stem from the lack of "instantaneous" measures of income flows. In any case I prefer to use the available annual measures of income flows and hours of work on the belief that these measures better represent the worker's state and labor search decisions than the instantaneous measures. Analytically, the disadvantage of the discretization is that it implicitly committs workers to fixed consumption and labor supply values for two year periods. I should point out, however, that even with two year time intervals, the worker is given 20 opportunities to revise his decisions between age 58 and the terminal age, 98. Such a model is quite a bit more flexible than existing models such as Burtless and Moffitt (1984) that don't allow workers to revise consumption and labor supply decisions at all. In any case I leave the analysis of the consequences of time aggregation to the actual estimation of the DP model in the third paper of this series.

The state variables y_t, w_t, and aw_t which are naturally treated as continuous must also be discretized in order to estimate the DP model. Similar to the discretization of time, there are theorems guaranteeing that one can approximate a continuous-state DP model abritrarily closely by a discrete-state DP model (Bertsekas and Shreve, 1976). In practice I have found that the DP solution is not very sensitive to the discretization of the state variables, and that one can obtain a good approximation using fairly coarse grids (Rust, 1987). In this study I use a grid size of \$1,000 (1968 dollars) which turned out to be more than adequate given the two year time discretization.

The most difficult problem, however, was construction of good measurements of health h_t, labor search, s_t, and consumption c_t, none of which are directly observable. Although the RHS asked respondents to list the amount spent on individual consumption items, in my opinion the list was too incomplete to construct reliable estimates of total consumption.[1] Since the RHS has very complete, detailed data on income and wealth,[2] my approach is to infer c_t from the budget equation

(2.1) $$w_{t+1}=w_t+y_t-c_t.$$

Unfortunately, the RHS only recorded income in the even-numbered years immediately preceding each survey date. Thus, in order to construct c_t I needed to impute income in odd-numbered years. This in turn necessitated construction of complete labor force histories for each worker, including total annual hours worked in each year.[3] Using hours worked together with annual wage earnings data from the SSER (available up to 1974), I was able to

[1] In fairness, I should mention that some authors such as Hammermesh (1982) have, with some success, used this data to impute total consumption c_t.

[2] The wealth and income data used in this study were "pre-cleaned" by the program IMPUTE written by Beth van Zimmerman and Phil Farrell, research associates of Michael Hurd, SUNY Stony Brook. Besides imputing missing values, the program estimated the value of service flows for owned assets such as autos and housing at a presumed opportunity cost of 3%..

[3] Constructing labor force histories turned out to be a major undertaking, requiring over 80 pages of FORTRAN code and over four months of full-time work to write and debug. The difficulties arose from the need to carefully track the survey skip patterns to extract the required variables from a battery of more than 130 questions in the "Work Experience" section of the RHS. Fine attention to detail was requied to avoid misclassifying 20% of the sample of workers with "non-standard" employment histories some of which involved multiple job transitions within the two-year period.

impute income in odd-numbered years and construct estimates of c_t over the two year sample interval. A limitation of the income data is absence of capital gains. I dealt with this problem by attributing 100% of the change in house value to capital gains (provided the respondent was a homeowner and had not moved within the interval) and by excluding workers who had substantial real estate or equity holdings. I faced equally difficult problems constructing h_t and s_t, but I will defer the details of their construction until later.

Good measurements of $\{x_t, d_t\}$ are absolutely critical to the success of the DP model since its is highly nonlinear in variables and there currently is no good theory of errors in variables for such models. Wherever possible, I have attempted to obtain independent measures of the variables to assess the magnitude of the measurement error. I have also constructed an array of associated "variable flags" to indicate the degree of confidence in each of the constructed state and control variables. By setting the appropriate flags, I can screen out questionable cases to obtain a core subsample for which confidence in the data is relatively high. To guard against the possibility that such screening could produce unpredictable sample selection biases I have compared the distribution of each variable to its distribution in the full sample using all available observations. Because presentation of tabulations of the flag variables takes us too much into the "guts" of the computer programs that generate the state variables, I have decided against presenting them. Instead I describe the nature of any special data or sample selection problems where appropriate.[4]

[4] Of course, I will be happy to provide the reader with the data and documented versions of all computer programs used to generate the variables so that other researchers can verify any of my results should they choose to do so.

I can state the major conclusions of the data analysis as follows:

1. At the aggregate level, the data show workers making a smooth transition from work into retirement, gradually reducing consumption and labor supply but maintaining wealth levels intact. This is consistent with the behavior of a neoclassical, risk-averse consumer who attempts to smooth consumption and leisure streams, and provide bequests to his heirs. However at the individual level, the data are anything but smooth: measured consumption shows erratic fluctuations and labor supply has an abrupt discontinuity, with the typical worker staying at his full-time job up until retirement age (62 to 65), at which time he applies for Social Security, quits his job, and remains out of the labor force for the rest of his life.

2. Constructing consumption expenditures from the budget equation, $c_t=w_t-w_{t+1}+y_t$, is susceptible to the frequent and often large measurement errors in wealth, possibly exacerbated by absence of good information on capital gains. The majority of the erratic variations in measured consumption appear to be attributable to response errors in wealth.

3. The distribution of real wealth changes is centered about 0, but with a large variance. On average, net worth is not very large, about 4 times annual income, and a substantial fraction of this wealth, 50-60%, is tied up in housing. These facts strongly support the view that the large swings in measured consumption are simply a result of response errors in wealth rather than erratic consumption/savings behavior. Although a simple test of the null hypothesis H_0: $c_t=y_t$ vs. H_A: $c_t \neq y_t$ rejects at the 5% level (but not at the 1% level), the fact that the average change in wealth is $-658 with a standard deviation of $47,015 makes it very hard to distinguish between alternative theories of consumption/savings behavior. Because of the problems involved in

accurately measuring wealth and therefore consumption, I have opted to start with a simpler DP model based on the hypothesis that $c_t=y_t$. In this model workers choose labor force participation strategies to maximize the expected discounted value of the utility of income, ignoring wealth and bequests.

4. Although respondent's total income is only recorded for even numbered years, the existence of independent income measures in the SSER and SSMBR datasets allowed me to construct reliable income imputations in odd numbered years. Thus, if wealth changes are indeed an insignificant component of consumption, total imputed income will be a good measure of actual consumption.[5]

5. The distribution of total annual hours worked is highly bimodal with most of its mass at either 0 or 2000. While some of this bimodality is likely an artifact of response error (with workers simply rounding their responses to 40 hours/week, 50 weeks/year), it does indicate that the tripartite classification of labor force status e_t into 1=full-time, 2=part-time, or 3=unemployed does not grossly misrepresent the data and that this measure is robust to fairly large variations in the hours cutoffs defining the three e_t states. Overall, the distributions provide little evidence to support the view that workers treat annual hours of work as a continuous decision variable.

[5]Biannual income was used only for purposes of constructing a measure of consumption. Based on conclusions 2 and 3 above, I have decided to exclude consumption/savings decisions and formulate a DP model with biannual time intervals, measuring workers' states over the preceding even-numbered survey years. Thus, the DP model will actually use the annual income flows that were recorded in the surveys. For further justification of this approach, see conclusion 6.

6. A systematic response error problem known as the *seam problem* produces exaggerated estimates of labor state transitions across the survey dates, or seams, of the RHS. This leads to artificial cyclical variations in the transition probabilities for "across-seam" transitions as compared to "between-seam" transitions. The variation is apparently due to imperfect recall of labor force history in the earlier year of the two year interview frame, leading to inconsistencies between recalled labor force status in the current interview and the labor force status reported in the last interview. One can ameliorate the seam problem by "skipping over the seams" and tracking transitions between the even-numbered years immediately preceding the odd-year survey dates in order to reduce the amount of recall on the part of respondents. This convinced me to formulate a DP model with a time period of two years rather than with a more fine-grained model with a one year time period.

7. There are three possible measures of the "job search" control variable: s_t= self-reported planned hours of work in the year following the survey, s_t=actual hours worked in the year following the survey, or $s_t=e_{t+1}$, actual hours worked in the second year following the survey. The last measure corresponds to a "perfect control" model wherein an unemployed worker who decides to go back to work is successful with probability 1. In my opinion the perfect control model is not a priori plausible, so I focus on the other two measures which correspond to "imperfect control" models where unemployed workers who decide to look for a full-time job have less than a 100% chance of being successful. Probably reflective of the fact that "talk is cheap", it appears that the first measure of s_t is a much more noisy measure of actual job search behavior than is the second measure. Since the data show that the second measure allows for a much more intuitive and predictable relationship between job search decisions and subsequent employment outcomes, I adopt it as the measure of s_t used to estimate the DP model.

8. The four-way classification of health status h_t into 1=good health, 2=health limitation but not disabled, 3=disabled, and 4=dead, seems to produce sensible results despite the inherently subjective nature of health status. Use of actual benefits paid from the SSMBR data was critical to the quality of h_t since self-reported measures of health significantly underestimate the occurrence of health state 3 due to systematic under-reporting of Social Security disability receipts by respondents. The Social Security requirement of doctor examination for disability qualification seems to be a significant factor in identifying individuals with substantially greater health problems as indicated by their significantly higher ex post mortality. An unfortunate aspect of the disability classification is the fact that no workers become disabled after age 62. This is an artifact of Social Security rules that automatically convert disability payments into OASI payments after age 62.

9. The SSMBR data allow me to identify when individuals actually apply for, and receive, OASI benefits. Twenty percent of eligible recipients apply for benefits as soon as they are able to receive them at the early retirement age 62, and another twenty per cent apply for benefits at the normal retirement age 65. Overall 60% of eligible workers apply for and receive OASI between the ages of 62 and 65. The implied retirement hazard and frequency distributions computed using the SSMBR data and a definition of "retirement" as the age of first receipt of OASDI differ significantly the distributions computed by other researchers using the RHS data and other definitions of retirement. In order to better understand the phenomenon of early retirement and the pronounced bimodal distribution of retirement dates, I have included a new control variable ss_t defined by

$$
(2.2)\ ss_t = \begin{cases} 0 \text{ if worker is not receiving OASI} \\ 1 \text{ if worker is receiving OASI and first applied for it} \\ \quad \text{before age 65 (early retirement)} \\ 2 \text{ if worker is receiving OASI and first applied for it} \\ \quad \text{after age 65 (normal retirement).} \end{cases}
$$

Including ss_t allows me to avoid ad hoc definitions of "retirement", separating the analysis of retirement behavior (i.e. collection of OASI) from labor supply behavior.

3. Estimation of Worker's Beliefs: Main Findings

Recall that workers' beliefs are represented by a Markov transition probability density $\pi(x_{t+1}|x_t,d_t)$. Under the assumption of homogeneous beliefs and rational expectations, one can "uncover" these beliefs from data on the realizations of $\{x_t,d_t\}$. Given the discretization of time and state variables proposed in section 2, π is a matrix with approximately 130 million elements. Clearly a non-parametric estimate of π is out of the question since nearly all cells of $\hat{\pi}$ would be estimated as identically zero even though we know that the corresponding transitions actually occur with positive probability. It is necessary, therefore, to find a parametric specification $\pi(x_{t+1}|x_t,d_t,\theta)$ that depends on a much lower-dimensional vector of unknown parameters θ in such a way that all relevant cells of π are assigned non-zero probabilities. It is also important to choose a specification that is parsimonious, yet sufficiently flexible so that the estimated model is consistent with the data. Direct parameterization of a 130 million element matrix seems out of the question, so a more clever approach must be employed. The strategy I have followed is to decompose π into a product of conditional and marginal densities and estimate each of the components separately. To see this more clearly, note that without loss of generality one can decompose a bivariate transition density f as follows:

$$(3.1) \quad f(x_{t+1},y_{t+1}|x_t,y_t)=f_1(y_{t+1}|x_{t+1},x_t,y_t)f_2(x_{t+1}|x_t,y_t)$$

$$=f_3(x_{t+1}|y_{t+1},x_t,y_t)f_4(y_{t+1}|x_t,y_t),$$

where f_1, f_2, f_3, and f_4 are defined from f in an obvious way. Although (3.1) shows that the ordering of the decomposition of f is irrelevant, it does make a difference when the functional form of f must be estimated from the data. For example, I have found empirically that future health h_{t+1} is much a more useful and interpretable variable for predicting future income y_{t+1} than the other way around. Having tried various decompositions of π, the one I found most plausible is given below

$$(3.2) \quad \pi(y_{t+1},e_{t+1},ms_{t+1},h_{t+1}|y_t,e_t,ms_t,h_t,a_t,d_t) =$$

$$\pi_y(y_{t+1}|e_{t+1},ms_{t+1},h_{t+1},y_t,e_t,ms_t,h_t,a_t,d_t) \times$$

$$\pi_e(e_{t+1}|ms_{t+1},h_{t+1},y_t,e_t,ms_t,h_t,a_t,d_t) \times$$

$$\pi_{ms}(ms_{t+1}|h_{t+1},y_t,e_t,ms_t,h_t,a_t,d_t) \times$$

$$\pi_h(h_{t+1}|y_t,e_t,ms_t,h_t,a_t,d_t).$$

Note that the decomposition (3.2) excludes the state and control variables c_t, w_t, aw_t from the original list presented in section 2. Consumption c_t and wealth w_t were excluded due to the measurement problems discussed in conclusion 3 of section 2. The Social Security average monthly wage aw_t (a complex average of the worker's historical earnings) was excluded since it turned out to sufficiently collinear with current income y_t that I could reduce the dimensionality of the model by making y_t do double duty as a proxy for aw_t. Finally future age a_{t+1} was excluded since it has a trivial non-stochastic transition rule: $a_{t+1}=a_t+2$ with probability 1.

The motivation for the decomposition of π given in (3.2) is that income y_t and employment status e_t are the most important state variables of the DP model, and therefore their evolution should be predicted as well as possible. If we view (3.2) as specifying π as a direct product of individual transition matrices, then π_y is the "innermost" component of the direct product, in the sense that income transitions are conditioned on the contemporaneously realized values of all the remaining state variables. From an empirical standpoint, including these contemporaneous values substantially improves the fit of the income regressions estimated in section 9.

The outermost component of the direct product, health status h_t, has additional structure resulting from the definition of health states $h_t=3$ and $h_t=4$. If I fix the values of the other variables $(y_t, e_t, ms_t, a_t, d_t)$, then π_h is represented by the following 4x4 transition probability matrix.

Figure 3.1: Structure of Health Transition Matrix

$$(3.3) \qquad \pi_h = \begin{bmatrix} \varphi_{11} & \varphi_{12} & \varphi_{13} & \varphi_{14} \\ \varphi_{21} & \varphi_{22} & \varphi_{23} & \varphi_{24} \\ 0 & 0 & \varphi_{33} & \varphi_{34} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

According to (3.1), death is treated as an absorbing state. Note that disability is also treated as an absorbing state in the sense that once a worker becomes disabled, he can only continue to stay disabled or die. This restriction was necessitated by data limitations. Although the Social Security SSMBR dataset includes the variable "date of termination of disability benefits", there were only a 2 or 3 cases where actual termination was observed. Perhaps this indicates problems in Social Security record-keeping, but it is more likely to an artifact of the definition of "disability" to which I alluded in section 2. According to Social Security rules, disabled workers who receive SSDI benefits past age 62 are automatically reassigned OASI benefits after turning 62. Thus, there is no real incentive for Social Security to keep track of the date when the actual physical disability terminates once the worker is older than 62. One can try to partially rectify the problem the following way: reclassify workers who received disability benefits prior to age 62 and who are now older than 62 and reporting that they are in good health as being in state $h_t=1$ rather than $h_t=3$. Unfortunately this reclassification scheme has its own problems: although it allows transitions from disability to good health ($h_t=3$ to $h_{t+1}=1$), there is no way to record transitions from $h_t=3$ to $h_{t+1}=2$ since the RHS variables do not allow us to distinguish between the states "existence of a health problem that limits one's ability to work" and "disability".

The remaining sections of the paper discuss the construction of the the state and control variables in more detail, and present estimation results for each of the four components of the of the decomposition of π given in (3.2). Having conducted an extensive specification search to find the appropriate functional form for π, I can summarize the main empirical findings below.

1. Age and income are relatively unimportant determinants of death rates after controlling for health, employment, and marital status. Death rates decrease slightly with income and actually decrease with age until age 67.[6] Not

[6]The latter conclusion disappears if I exclude the variable ss_t distinguishing respondents who are receiving OASDI. Since men who collect OASDI have higher death rates, excluding ss_t produces a model where death rates increase slightly with age.

surprisingly, single workers are significantly more likely to die than married workers. However even this variable has small impact relative to health h_t, labor supply and retirement decisions (s_t, ss_t). Workers who are in poor health (h_t=2,3) are 2 to 4 times more likely to die than healthy workers. There is an equally strong association between the job search decision s_t and the probability of death, but the nature of the relation depends critically on the worker's health and retirement status. If the worker is retired or disabled (h=3 or $ss \in \{1,2\}$), any attempt to return to work on either a full or part time basis is extremely hazardous, significantly increasing the risk of death. However if the worker has not already retired and is in relatively good health ($h \in \{1,2\}$ and ss=0), the decision to quit work is associated with significantly higher death rates. Although this latter finding may represent spurious causality due to failure to completely control for all dimensions of health status, from the standpoint of a worker behaving according to the DP model the association is necessarily interpreted as cause and effect.

2. The probability of becoming disabled is a sharply decreasing function of age: a result that is an artifact of the definition of disability discussed above. It is very likely that disability is an endogenous state variable (i.e the outcome of an underlying decision process), as evidenced by the fact that the probability of becoming disabled decreases well before age 62. The explanation is that the process involved in applying and qualifying for SSDI imposes significant costs on the worker, including doctor examination at the worker's expense. Naturally, the closer one is to the early retirement age of 62, the less incentive one has to incur the costs of applying for SSDI, especially when the probability of qualification is less than 1. Despite the difficulties with the disability classification, it is still a worthwhile distinction since the required doctor certification appears quite successful in identifying a group of workers who suffered serious health problems before early retirement age, as confirmed by their signifcantly higher mortality rates. Workers who become disabled after age 62 are covered under health state

h_t=2 (existence of a health problem that limits the ability to work or get around), so that the main problem with the health status variable is that older workers classified in health state h_t=2 are likely to include a greater percentage who would have otherwise been classified as disabled. Other findings of interest are the fact that both single and higher income workers are significantly less likely to become disabled. The finding for single workers might be partly a result of sample selection bias: single workers are presumably less likely to have a family support network to rely on, so they are more likely to become institutionalized if they have serious health problems. Such workers are lost from the sample since the RHS did not attempt to interview institutionalized individuals.

3. The probability of being in good health is a declining function of age and an increasing function of income. Whether the worker is single or married has no significant impact on the probability of being in good health. By far the most important determinant of future health is current health. Currently healthy workers are three times more likely to be in good health than currently unhealthy workers (h_t=2). There is weak evidence that continuing to work on a full or part-time basis increases the probability of being in good health. Conversely, the decision to quit working is associated with a deterioration in health. This result is corroborated by the fact that retired workers, $ss \in \{1,2\}$, are significantly less likely to be in good health. As with my comments in point 1, the association might indicate spurious causality due to imperfections in the measure of health status: healthier workers continue working while unhealthy workers quit and retire.

4. By far the most important variable predicting future marital status is current marital status: once single, a worker has less than a 7% chance of finding a new mate. Older workers are more likely to lose their spouse, while higher income workers are less likely to become single, at least up to an

income of $30,000. There is weak evidence that among single workers, the worse one's health, the more likely one is to remain single. Among married workers, being in worse health increases the chance of remaining married. Economic decisions such as the labor search decision s_t or the retirement decision ss_t appear to have little or no effect on future marital status.

5. As one would expect, future employment status e_{t+1} is most strongly affected by the employment search decision s_t[7]. In addition, the worker's previous employment state e_t has a significant impact on probability that the search decision s_t is realized. Thus, currently employed workers have a significantly higher chance that $s_t=1$ will result in $e_{t+1}=1$ than do part-time or unemployed workers, but interestingly unemployed workers have a significantly higher chance of being successful in gaining a full-time job than part-time employed workers. Conversely, if a worker decides to quit, he is more likely to realize his decision if he is currently unemployed than if he had a full or part-time job. Full-time workers are more likely to realize their quit decisions than part-time workers. Health status also has a very strong impact on employment status. Workers who become disabled are 2.5 times more likely to be out of the labor force and their chances of staying in a full-time job are less than 1/3 that of non-disabled workers. There are clear aging effects on the ability to continue working full-time; for example, the probability that a 67 year old worker will be successful in keeping or finding a full-time job is only 1/3 as high as an equivalent worker under 60. Income appears to be a statistically significant proxy for employability, with high income workers being 60% more likely to keep or obtain a full-time job than low income workers. Somewhat surprisingly, changes in marital status have no

[7]Recall that s_t is proxied by the actual realized employment state in the odd-numbered years between the even-numbered years in which e_t and e_{t+1} are recorded.

significant impact on employment status. Less surprising is the fact that retired workers are less likely to be fully employed and more likely to be unemployed, all other things equal.

6. In order to match the long-tailed cross-sectional income distributions, the stochastic process for income was assumed to have a transition density with a conditionally heteroscedastic lognormal distribution. Income is strongly autocorrelated with an autoregressive coefficient of .95, and there is evidence of nonlinearity in this relation in the sense that higher powers of current income y_t enter the model with highly significant coefficients. The variance of future income y_{t+1} is an increasing function of current income, but the relation is far from proportional: a worker earning $50,000 has a standard deviation in y_{t+1} of $12,000 whereas a worker earning $5,000 has a standard deviation in y_{t+1} of $2,000. Health status has a significant impact on income prospects: healthy workers expect a 3% increase in real income, and disabled workers expect a 5% increase in income. However currently healthy workers who become disabled expect a 20% drop in income. Changes in marital status have large and statistically significant impacts on income. A worker who loses his wife expects a 25% drop in income, and a bachelor who has no prospects of remarriage expects his income to fall by about 20%. Of course, realized employment states are very strong predictors of income. Workers who keep working at their full-time jobs expect a 25% increase in income, while workers who exit from the labor force expect an 18% decrease in income (in the absence of Social Security). The estimated income process successfully captures the main features of OASDI benefit rules, including the regressive nature of the payoffs, the extra benefits to a spouse, the early retirement penalty, and the effect of the "earnings test" for workers under 70.

7. It's possible that there exist unmeasured differences or *heterogeneity* among workers that create systematic differences in workers' beliefs but which

are not captured in the list of state and control variables set forth in this paper. In order to assess the potential magnitude of this problem I included several demographic variables in the estimation of workers' beliefs π, including a variable classifying the respondent as a "work lover" or "leisure lover" as well as his education, race, and the industry and occupation of his longest held job. Surprisingly, except for the finding that blacks expect significantly lower incomes than whites, none of these variables had a major impact on the estimated transition probability $\hat{\pi}$. Thus, the available evidence indicates that the list of state and control variables set forth in this paper provides a reasonably complete set of "sufficient statistics" for the states and decisions of my sample of workers. In particular, there is no strong evidence that the failure to account for unmeasured hetereogeneity leads to a gross misrepresentation of workers' beliefs.[8]

4. Age, Marital Status, and Demographic Variables

A set of variables that we ought to be able to measure accurately are the identity of the respondent, his or her age, and basic demographic variables such as race, education, and the occupation/industry of the respondents' longest job. By in large this is true of the RHS data, although cross-checks of self-reported values with Census and Social Security records do indicate discrepancies. For example, out of an initial 1969 sample of 8,131 males, reported and recorded Census date of birth differed by more than 1 year in 53 cases, in some cases by more than ten years. In order to estimate the DP model, I need to track each of the 8,131 original male respondents over the

[8]In order to keep the length of this article within bounds, I have chosen not to present the estimation results that lead to this conclusion. I defer the presentation of the results to the third paper of the series which will examine the heterogeneity issue in more detail.

ten year survey period. RHS respondent identifiers allowed me to distinguish the original male respondent from his surviving spouse (or other household members), and in conjunction with comprehensive death records compiled by Paul Taubman I was able to determine whether or not the original respondent died even if he was no longer responding to the survey. Table 4.1 provides a response summary that shows that the basic sample of original male respondents decreased from 8,131 in 1969 to 4,298 in 1979. There was significant attrition of the original 1969 male respondents over the survey. Table 4.1 shows that the attrition was due to the respondent's death in 2327 cases, and non-response in 1506 cases.

Table 4.1: RHS Response Summary

	71	73	75	77	79
original 69 male respondent	7054	6239	5541	4811	4298
no-response, 69 respondent still alive	534	889	1104	1315	1426
no response, 69 respondent dead	152	361	610	917	1245
surviving spouse responds, 69 respondent dead	244	488	722	908	1075
other relation responds, 69 respondent dead	37	58	54	60	7
other relation responds, 69 respondent alive	110	96	100	120	80
	8131	8131	8131	8131	8131

A discrepancy exists between the individual sub-record identifier in the SSER tapes and the respondent identifiers on the original RHS tapes: the former showed 8091 original respondents in 1971 versus 7054 in the RHS. The former figure could not possibly be right given that 433 respondents had died by the 1971 interview. Indeed, a second cross-check using the Census Non-Response file[9] agreed with the RHS identifiers. This provided a sobering reminder that one cannot necessarily trust the SSA's internal accounting data more than the RHS interview data.

Relatively minor discrepancies exist in the data on marital status. For example, nine individuals reported being married with spouse not present in 1969, but reported having never been married in 1971; two cases reported having a deceased spouse in 1969 and never having been married in 1971. Thirty-five cases classified themselves as being a surviving spouse in 1971, but listed themselves as having a "spouse in 69 but not in 71" instead of the correct response "69 spouse deceased, no 71 spouse". Using the corrected marital status data, I defined the marital state variable ms_t as follows:

$$ms_t = \begin{cases} 1 & \text{if respondent is married} \\ 2 & \text{if respondent is widowed, separated, divorced, or never married} \end{cases}$$

Table 4.2 presents the computed 2 state Markov transition matrices for marital status (where "M" denotes cases which are missing due to death or non-response). The transition matrices change in the expected way over time: the probability of becoming a widower over the two year survey frame increases from 6% in 1969 to 9% in 1977. The probability of remarriage decreases over time from 7% in 1969 to 2% in 1977.

[9]The Non-Response File was compiled by the Census in the process of conducting the RHS interviews and was used by SSA as part of an internal auditing system to remove cases in which the interviewer was unable to contact the original 1969 respondent or related household members. For some reason the non-response data was not included on the RHS tapes, and is only available separately as a subfile of the SSMBR tape. The Non-Response file will also be used in section 5 to identify men who were institutionalized after the 1969 interview.

BLANK

Table 4.2: Markov Transition Matrices for Marital Status

Year of Transition		Cell Counts					Transition probabilities	
1969-1971		1	2	M	total	%	1	2
	1	6180	386	512	7078	87.05	0.9412	0.0588
	2	65	814	174	1053	12.95	0.0739	0.9261
					8131	100.00		
1971-1973		1	2	M	total	%	1	2
	1	5434	405	406	6245	83.88	0.9306	0.0694
	2	66	976	158	1200	16.12	0.0633	0.9367
					7445	100.00		
1973-1975		1	2	M	total	%	1	2
	1	4760	410	330	5500	79.93	0.9207	0.0793
	2	63	1140	178	1381	20.07	0.0524	0.9476
					6881	100.00		
1975-1977		1	2	M	total	%	1	2
	1	4141	391	312	4844	75.49	0.9137	0.0863
	2	48	1310	215	1573	24.51	0.0353	0.9647
					6417	100.00		
1977-1979		1	2	M	total	%	1	2
	1	3602	372	220	4194	71.10	0.9064	0.0936
	2	32	1434	239	1705	28.90	0.0218	0.9782
					5899	100.00		

Table 4.3 presents the estimation results for π_{ms}, the marital status component of the decomposition of π given in (3.2). The elements of π_{ms} were estimated by maximum likelihood, using a linear-in-parameters, binomial logit specification of the probability that $ms_{t+1}=2$. Note that the parameter standard errors and t-statistics have been corrected using White's (1982) formula; for purposes of comparison I also present the usual t-statistics computed from the diagonal of the inverse hessian. The estimation results in table 4.3 are based on a smaller subsample than table 4.2 (18,833 versus 34,773 observations) as a result of conditioning on the availability of complete observations for the state and control variables entering π_{ms}, and conditioning on a sample boolean variable. The boolean excludes respondents who are not the original 1969 male respondents, and further excludes respondents who are farmers or farm owners, respondents with significant pension wealth, and respondents who made sufficiently erroneous or suspicious responses as determined from the flag variables described in section 2. Overall, the estimation results in table 4.3 support the conclusions drawn in point 4 of section 3.

TABLE 4.3: Estimates of Marital Status Transition Probability

Dependent variable: $I\{ms_{t+1}=2\}$

variable	parameter estimates	corrected std. error	uncorrected t-statistic	corrected t-statistic
$s_t=1$	-0.05628402	0.14898058	-0.37838900	-0.37779435
$s_t=3$	-0.08757423	0.14583786	-0.59606925	-0.60049036
$ms_t=2,\ h_{t+1}=1$	-1.87853000	1.11208414	-1.74781953	-1.68919772
$ms_t=2,\ h_{t+1}=2$	-2.05862795	1.12510188	-1.89503827	-1.82972582
$ms_t=2,\ h_{t+1}=3$	-2.68821465	1.18536881	-2.30413183	-2.26782976
$ms_t=1,\ h_{t+1}=1$	4.06035625	1.10054546	3.80807860	3.68940349
$ms_t=1,\ h_{t+1}=2$	4.15716766	1.11334007	3.88661200	3.73396035
$ms_t=1,\ h_{t+1}=3$	4.67178723	1.08704320	4.43536638	4.29770154
a_t	-0.01895495	0.01755366	-1.11973463	-1.07982864
y_t	0.15003823	0.01758073	8.37073953	8.53424374
y_t*y_t	-0.00267115	0.00037456	-6.57936304	-7.13140351
$ss_t\in\{1,2\}$	-0.13039682	0.15527776	-0.84612070	-0.83976492
$ss_t\in\{1,2\},\ ms_t=2$	0.02793904	0.23740802	0.11844131	0.11768364

log likelihood -2347.85020571
grad*direc 6.46274344E-025

percent correctly predicted 97.09
total observations 18833

5. Health status

A key variable in the DP model is the worker's health status. This variable shifts the worker's mortality hazard, and affects his ability to work and enjoy leisure. In order to construct the health status variable, I used mortality data from Paul Taubman's "death tape", and a battery of over 75 questions on health status in the RHS. It turned out, however, that two of the 75 RHS health variables were most relevant for classifying health status: HLIM: "do you have any health condition, physical handicap, or disability that limits how well you get around?" and HWRK: "does your health limit the kind or amount of work or housework you can do?". Originally I used these variables, together with 15 other health-related questions and the respondent's report as to whether or not he received SSDI benefits to classify health status h_t into one of four states: 1=respondent in good health, 2=respondent has a health problem that limits his ability to work or get around, but is not severe enough for the worker to qualify for SSDI, 3=respondent has a health problem severe enough for him to qualify for SSDI, and 4=respondent is dead. My original construction of this variable yielded significantly lower estimates of the probability of being on SSDI than those of Bound (1986): 1.17% in 1969 versus Bound's estimate of 7.1% for men aged 55-64 in 1970. In addition, the data appeared to show an unexpected mass outbreak of poor health in 1975, with only 1254 respondents classified as $h_t=1$ and 3958 classified as $h_t=2$. By using the SSMBR OASDI payments data I was able to directly verify whether a worker was classified as disabled by SSA by determining whether he was receiving SSDI payments. Furthermore, analysis of the health input variables revealed that the HWRK variable for some unknown reason had 5956 missing values in 1975, and the remaining cases contained a disproportionate percentage of workers reporting a health limitation (1476 out of 2200). Either there is a coding error problem in HWRK75 or else the 1975 RHS survey had for some unknown reason primarily recorded HWRK for a subsample of individuals with health problems.[10] I fixed the problem by using only the HLIM variable to classify workers into health state h=1 or h=2, and merging the disability data from the SSMBR to classify disabled workers, h=3.

Another problem arose from the fact that the RHS survey did not attempt to track workers who became institutionalized, instead simply recording them as missing. There is good reason to believe that the failure to track institutionalized workers induces a sample selection bias since single workers are less likely to have a family support network to rely on, and are therefore more likely to become institutionalized and be lost from the sample. To correct this problem I merged data from the Census Non-Response file which records the reasons for non-response, including institutionalization. Analysis of health status of the institutionalized workers showed that among the sample of 113 institutionalized workers (36% of who were single in 1969 as compared to 13% for the sample as a whole), in only 1 case did the worker return to the RHS sample with improved health: the preponderant majority of institutionalized workers died within a few years after entering the institution. Based on this evidence I decided to redefine health state 4 as workers who are either dead or institutionalized.

A final problem was more difficult to resolve. Although we have fairly complete data on the month and year that a worker died, in order to be included in the estimation of the health transition probability matrix, we must observe the worker's state and control vector (x_t, d_t) in the survey period immediately preceding his death. Unfortunately there are many cases were the worker failed to respond to the survey for two or more survey periods preceding his death. Analysis of these cases shows that a disproprtionate

[10] Conversations with Gary Burtless of the Brookings Institution revealed that he encountered similar problems with HWRK75. Thus the problem is not likely to be due to a read error on my copy of the variable.

number consist of single men. One solution is to "remove" the intervening periods of missing data by treating the death as occuring just after the last survey to which the worker responded. Unfortunately this approach has the effect of "accelerating" the deaths of a fairly large group of workers, distorting the estimates of age-death profiles. I decided, therefore, to leave the data as they were, and simply acknowledge the possibility of sample selection bias that might lead to an underestimate of mortality rates for single workers.

Table 5.1 displays the transition probability matrices for my final definition of h_t. The data show a much more reasonable rate of disability receipt, 8.1% in 1969, which is much closer to Bound's estimate. The transition matrices generally appear to be quite reasonable, with workers in worse health states having signficantly higher risk of death and disability. Mortality rates appear fairly stable over time, and are in rough agreement with independent estimates calculated by Mott and Haurin (1985) using NLS data. Note that the transition probabilities in table 5.1 imply that disability is an absorbing state: once a worker becomes disabled he either remains disabled, becomes institutionalized, or dies. This is simply a reflection of the data limitations discussed in section 3: the SSMBR data do not record the date of termination of disability. As a result, in each survey year there are approximately 100 workers who report that they have no health problem that limited their ability to work or get around despite the fact that Social Security records indicate that they are disabled. Because the existing classification of disability confirms my a priori belief that disabled workers have significantly higher mortality rates, and more importantly, because this classification matches the aggregate disability rates compiled by Bound, I decided not to reclassify these workers as $h_t=1$.

Table 5.1: Health Transition Probabilities

Year		Cell Counts							Transition probabilities			
69-71		1	2	3	4	M	total	%	1	2	3	4
	1	4347	630	111	211	470	5769	70.99	0.8203	0.1189	0.0209	0.0398
	2	562	790	84	116	147	1699	20.91	0.3621	0.5090	0.0541	0.0747
	3	0	0	506	106	46	658	08.10	0.0000	0.0000	0.8268	0.1732
	4	0	0	0	0	0	0	00.00	0.0000	0.0000	0.0000	1.0000
							8126	100.00				
71-73		1	2	3	4	M	total	%	1	2	3	4
	1	3629	730	76	181	296	4912	69.82	0.7862	0.1581	0.0165	0.0392
	2	449	688	51	126	107	1421	20.20	0.3417	0.5236	0.0388	0.0959
	3	0	0	533	113	56	702	09.98	0.0000	0.0000	0.8251	0.1749
	4	0	0	0	0	0	0	00.00	0.0000	0.0000	0.0000	1.0000
							7035	100.00				
73-75		1	2	3	4	M	total	%	1	2	3	4
	1	2975	707	20	177	240	4119	66.16	0.7670	0.1823	0.0052	0.0456
	2	371	831	12	149	76	1439	23.11	0.2722	0.6097	0.0088	0.1093
	3	0	0	541	89	38	668	10.73	0.0000	0.0000	0.8587	0.1413
	4	0	0	0	0	0	0	00.00	0.0000	0.0000	0.0000	1.0000
							6226	100.00				
75-77		1	2	3	4	M	total	%	1	2	3	4
	1	2495	510	0	164	217	3386	61.19	0.7873	0.1609	0.0000	0.0518
	2	422	877	2	171	87	1559	28.17	0.2867	0.5958	0.0014	0.1162
	3	0	0	454	92	43	589	10.64	0.0000	0.0000	0.8315	0.1685
	4	0	0	0	0	0	0	00.00	0.0000	0.0000	0.0000	1.0000
							5534	100.00				

77-79	1	2	3	4	M	total	%	1	2	3	4
1	2133	540	0	131	130	2934	61.15	0.7607	0.1926	0.0000	0.0467
2	316	867	0	158	60	1401	29.20	0.2356	0.6465	0.0000	0.1178
3	0	0	359	73	27	459	09.57	0.0000	0.0000	0.8310	0.1690
4	0	0	0	0	0	4	00.00	0.0000	0.0000	0.0000	1.0000
						4798	100.00				

Another apparent contradiction exists between Census/Social Security death records, the RHS death records, and the death records independently compiled by Paul Taubman. The RHS date of death differs from Taubman's data in 36 cases, which in turn differs from the Census and Social Security death date (from the SSMBR tape) in 302 cases. Case-by-case cross checks resolved the discrepancies between Taubman's data and RHS, and cross-checks of Taubman's data with the Census data reveal that in 285 cases Taubman's data recorded the respondent as dead while Census and SSA had no record of death. Individual cross-checks reveal that Taubman's data are probably right in these cases. In fact one can identify at least 26 cases of apparently fraudulent behavior involving a surviving spouse who continued to collect both her and her husband's OASI benefits even though the husband had been deceased for several years![11] The final death data that I used to construct the health variable are Taubman's original data, edited in approximately 60 cases where case-by-case examinations revealed that either the RHS or SSMBR death date was correct.

I conclude this section with tables 5.2, 5.3, and 5.4 which present the estimates of the transition probabilities for health, disability, and death respectively. Each of the transition probabilities were specified to have linear-in-parameters binomial logit functional forms. Products of the

[11]Although the total number of cases seems small, think of the millions of unnecessary tax dollars spent if this error rate exists in the population at large.

estimated probability functions can be multiplied out to compute the estimated health transition matrix, $\hat{\pi}_h$. The interpretation of the estimation results has been listed in points 1, 2, and 3 of section 3 and will not be repeated here. However in order to get more intuition about how workers believe their health declines with age, I present figures 5.1 to 5.3, which show $\Pr\{h_{t+1}=1|a_t\}$, $\Pr\{h_{t+1}=3|a_t\}$, and $\Pr\{h_{t+1}=4|a_t\}$, respectively.

Table 5.2: Estimates of Health Transition Probability

Dependent variable $I\{h_{t+1}=1\}$

variable	parameter estimates	corrected std. error	uncorrected t-statistic	corrected t-statistic
$h_t=1,\ s_t=1$	-2.28309814	0.47881691	-4.80449217	-4.76820703
$h_t=1,\ s_t=2$	-2.15303544	0.49528616	-4.40304497	-4.34705351
$h_t=1,\ s_t=3$	-1.84775982	0.49618576	-3.76289898	-3.72392752
$h_t=2,\ s_t=1$	-0.13261981	0.48134990	-0.27822978	-0.27551644
$h_t=2,\ s_t=2$	-0.21572087	0.49269514	-0.44160374	-0.43783844
$h_t=2,\ s_t=3$	0.20683969	0.49534966	0.42169997	0.41756300
a_t	0.01278949	0.00778216	1.65308926	1.64343610
y_t	-0.05423888	0.00752685	-7.09584577	-7.20604944
y_t*y_t	0.00088587	0.00016736	5.00256272	5.29325400
$ms_t=2$	0.02113699	0.06228325	0.34705959	0.33936883
$ss_t\in\{1,2\}$	0.17005037	0.05776144	2.96758256	2.94401187

log likelihood	-8470.70743899	percent correctly predicted	78.88
grad*direc	1.97088242E-028	total observations	17536

Table 5.3: Estimates of Disability Hazard Function
Dependent Variable: $I\{h_{t+1}=3\}$

variable	parameter estimates	corrected std error	uncorrected t-statistic	corrected t-statistic
constant	-27.83091400	2.17638516	-9.18660618	-12.78767861
$h_t=1,\ s_t=1$	0.68277943	0.30811521	2.24293956	2.21598743
$h_t=1,\ s_t=3$	-0.39444705	0.36727688	-1.07755978	-1.07397734
$h_t=2,\ s_t=1$	-0.65686909	0.31441274	-2.10694736	-2.08919366
$h_t=2,\ s_t=2$	-0.73440502	0.42346985	-1.74560322	-1.73425573
$h_t=2,\ s_t=3$	-0.95738762	0.35761347	-2.69423827	-2.67715762
a_t	0.51877505	0.03481550	10.46765510	14.90069085
y_t	0.41976624	0.13932533	1.81436962	3.01284938
a_t*y_t	-0.00674512	0.00221385	-1.78433113	-3.04678841
$ms_t=2$	0.73731675	0.28755312	2.58747453	2.56410626

log likelihood	-1048.45610014	percent correctly predicted	98.72
grad*direc	6.68166236E-027	total observations	17763

Table 5.4: Estimates of Mortality Hazard Function
Dependent Variable: $I\{h_{t+1}=4\}$

variable	estimates	std error	t-stat u	t-stat c
$h_t=1,\ s_t=1$	3.02142632	0.21469528	14.74075784	14.07309152
$h_t=1,\ s_t=2$	1.67576602	0.46474387	3.67207026	3.60578400
$h_t=2,\ s_t=1$	2.47115478	0.34807814	7.22980179	7.09942529
$b_t=2,\ s_t=2$	2.35550786	1.00813087	2.33931275	2.33650999
$h_t=2,\ s_t=3$	0.20975890	0.35622039	0.62215231	0.58884584
$h_t=3,\ s_t=1$	-1.83437865	0.24681647	-9.08178222	-7.43215651
$h_t=3,\ s_t=2$	-0.93274864	0.27379854	-4.32434492	-3.40669698
$h_t=3,\ s_t=3$	-0.33218442	0.12672294	-3.01163747	-2.62134402
y_t	0.00670965	0.00435412	1.47047807	1.54098777
$a_t\in[0,60)$	2.15267752	0.13634172	16.57926607	15.78883916
$a_t\in[60,62)$	2.35948514	0.11708723	22.41837696	20.15151503
$a_t\in[62,65)$	2.71931875	0.10230694	31.08372566	26.58000347
$a_t\in[65,68)$	3.03076021	0.10767281	31.97119315	28.14787014
$a_t\in[68,71)$	2.92505558	0.11992181	26.94316060	24.39135665
$a_t\geq 71$	2.79733898	0.18388376	15.54234305	15.21253929
$ms_t=2$	-0.31691544	0.07285017	-4.21539501	-4.35023629
$h_t=1,\ s_t=1,\ ss_t\in\{1.2\}$	-3.15683370	0.21015059	-15.12909470	-15.02176920
$h_t=1,\ s_t=2,\ ss_t\in\{1,2\}$	-1.82248054	0.46700518	-3.92339216	-3.90248468
$h_t=1,\ s_t=3,\ ss_t\in\{1,2\}$	-0.27111012	0.10484321	-3.14802854	-2.58586240
$h_t=2,\ s_t=1,\ ss_t\in\{1,2\}$	-3.40066483	0.35127786	-9.67327528	-9.68084012
$h_t=2,\ s_t=2,\ ss_t\in\{1,2\}$	-3.14843443	1.01031812	-3.11117919	-3.11628027
$h_t=2,\ s_t=3,\ ss_t\in\{1,2\}$	-0.84830371	0.35539269	-2.49706044	-2.38694759

log likelihood -4713.91257216
grad*direc 2.01058679E-027

percent correctly predicted 93.97
number of observations 24233

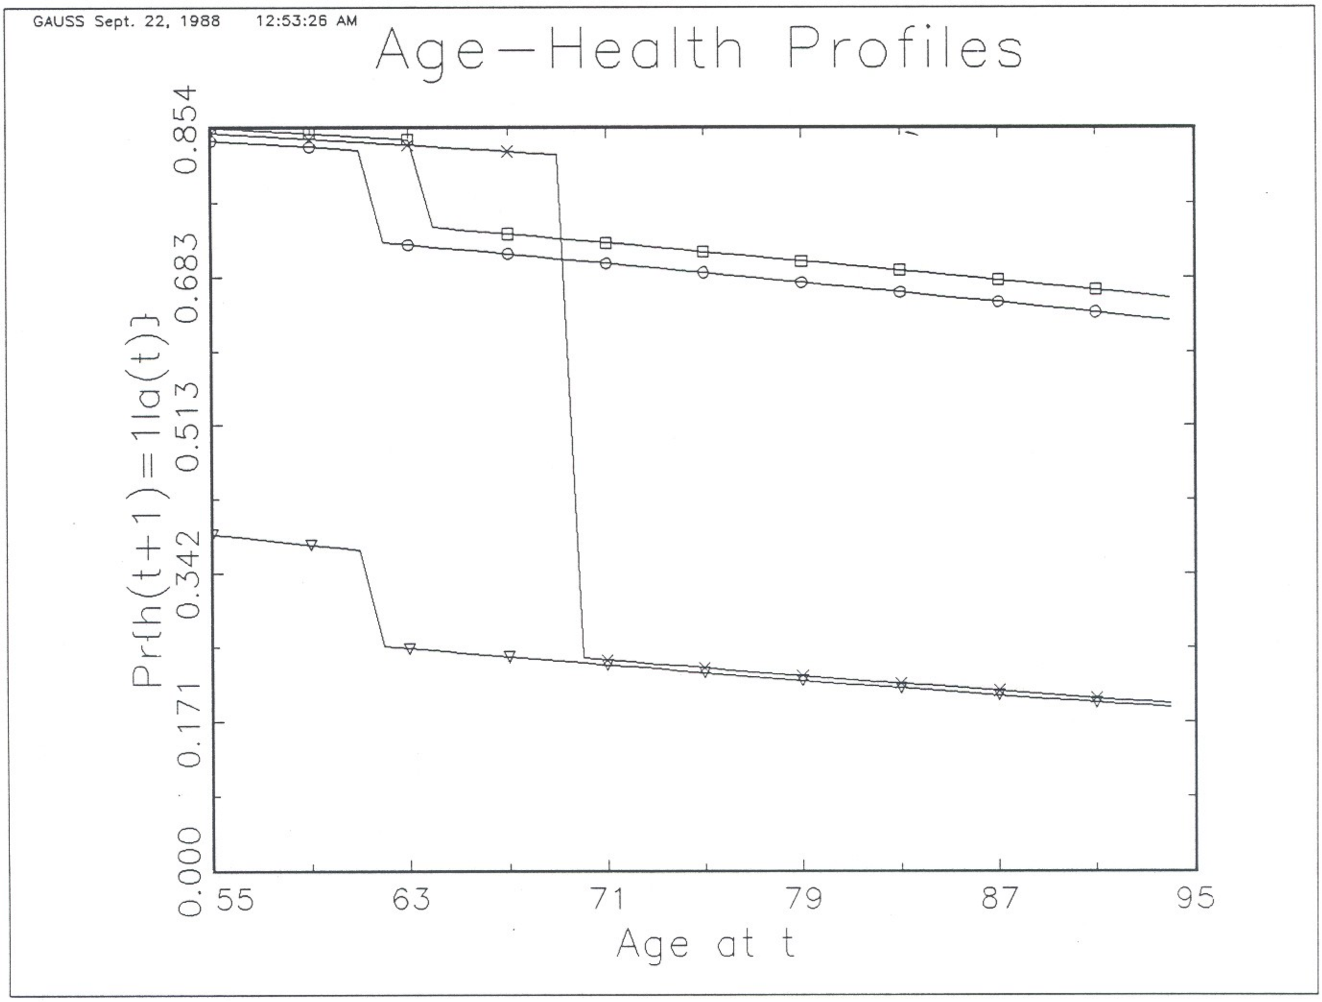
GAUSS Sept. 22, 1988 12:53:26 AM
Age-Health Profiles
Pr{h(t+1)=1|a(t)}
0.000
0.171
0.342
0.513
0.683
0.854
55
63
71
79
87
95
Age at t

Figure 5.1

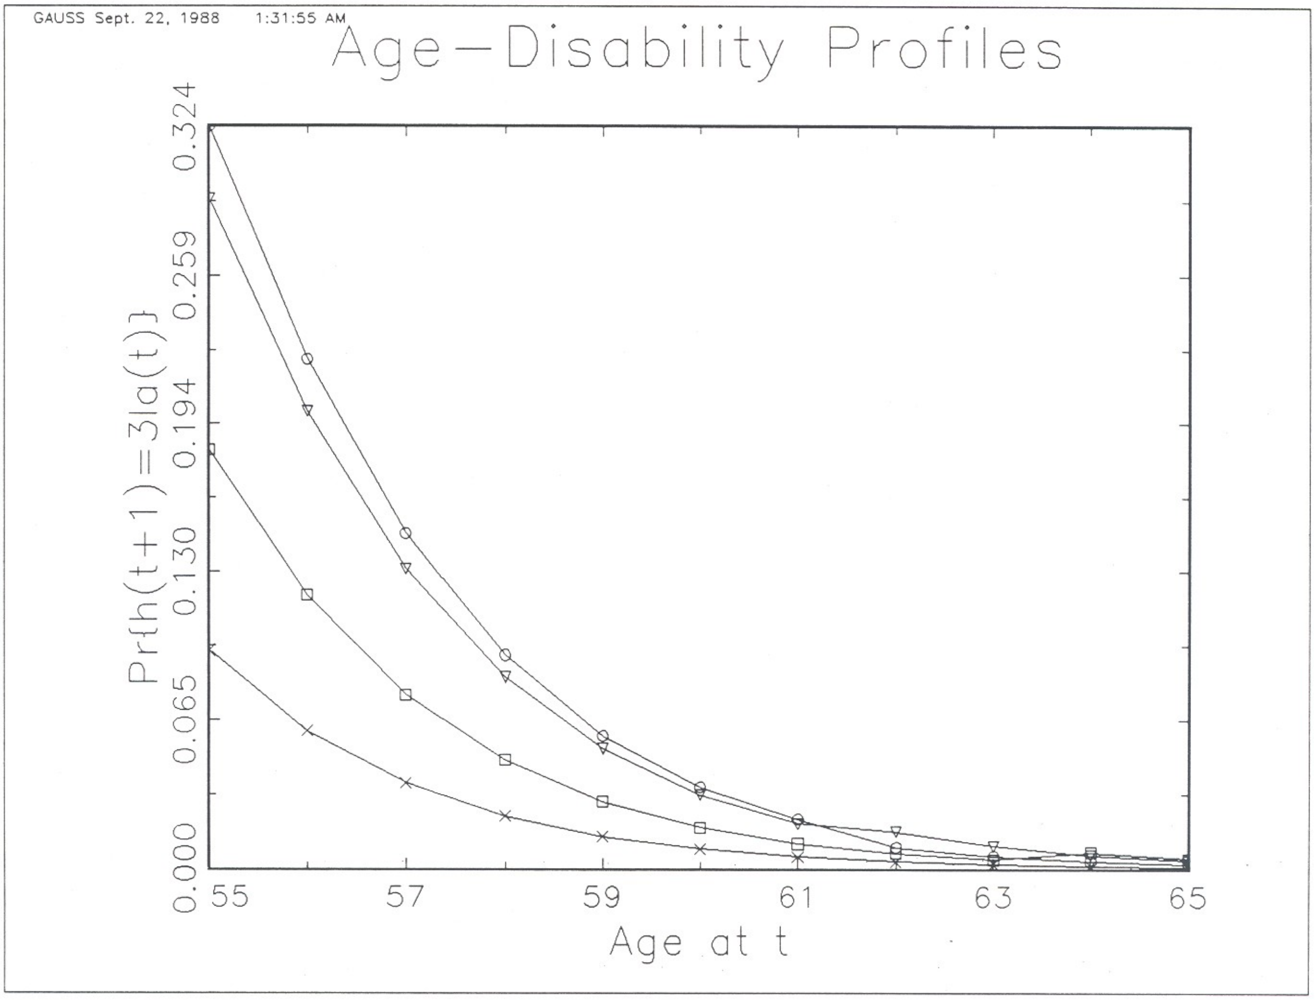

GAUSS Sept. 22, 1988 1:31:55 AM
Age-Disability Profiles
Pr{h(t+1)=3|a(t)}
0.000
0.065
0.130
0.194
0.259
0.324
55
57
59
61
63
65
Age at t

Figure 5.2

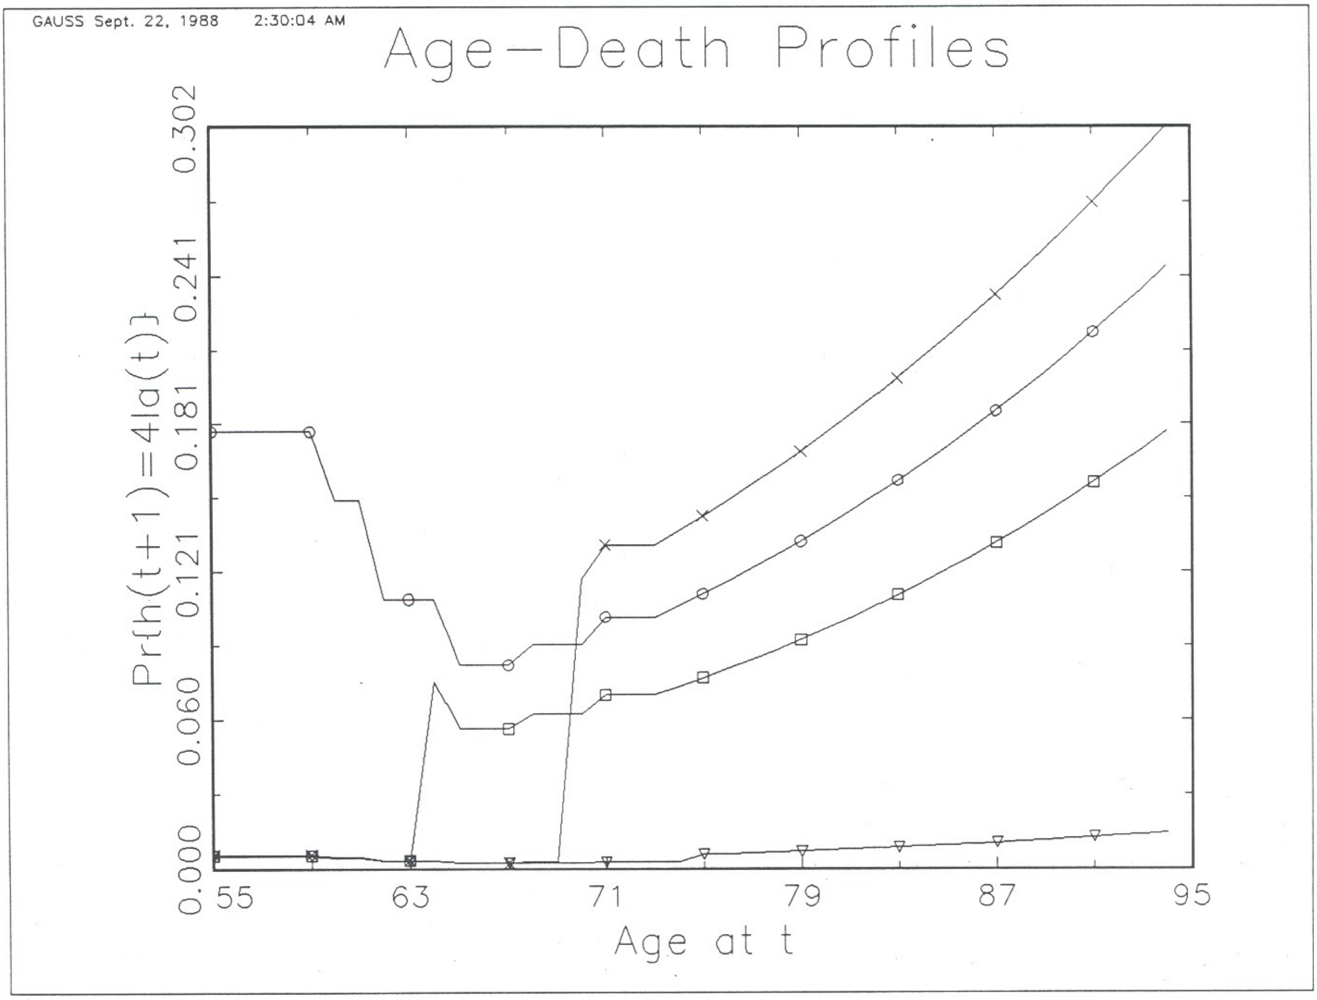
GAUSS Sept. 22, 1988 2:30:04 AM
Age-Death Profiles
Pr{h(t+1)=4|a(t)}
0.000
0.060
0.121
0.181
0.241
0.302
55
63
71
79
87
95
Age at t

Figure 5.3

Figure 5.1 shows the probability of remaining in good health as a function of age for four configurations of the remaining state variables. All four curves show that health declines with age, however changes in the other variables have a stronger impact on health than age alone. The top two curves (marked with circles and squares), represent the health expectations of workers who are already in good health, and whom retire at ages 62 and 65, respectively (the latter worker also has 10% higher income). The bottom curve represents the health expectations of a worker who is in poor health, $h_t=2$, and who retires at age 62. The remaining curve, marked with x's, shows the health expectations of a healthy worker who continued to work up until age 70, at which time he fell ill ($h_t=2$), quit his job, and began collecting OASI. The combination of all these events at age 70 produced a dramatic downturn in the worker's health expectations.

Figure 5.2 shows the probability of becoming disabled as a function of age. In this case age effects dominate, reflecting sharp declines in workers' incentives to incur the costs of applying for disability benefits as they approach the early retirement age, 62. The topmost curve corresponds to a low income married worker who is currently in poor health ($h_t=2$), while the lowest curve corresponds to a high income single worker who is in good health.

Finally, figure 5.3 plots the estimated death hazard function. As discussed in section 3, it was difficult to identify the independent effect of age on death rates. Both linear and quadratic specifications of age effects produced ultimately falling death hazards, a result I found implausible. Using age dummies I discovered the explanation: the age dummies reveal that workers' death rates decrease until age 67 after which they begin rising with age. However because the RHS surveyed men between 58 and 63 in 1969, the oldest possible age reached during the survey is 73. This implies that there are relatively few observations beyond age 67, so that both the linear and quadratic specifications attempted to fit the downward sloping part of the death hazard function from age 58 to 67, ignoring the upturn that occurred afterwards owing to a lack of observations. Unfortunately, while the age dummies allow the model to fit the data well, it implies that the risk of death is constant after the worker reaches his early 70's. Aggregate mortality statistics show (unconditional on health and employment status), that death rates increase with age, which implies that a model containing only age dummies from 58 to 73 will ultimately undepredict death rates. To correct this, I added a pure age-trend to the model in order to match the aggregate mortality statistics from ages 74 to 95.[12] Figure 5.3 (which incorporates the age-trend after age 73) displays mortality expectations for four different workers. The "v"-shaped curve marked with circles corresponds to a single, low income disabled worker. While his death rate is much higher than average, it shows significant improvement until age 67 after which it begins to steadily worsen. The bottom curve, marked with triangles, shows the mortality expectations of a high income "workaholic" who is in good health and who continued working full-time until his health deteriorated to $h_t=2$ at age 75, after which he started working part-time. In spite of his health problems, the workaholic never retired, in the sense of collecting OASI. The remaining two curves (marked with a "□" and "x", respectively) show the death rates of two average income workers who retire at 64 and 70, respectively. The latter worker retired at 70 owing to the fact that his wife died and his health

[12]The aggregate mortality rates were obtained from the Statistical Abstract of the U.S., (U.S. Government Printing Office, 1979). In future work I would like to formally incorporate auxiliary mortality data for very old men into a pooled maximum likelihood estimation of the death hazard model. A difficulty of this approach is the likely absence of associated health and employment status in any auxiliary data set. This will require me to "integrate out" these variables, which in turn requires further distributional assumptions on the cross-sectional distributions of health and employment status for very old men. Given these problems, I decided to use the short-cut described above.

deteriorated from h=1 to h=2; this explains the dramatic increase in his death rate.

There are two features of figure 5.3 that seem implausible: 1) the sharp v-shaped death hazard for the disabled worker, and 2) the significantly lower death rates for the high-income "workaholic" in comparison to the two average income workers who retired at 65 and 70. Looking back to the estimation results in table 5.4, it appears that these predictions result from the fact that Social Security recipients ($ss_t \in \{1,2\}$) have significantly higher death rates. As discussed in conclusion 1 of section 3, this is probably due to the fact that h_t does not capture all dimensions of health status. Workers who are in worse health are probably more likely to retire than healthy workers. However from the standpoint of the DP model the relation is necessarily treated as cause and effect: collecting OASI can be hazardous to your health. To avoid this problem I re-estimated the model without the ss_t interactions. While there was a significant drop in the likelihood (from -4714 to -5045), figure 5.4 shows that the resulting model seems to produce more reasonable predictions. In particular, the age effects now show a slightly increasing rather than decreasing hazard rate, and the gross disparities between the workaholic (who never collected OASI) and his average-income colleagues has disappeared. Based on these results, I have decided to exclude the ss_t interactions in the specification of the mortality hazard, even though they clearly improve the fit of the model.

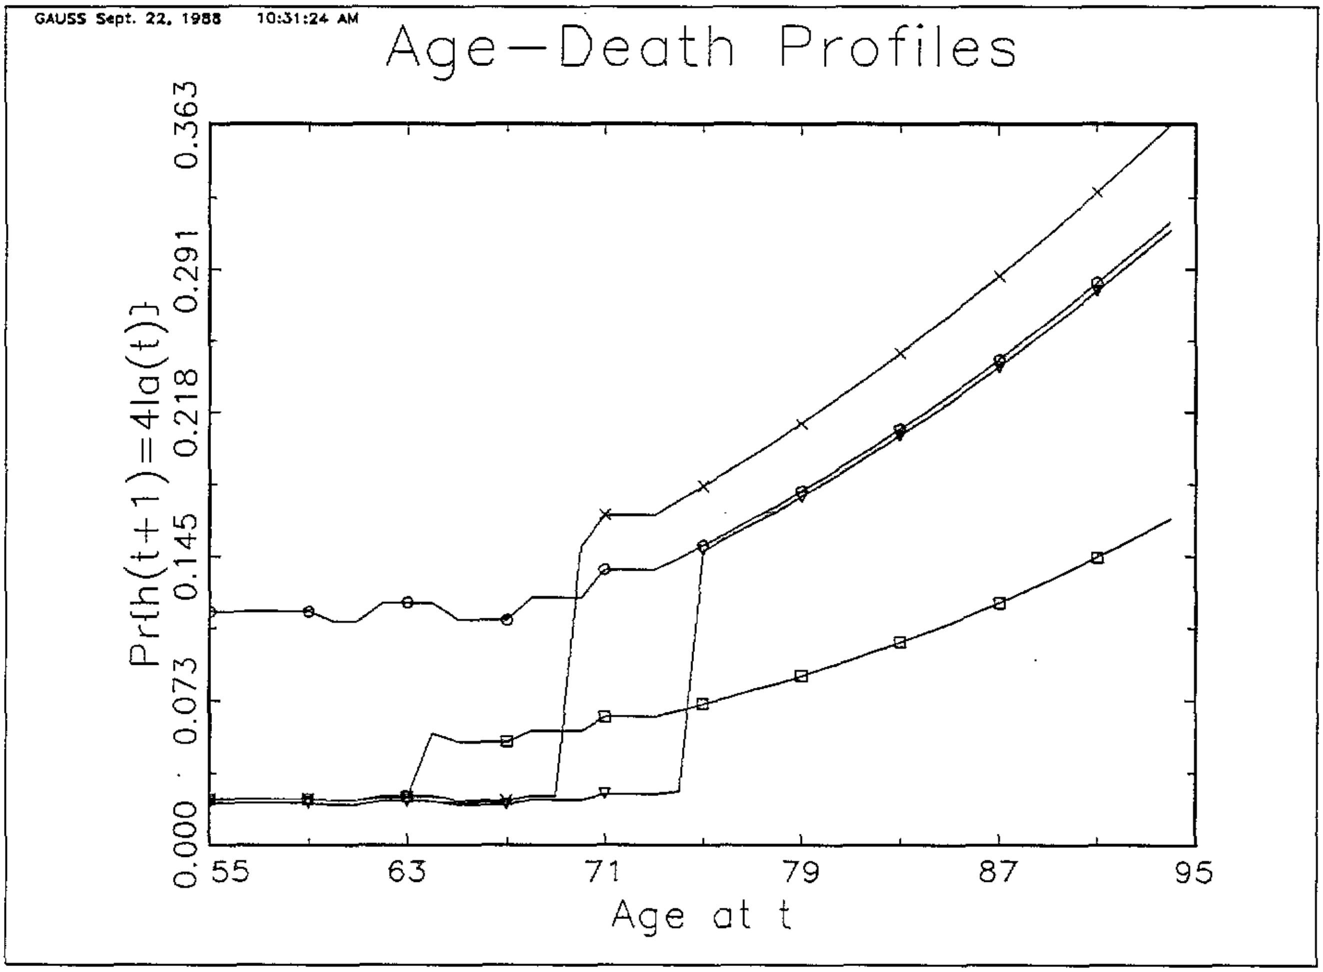
GAUSS Sept. 22, 1988 10:31:24 AM
Age–Death Profiles
Pr{h(t+1)=4|a(t)}
0.000
0.073
0.145
0.218
0.291
0.363
55
63
71
79
87
95
Age at t

Figure 5.4

6. **Employment Status**

Accurate classification of employment status e_t is the key to the entire undertaking: employment status is the most important variable affecting income and utility levels in the DP model, and is a crucial input into the income imputation routines that construct biannual income. They are also key inputs for the construction of biannual consumption expenditures in section 8. The RHS dataset allowed me to construct three independent measures of labor force status: self-classification of employment status (SE), instantaneous employment status (IE), and historical employment status (E). Each of the measures assume three values, 1=full-time, 2=part-time, and 3=not-employed. The SE variable was directly recorded in a trichotomous format from the survey question "do you consider yourself partly retired, completely retired, or not retired at all?". The IE measure was determined from the survey question, "how many hours per week do you usually work on your current job". Using this response I defined IE=1 if the worker worked more than 25 hours per week, IE=2 if the worker worked between 5 and 25 hours per week, and IE=3 if the worker was not currently employed or worked less than 5 hours per week. The historical employment status measure E is an annual measure based on the total number of hours worked in the preceding year. I defined E=1 if the respondent worked more than 1300 hours in the past year, E=2 if the respondent worked between 200 and 1300 hours, and E=3 otherwise. Because the worker might have had multiple jobs in the two years preceding the RHS interview, computation of total hours worked required direct reconstruction of the underlying continuous-time labor force histories from a battery of more than 130 questions in the "Work Experience" section of the RHS survey. Previous studies have used the IE and SE measures of employment status, probably because they were among the easiest variables to pull off the RHS tapes. Constructing retrospective labor force histories is a considerably more complicated undertaking due to the existence of complicated skip patterns in the survey questionnaire and the need to carefully account for the beginning and ending dates of jobs when there are multiple transitions within the interview frame. To my knowledge, this is the first study to construct complete labor force histories using the RHS data.

Table 6.1 presents aggregate employment distributions using each of the definitions of employment status. Although there are significant differences between the measures, all three confirm conclusion 1 of section 2 that the aggregate data show workers making a smooth transitions from work into retirement. The main differences are that SE appears to substantially over-estimate the occurrence of part-time work relative to the E and IE measures, and E appears to slightly over-estimate part-time work relative to the IE measure. The latter effect is to be expected from the nature of the definition of E: a worker who worked at a full-time job until mid-year and then retired would be classified as being in state 2 by the E measure and in state 3 by the IE measure.

Table 6.1: Cross-sectional Distributions of Measures of Employment Status

Historical Employment Status

	E68	E69	E70	E71	E72	E73	E74	E75	E76	E77	E78
1	70.77	72.06	60.70	54.38	39.89	33.59	23.12	19.49	13.83	12.69	10.10
2	9.03	9.15	12.37	12.81	14.99	15.55	15.85	14.95	14.07	13.96	13.81
3	20.20	18.80	26.93	32.81	45.12	50.86	61.03	65.56	72.10	73.35	76.08
Obs	8117	7379	7379	6837	6837	6392	6392	5871	5871	5415	5415

Instantaneous Employment Status

	IE69	IE71	IE73	IE75	IE77	IE79
1	74.10	61.04	39.51	23.65	15.91	12.15
2	3.71	5.10	7.08	8.90	10.39	10.73
3	22.19	33.86	53.41	67.44	73.70	77.12
Obs	8117	7434	6897	6392	5871	5415

Self-Reported Employment Search Decision

	SR69	SR71	SR73	SR75	SR77	SR79
1	72.20	55.86	36.41	21.46	13.46	10.29
2	4.99	11.93	12.50	15.35	13.06	12.11
3	22.71	32.20	51.09	63.19	73.48	77.60
Obs	7894	7434	6897	6392	5871	5415

Self-Employment Status

	SE69	SE71	SE73	SE75	SE77	SE79
1	77.42	59.49	36.08	20.64	11.72	9.27
2	8.12	12.53	16.26	18.40	18.70	17.59
3	14.46	27.98	47.65	60.97	69.58	73.15
Obs	8070	7431	6881	6387	5861	5407

The SE variable variable seems like the poorest candidate for use as a measure of employment status owing to the ambiguity of the term "retired". Some people may interpret being "retired" as quitting their career job and will report being fully or partly retired even though they are working full-time at a new job. Other people may interpret "retired" as meaning "are you working now?", and will report that they are not retired if they had quit their main career job but are currently working at a new full-time job. Still others may may report being partly retired even though they are not working because they like to think they have the virility to return to work at some unspecified future date. The latter problem seems to be reflected in table 6.1, which shows that the SE measure substantially overestimates the incidence of part-time work, sometimes as much as 200% in comparison to the IE measure. I decided not to use the SE measure because of the problems of ambiguity and subjective interpretation, and also, for reasons I elaborate below, because SE is an "instantaneous" measure that doesn't correspond well to the time intervals of the DP model.

The instantaneous employment status variable completely avoids the subjective definition of the concept "retired". Like SE, IE has a high response rate and is easy to pull from the tapes. However it too has certain drawbacks from the standpoint of estimating the DP model. Since I am using a relatively coarse two year time interval (for computational reasons discussed in section 2), the instantaneous IE measure would not provide a good measure of the worker's actual state over the whole time period. In principle, workers may have changed jobs many times in the two year time interval or may have only recently retired, so there may be only a weak association between IE and the respondent's actual labor force status over the last two years.

From the standpoint of the discrete-time DP model, the most appropriate measure of labor force status is the historical employment status measure, E. The main drawbacks of this measure are 1) it requires the worker to recall his employment history (which may be especially difficult in the cases where the worker had multiple job transitions), and 2) since E is a flow measure, it may overestimate the occurrence of part-time work by misclassifying full-time workers who retire in mid-year. Table 6.2 sheds some light on the last problem by summarizing the distribution of employment histories (using the E measure of e_t) over the 11 years of the RHS survey. To keep the table manageable, the 11^4 possible employment sequences have been "collapsed", i.e. the sequence (1,1,1,2,2,3,3,3,3,M,M) (where "M" represents missing data) is classified as a "1-2-3" sequence. The first thing to notice is that in contrast to the aggregate employment statistics, the individual employment sequences are far from smooth: only 18% the sample is observed to gradually phase out of work in a "1-2-3" employment sequence. If I reclassify all "1-2-3" sequences with only 1 intervening year in state 2 as actually being a mis-classified "1-3" sequence, then only 3% of the sample is observed to follow a smooth employment transition; a plurality of the sample, 33%, are observed to follow a discontinuous "1-3" sequence. Another 28% of the sample have complex "non-monotonic" employment histories, with periods of unemployment followed by subsequent re-employment. Of course many of the "1" and "1-2" sequences may actually be part of an ultimate "1-2-3" sequence; however since these sequences only account for 14% and 4% of the sample, respectively, the basic story would remain unchanged.

For comparison, table 6.2 presents the distribution of employment sequences for the IE and SE measures of e_t, and also an annual measure of e_t similar to the E measure but computed from the NLS data by Berkovec and Stern, 1988. Notice that in all of the tables, only 3-4% of all workers are observed to follow a "1-2-3" sequence. The NLS data show a somewhat higher fraction of workers following a "1" sequence, but this is to be expected given that the NLS sample follows a younger group of men who were initially aged 45-59 in the first year of the survey, 1966[13]. Based on the comparison of the employment measures presented in table 6.2, I decided to reclassify all "1-2-3" employment sequences with only 1 intervening year in state 2 as a "1-3" sequence by reassigning the state $e_t=2$ as either $e_t=1$ or $e_t=3$ depending on whether or not hours worked in that year are greater than 1,000.

[13]The NLS contained an enriched sample of black respondents, who are presumably more likely to be unemployed. Apparently the effect of a more youthful sample in the NLS dominated the effect of a larger proportion of blacks, leading to the discrepancies noted above.

Table 6.2: Distribution of Employment Sequences

Measure of Employment State	Sequence	number of cases		percent of cases	
E	1.....	1174		14.44	
	2.....	91		1.12	
	3.....	1033		12.70	
	1-3...	1488	(2700)	18.30	(33.21)[14]
	2-3...	255		3.14	
	1-2...	306		3.76	
	1-2-3.	1450	(238)	17.83	(2.93)
	others	2334		28.70	
	Total	8131		100.00	
IE	1.....	1321		16.25	
	2.....	28		0.34	
	3.....	1337		16.44	
	1-3...	3269		40.20	
	2-3...	112		1.38	
	1-2...	276		3.39	
	1-2-3.	308		3.79	
	others	1480		18.20	
	Total	8131		100.00	
SE	1.....	1239		15.24	
	2.....	131		1.61	
	3.....	897		11.03	
	1-3...	2642		32.49	
	2-3...	298		3.66	
	1-2-3.	748		9.20	
	others	1575		19.37	
	Total	8131		100.00	

[14] Numbers in parentheses obtained by reclassifying all "1-2-3" sequences with only one intervening year in state e_t=2 as a "1-3" sequence. See page 28 for further explanation.

NLS[15]	1.....	585	23.43
	2.....	13	0.52
	3.....	187	7.49
	1-3...	1052	42.13
	2-3...	29	1.16
	1-2...	90	3.60
	1-2-3.	89	3.56
	others	452	18.10
	Total	2497	100.00

Overall, table 6.2 casts doubt on the notion that most workers gradually phase out of their full-time jobs through a spell of "partial retirement", a view promoted by Gustman and Steinmeier (1984) and suggested from casual interpretation of the macro data in table 6.1. Even if we counted all "1-2" and "1" sequences as forming part of an eventual "1-2-3" sequence, the number of "smooth" employment transitions would be at most 23%. In reality, most of the "1" sequences will form part of an eventual "1-3" sequence, and a large fraction of the "3" sequences are actually left-truncated "1-3" sequences. If I count all these sequences as "1-3" sequences, I obtain an estimate that approximately 75% of all retirement sequences involve discontinuous transitions from a full-time job into unemployment. Table 6.2 also shows that a significant fraction of the sample, over 18%, follow "non-monotonic" sequences involving some form of "unretirement", i.e. a return to employment from a state of unemployment or partial employment. Table 6.3 provides more

detail on the structure of the non-montonic employment sequences for the E, IE, and SE measures of employment status. The structure of these transitions

[15]An annual measure of employment status similar to E. This measure was constructed by Berkovec and Stern (1988) who wrote more than 2,000 lines of Fortran code to accurately follow NLS skip patterns to accurately reconstruct the employment histories.

is extremely complex, as you can see from table 6.3. The most common non-monotonic sequences are "3-1-3", "1-3-2", "1-3-1", "1-3-2-3", and "1-3-1-3". Even though a majority of workers follow the "1-3" sequence, the traditional approach to modelling retirement behavior as an ex ante choice of a fixed retirement date after which the worker ceases to work is incapable of explaining the labor force history of 20% of the sample.

Table 6.3: Distribution of Non-Monotonic Employment Sequences

	IE			E			SE		
Sequence	#	% of 1480	% of 8131	#	% of 2334	% of 8131	#	% of 1575	% of 8131
1313	99	6.69	1.22	36	1.54	.44	39	2.48	.48
1312	28	1.89	.34	20	.86	.25	15	.95	.18
1213	47	3.18	.58	75	3.21	.92	46	2.92	.57
1212	33	2.23	.41	83	3.56	1.02	71	4.51	.87
1231	10	.68	.12	21	.90	.26	24	1.52	.30
1232	28	1.89	.34	101	4.33	1.24	85	5.40	1.05
1321	30	2.03	.37	30	1.29	.37	18	1.14	.22
1323	100	6.76	1.23	115	4.93	1.41	158	10.03	1.94
2131	1	.07	.01	2	.09	.02	0	.00	.00
2132	1	.07	.01	11	.47	.14	7	.44	.09
3131	12	.81	.15	6	.26	.07	0	.00	.00
3132	8	.54	.10	9	.39	.11	5	.32	.06
3231	2	.14	.02	3	.13	.04	0	.00	.00
3232	6	.41	.07	10	.43	.12	8	.51	.10
3121	5	.34	.06	5	.21	.06	2	.13	.02
3123	17	1.15	.21	90	3.86	1.11	8	.51	.10
3213	2	.14	.02	15	.64	.18	3	.19	.04
3212	0	.00	.00	7	.30	.09	3	.19	.04
2121	3	.20	.04	10	.43	.12	2	.13	.02
2123	15	1.01	.18	77	3.30	.95	26	1.65	.32
2321	2	.14	.02	3	.13	.04	1	.06	.01
2323	7	.47	.09	34	1.46	.42	34	2.16	.42
2312	1	.07	.01	0	.00	.00	1	.06	.01
2313	3	.20	.04	3	.13	.04	1	.06	.01
131	122	8.24	1.50	35	1.50	.43	78	4.95	.96
121	105	7.09	1.29	122	5.23	1.50	109	6.92	1.34
132	201	13.58	2.47	137	5.87	1.68	204	12.95	2.51
321	12	.81	.15	12	.51	.15	7	.44	.09
312	19	1.28	.23	25	1.07	.31	8	.51	.10
323	57	3.85	.70	80	3.43	.98	68	4.32	.84
313	176	11.89	2.16	102	4.37	1.25	35	2.22	.43
213	34	2.30	.42	75	3.21	.92	46	2.92	.57
231	2	.14	.02	5	.21	.06	4	.25	.05
232	5	.34	.06	22	.94	.27	24	1.52	.30
212	5	.34	.06	20	.86	.25	25	1.59	.31

31	77	5.20	.95	51	2.19	.63	60	3.81	.74
32	40	2.70	.49	45	1.93	.55	43	2.73	.53
21	9	.61	.11	30	1.29	.37	24	1.52	.30
others	156	10.54	1.92	807	34.58	9.92	283	17.97	3.48
Total	1480	100.00	18.20	2334	100.00	28.70	1575	100.00	19.37

The above discussion suggests the possibility that the discretization of labor force status into only three states variable could seriously misrepresent the labor force participation decision. Other researchers (e.g. MaCurdy, 1983) have suggested that the labor force participation decision can be modelled as a continuous choice variable, say, as choice of annual hours of work. There are strong practical reasons for maintaining this viewpoint: an interior solution allows one to derive stochastic Euler orthogonality conditions that permit estimation of identified parameters by the method of moments (Hansen, 1982). Figure 6.1, which displays the distribution of annual hours of work over the period 1968 to 1978, provides convincing evidence against this view. The distribution has almost all of its mass at two spikes, one at 0 and the other at 2000. The distributions are almost excessively concentrated at the two spikes, suggesting a systematic tendency of

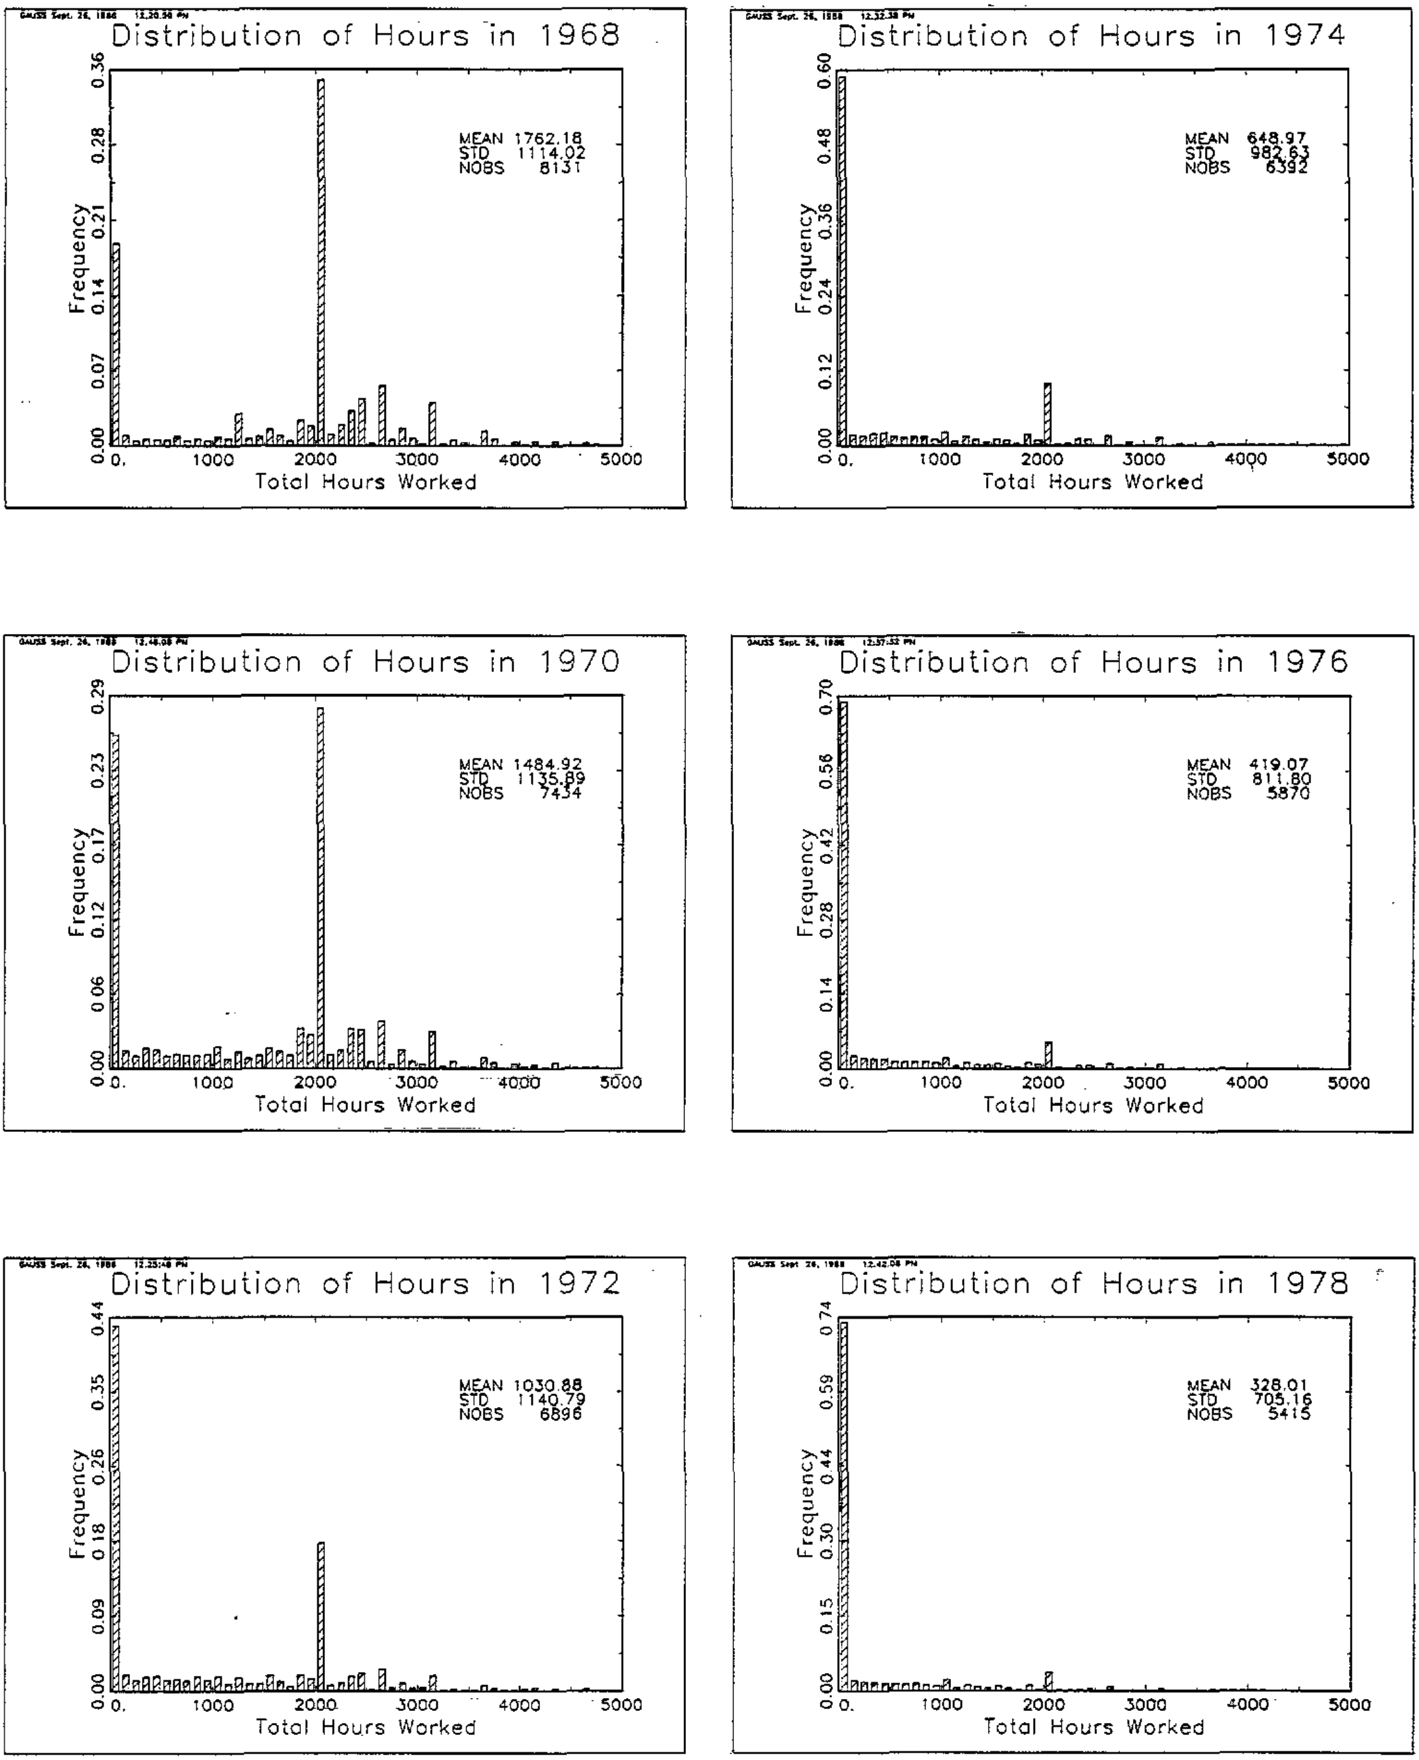

Distribution of Hours in 1968
MEAN 1762.18
STD 1114.02
NOBS 8131
Frequency
Total Hours Worked
Distribution of Hours in 1974
MEAN 648.97
STD 982.63
NOBS 6392
Frequency
Total Hours Worked
Distribution of Hours in 1970
MEAN 1484.92
STD 1135.89
NOBS 7434
Frequency
Total Hours Worked
Distribution of Hours in 1976
MEAN 419.07
STD 811.80
NOBS 5870
Frequency
Total Hours Worked
Distribution of Hours in 1972
MEAN 1030.88
STD 1140.79
NOBS 6896
Frequency
Total Hours Worked
Distribution of Hours in 1978
MEAN 328.01
STD 705.16
NOBS 5415
Frequency
Total Hours Worked

Figure 6.1

respondents to round their reponses (e.g. 40 hours/week, 50 weeks/year). Nevertheless, I believe the distributions provide solid evidence against the notion that annual hours of work is best modelled as a continuous choice variable that satsifies an interior first order condition. This notion is also supported by the work of Gustman and Steinmeier (1983), (1984) who present convincing evidence of widespread minimum hours restrictions and signficant wage cuts associated with transitions from full-time work to part-time work. The combination of these constraints and the Social Security "earnings test" are probably the key factors that lead the majority of workers to follow a "bang-bang" work/no work decision rule. Figure 6.1 also shows that the definition of the E variable is robust to fairly large changes in the cutoffs defining the three employment states: there is a small amount of probability mass uniformly distributed between 0 and 2000 hours of work, so that changes in the cutoff in this range will not significantly alter the distribution of e_t.

I conclude this section by presenting markov transition probability matrices for the IE and SE measures of e_t in table 6.4 and the E measure in table 6.5. These are "uncontrolled" transition probabilities because I haven't conditioned on a measure of the respondent's search decision, s_t. Nevertheless, the resulting transition matrices are quite illuminating. The matrices show a clear pattern of age-effects: for example table 6.4 shows that the probability of re-employment (i.e. a transition from e_t=3 to e_{t+1}=1), is 16% in 1969, but falls to 2% by 1977, and the probability of remaining fully employed declines from 75% in 1969 to 55% in 1977. Interestingly, the probability of retiring from a full or part time job peaks at approximately 40% between 1971 and 1973, declining to 30% in 1977.

Table 6.4: Transition Matrices for IE and SE Measures of Employment Status

IE69 to IE71 (8117 obs)

75.46	4.12	20.41
29.64	37.50	32.86
16.45	2.81	80.74

SE69 to SE71 (8070 obs)

73.61	10.70	15.69
19.49	41.53	38.98
6.39	6.39	87.23

IE71 to IE73 (7434 obs)

59.23	6.36	34.41
19.15	38.87	41.97
5.51	3.44	91.05

SE71 to SE73 (7431 obs)

55.51	14.51	29.98
12.62	44.74	42.64
3.68	7.15	89.16

IE73 to IE75 (6897 obs)

52.39	8.93	38.68
14.63	46.29	39.08
3.17	3.82	93.02

SE73 to SE75 (6881 obs)

50.17	17.39	32.44
8.60	54.06	37.33
2.17	6.56	91.27

IE75 to IE77 (6392 obs)

52.95	11.32	35.73
14.63	51.03	34.33
2.85	4.55	92.60

SE75 to SE77 (6387 obs)

45.19	23.05	31.77
6.81	54.22	38.96
1.78	6.13	92.09

IE77 to IE79 (5871 obs)

54.67	15.34	29.99
17.44	48.36	34.20
2.09	4.23	93.68

SE77 to SE79 (5861 obs)

51.85	21.99	26.16
9.92	56.36	33.72
2.01	6.10	91.89

Table 6.5: Markov Transition Matrices for Historical Employment Status

E68 to E69 (seam: 8117 obs)

91.45	3.07	5.48
30.12	42.51	27.37
22.61	7.86	69.54

E69 to E70

82.50	6.94	10.56
17.87	58.12	24.01
1.69	1.62	96.68

E70 to E71 (seam: 7379 obs)

82.41	4.65	12.94
29.41	50.87	19.71
4.78	6.65	88.57

E71 to E72

71.44	10.07	18.49
11.25	58.71	30.05
0.66	2.28	97.06

E72 to E73 (seam: 6837 obs)

74.14	7.52	18.33
21.05	56.57	22.39
2.86	6.66	90.47

E73 to E74

65.16	14.22	20.62
8.97	61.33	29.70
0.36	1.92	97.72

E74 to E75 (seam: 6392 obs)

69.71	9.71	20.57
13.70	58.52	27.78
2.19	4.92	92.90

E75 to E76

63.24	15.34	21.42
8.23	64.56	27.22
0.46	2.38	97.16

E76 to E77 (seam: 5871 obs)

75.73	9.10	15.17
10.82	64.11	25.07
1.17	4.89	93.94

E77 to E78

65.52	22.17	12.30
10.43	67.35	22.22
0.33	2.56	97.12

A strange pattern appears in the historical employment state transition matrices in table 6.5. Notice how the transition matrices appear to cycle in two-year intervals: for example, the (1,1) elements appear significantly higher in even numbered years, while the (3,3) elements appear significantly higher in the odd-numbered years. For a long time I was convinced that these regular fluctuations had to be an artifact of my FORTRAN code for processing the observations. I labored for many weeks to make ever more detailed and accurate correspondence with the survey questionnaire but had no success in eliminating the strange fluctuations in the transition matrices. Only recently have I become aware of work by Daniel Hill (1988) that has convinced me that the fluctuations are not artifacts of my computer programs, but rather are symptoms of a systematic reponse error problem known as the *seam problem.* The seam problem arises from the way the RHS collects data on retrospective labor force history in successive two year survey frames. Each of the odd-numbered survey years represents a seam, and the survey questionnaire required the respondent to recall his labor force history in the two year survey frame prior to the interview. It appears that while respondents offer an internally consistent view of the preceding two years, their view of history changes between survey dates in a way that generates inconsistent labor force transitions across seams. For example, to compute the across-seam transition probability matrix from 68 to 69 I needed data from 2 different surveys: the 1969 survey gave me retrospective data on labor force states in 68 and the 1971 survey gave me retrospective data on labor force states in 69. On the other hand the between-seam transition probability matrix from 69 to 70 was computed entirely from retrospective data obtained at the 1971 interview. The pattern of fluctuations in the transition matrices indicates that men in state $e_t=1$ are more likely to remain in state 1 for across-seam transitions than for between-seam transitions, whereas men in state $e_t=3$ are less likely to remain in state 3 for across-seam transitions than for between-seam transitions.

I have recomputed table 6.5 using the flag variables to eliminate observations that showed any evidence of internally inconsistent responses. While the sample sizes were significantly reduced, the seam problem persisted. Although an analysis of the perceptual/psychological factors underlying the seam problem is beyond the scope of this paper, it appears that by using between-seam transitions based on data from a single survey frame one is much more likely to obtain consistent transition probabilities. Indeed, looking at the between-seam transition matrices in table 6.5 one can see that they change in a sensible way over time, with no suspicious patterns indicative of further inconsistencies. In particular, while the transition matrices do not closely match the IE or SE transition matrices (the latter two are two year transition matrices while the E transition matrix is for one year intervals), the matrices follow the same general pattern as the IE and SE transition matrices, namely a probability of re-employment and continued employment that gradually declines over time. I conclude that the seam problem is sufficiently severe to make it inadvisable to build a DP model based on annual data even though such a model is superior from a theoretical viewpoint since it has "finer grain" and thus suffers less from problems of time aggregation. Instead I will focus on constructing of model of biannual transitions, using consistent data on employment transitions between seams rather than across seams.

7. Job Search Decision

The DP model requires a control variable s_t that represents the respondent's labor force search/participation decision. In a discrete-time model the agent is in labor force state e_t at time t, and conditional on e_t and his search decision s_t, he makes a transition to a new labor force state e_{t+1} at time t+1. Thus, the DP model gives an employed worker the option of quitting (e_t=1 or e_t=2, and s_t=3), and an unemployed worker the option of returning to work (e_t=3, and s_t=1 or s_t=2). Unfortunately, while it is convenient to trichotomize s_t into 3 values (1=search for full-time job, 2=search for part-time job, or 3=quit the labor force) the "true" search decision s_t is essentially a latent variable; a complicated, possibly multidimensional function of the variety and intensity of the worker's search activities over the period. The RHS has three possible variables from which to construct a measure of s_t

SR : self-reported planned hours of work in the year following the survey,

NE : actual hours worked in the year following the survey,

e_{t+1}: actual hours worked in the second year following the survey.

The latter measure corresponds to a "perfect control" DP model where workers' search decisions are successful with probability 1. I find this latter measure implausible given the well-known labor market problems of older workers, and will focus hereafter on the other two measures.

Using the SR and and NE measures, I constructed a trichotomous estimate of s_t using the same cutoffs that I used to construct the e_t variable described in section 6. Table 6.1 summarizes aggregate distribution of the self-reported measure of s_t. This measure follows very much the same trends as the E, SE, and IE measures of e_t: a gradual phase-out from full-employment into unemployment. The NE measure of s_t is recorded in the odd-year columns of the E distribution at the top of table 6.1. At least on the aggregate level, the two measures appear to track each other fairly closely.

To get a better handle on the issue of which measure s_t better approximates the underlying latent employment search decision, I computed the controlled transition probability matrices which predict the probability of e_{t+1} conditional on e_t and s_t. Table 7.1 presents the controlled transition matrices using the E measure for e_t and the SR measure for s_t. These matrices show a very weak relation between employment search decisions and ex post realized employment states. If control were perfect, the transition matrix should have 1's in the column corresponding to the value s_t assumes. However in table 7.1 we see that under the SR measure control is highly imperfect. For example, a full-time worker who reported an intention to quit working in 1969 still has a 25% chance of remaining at work in 1971. A worker who had a full-time job in 1968 and who reported an intention to start working part-time in 1969 has only a 20% chance of actually realizing his intentions by 1970. An unemployed worker in 1974 who reports the intention to return to work full-time in 1975 only has a 13% chance of actually being employed in 1976. Thus the SR measure of s_t leads to a DP where control is too *imperfect*, in the sense that there is little correspondence between employment search decisions and subsequent labor market outcomes.

Table 7.1: Controlled Transition Probabilites
SR Measure of Job Search Variable, s_t

E68 to E70 given SR69=1
(5707 observations)

80.82	7.08	12.10
57.74	20.08	22.18
50.23	17.37	32.39

E68 to E70 given SR69=2
(394 observations)

33.10	13.38	53.52
24.14	50.57	25.29
2.50	20.00	77.50

E68 to E70 given SR69=3
(1793 observations)

25.63	12.18	62.18
13.66	20.22	66.12
14.94	5.38	79.68

E70 to E72 given SR71=1
(4150 observations)

71.30	9.41	19.29
53.63	25.14	21.23
28.95	28.95	42.11

E70 to E72 given SR71=2
(884 observations)

3.89	20.44	75.67
13.94	53.94	32.12
4.95	33.66	61.39

E70 to E72 given SR71=3
(2345 observations)

0.0458	0.0634	0.8908
0.0833	0.1667	0.7500
0.0211	0.0380	0.9409

E72 to E74 given SR73=1
(2508 observations)

60.01	13.61	26.38
39.53	34.11	26.36
26.92	28.21	44.87

E72 to E74 given SR73=2
(859 observations)

11.67	20.56	67.78
14.19	54.62	31.18
4.17	35.12	60.71

E72 to E74 given SR73=3
(3470 observations)

3.33	6.00	90.67
5.26	19.74	75.00
1.58	4.29	94.13

E74 to E76 given SR75=1
(1372 observations)

58.14	14.38	27.49
30.84	36.45	32.71
12.50	16.07	71.43

E74 to E76 given SR75=2
(981 observations)

7.01	20.38	72.61
9.95	55.50	34.55
5.64	22.56	71.79

E74 to E76 given SR75=3
(4039 observations)

1.46	10.95	87.59
3.08	16.92	80.00
1.55	4.22	94.22

Table 7.2: Controlled Markov Transition Probabilities

NE Measure of Employment Search Decision, s_t

E68 to E70 given E69=1
(5348 observations)

83.06	6.30	10.64
76.14	15.23	8.63
78.05	11.28	10.67

E68 to E70 given E69=2
(554 observations)

13.58	58.02	28.40
22.66	57.55	19.78
12.28	59.65	28.07

E68 to E70 given E69=3
(1477 observations)

3.11	1.73	95.16
1.12	3.91	94.97
1.39	1.19	97.42

E72 to E74 given E73=1
(2183 observations)

66.82	12.11	21.07
47.77	35.67	16.56
61.18	21.18	17.65

E72 to E74 given E73=2
(817 observations)

6.66	53.81	39.59
11.37	60.66	27.96
6.57	69.19	24.24

E72 to E74 given E73=3
(3335 observations)

.83	1.04	98.13
1.20	9.58	89.22
.22	1.60	98.18

E70 to E72 given E71=1
(3767 observations)

72.44	8.97	18.60
55.14	25.41	19.46
67.39	19.57	13.04

E70 to E72 given E71=2
(645 observations)

9.64	54.31	36.04
14.06	56.25	29.69
6.25	71.09	22.66

E70 to E72 given E71=3
(2376 observations)

1.09	2.74	96.17
4.03	12.90	83.06
.23	1.35	98.42

E74 to E76 given E75=1
(1167 observations)

66.39	12.91	20.70
37.84	38.74	23.42
60.00	12.50	27.50

E74 to E76 given E75=2
(790 observations)

3.68	49.26	47.06
9.92	66.88	23.21
7.22	70.00	22.78

E74 to E76 given E75=3
(3914 observations)

1.04	4.51	94.44
2.22	8.44	89.33
.29	1.79	97.91

Table 7.2 presents controlled transition probabilities for the E measure of e_t and the NE measure of s_t. Comparing tables 7.1 and 7.2 we can see that while the NE measure of s_t does reflect imperfect control, the relation between s_t and e_{t+1} is much stronger than for the SR measure of s_t. For example, consider the probability that a worker who intends to quit his full-time job is successful, (i.e. the transition from $e_t=1$ to $e_{t+1}=3$ given $s_t=3$). In 1968 the NE measure gives a 95% chance that the decision will be realized compared to only 62% for the SR measure of s_t. In the case of an unemployed worker who intends to return to work, the data for 1974 show that according to the NE measure of s_t the worker will have a 60% chance of success compared to only a 13% chance for the SR measure of s_t. The very weak correspondence between the SR measure of s_t and subsequent employment outcomes e_{t+1} may be an indication of the fact that "talk is cheap": it is one thing to say you intend to remain employed or return to work, but quite another thing to actually go out and do it. The NE measure is a compromise between the perfect control model implied by the e_{t+1} measure of s_t and the highly imperfect control model implied by the SR measure of s_t. Because a strong, interpretable relationship between search decisions and subsequent employment states is key to obtaining a sensible DP solution, I have adopted the NE measure of s_t for use in estimating the DP model.

Tables 7.3 and 7.4 present the maximum likelihood estimates of the controlled transition probabilities using the E measure of e_t and the NE measure of s_t. The estimates correspond to the component π_e in the decomposition of π given in (3.2). The probabilities were estimated using a linear-in-parameters specification of a trinomial logit model of the probability that e_{t+1} assumes the three values $\{1,2,3\}$. Table 7.3 presents the parameter estimates corresponding to the event $I\{e_{t+1}=1\}$ (full-time work), while table 7.4 presents the parameter estimates corresponding to the event

$I\{e_{t+1}=3\}$ (unemployment).[16] The interpretation of the estimation results in has already been summarized in conclusion 5 of section 3 and will not be repeated here. Instead I conclude this section with table 7.5 which presents the implied transition matrices for various configurations of the state and control variables.

[16]Since probabilities sum to one, it is not necessary to present parameter estimates for the event $I\{e_{t+1}=2\}$. In other words, identification of the parameters required me to normalize the parameters corresponding to the event $I\{e_{t+1}=2\}$ to zero.

Table 7.3: Estimates of Employment Status Transition Probability

Dependent Variable: $I\{e_{t+1}=1\}$

variable	parameter estimates	corrected std error	uncorrected t-statistic	corrected t-statistic
$e_t=1,\ s_t=1$	3.64047276	0.39767737	9.45528680	9.15433717
$e_t=2,\ s_t=1$	2.50066683	0.41218308	6.23215652	6.06688373
$e_t=3,\ s_t=1$	2.96647111	0.42303720	7.19813528	7.01231742
$e_t=1,\ s_t=2$	-0.92193667	0.44636590	-2.09671348	-2.06542812
$e_t=2,\ s_t=2$	-0.22332551	0.43614838	-0.52611010	-0.51204021
$e_t=3,\ s_t=2$	-1.12311922	0.47166279	-2.44907306	-2.38119107
$e_t=1,\ s_t=3$	1.66201979	0.76052861	2.24386341	2.18534817
$e_t=2,\ s_t=3$	0.85628243	0.74420975	1.16201613	1.15059287
$e_t=3,\ s_t=3$	1.42660101	0.71859281	2.01265302	1.98527037
$h_t=1,\ h_{t+1}=1$	0.04182958	0.09095360	0.45864771	0.45990021
$h_t=1,\ h_{t+1}=2$	-0.09248139	0.11701242	-0.79525628	-0.79035528
$h_t=1,\ h_{t+1}=3$	-0.86195143	0.33062624	-2.69698907	-2.60702671
$h_t=2,\ h_{t+1}=1$	-0.01050426	0.13169805	-0.08102400	-0.07976021
$h_t=2,\ h_{t+1}=3$	-0.50116877	0.37841526	-1.34720139	-1.32438835
$h_t=3,\ h_{t+1}=3$	-0.82519273	0.33886383	-2.36096966	-2.43517503
$a_t\in[0,60)$	0.46894279	0.14771651	3.16557318	3.17461331
$a_t\in[60,62)$	0.37064936	0.12279250	3.04499184	3.01850159
$a_t\in[62,65)$	0.16940925	0.10848216	1.58353029	1.56163226
$a_t\in[65,68)$	-0.05630496	0.11120326	-0.50766144	-0.50632476
$y_t\in[0,4)$	-0.52162418	0.19991055	-2.67741816	-2.60928788
$y_t\in[4,7)$	-0.55721500	0.17250413	-3.23698482	-3.23015450
$y_t\in[7,10)$	-0.45827949	0.17025525	-2.68879219	-2.69172016
$y_t\in[10,13)$	-0.43955168	0.17563187	-2.49588461	-2.50268740
$y_t\in[13,21)$	-0.34465953	0.17880938	-1.93446250	-1.92752493
$y_t\in[21,31)$	-0.15707978	0.21297779	-0.73426161	-0.73754066

$ms_t=2,\ ms_{t+1}=2$	-0.10073099	0.34258973	-0.31131817	-0.29402804
$ms_t=1,\ ms_{t+1}=2$	-0.17166326	0.38546353	-0.46566385	-0.44534242
$ms_t=1,\ ms_{t+1}=1$	-0.21573198	0.33203186	-0.69209921	-0.64973277
$ss_t\in\{1,2\},\ s_t=1$	-1.50403909	0.09419539	-15.13231598	-15.96722584
$ss_t\in\{1,2\},\ s_t=2$	-0.76541385	0.22012150	-3.51805524	-3.47723356
$ss_t\in\{1,2\},\ s_t=3$	-2.09065804	0.60984560	-3.42590823	-3.42817599

log likelihood -9154.95723806
grad*direc 7.13617145E-028

percent correctly predicted 82.20
number of observations 18778

Table 7.4: Estimates of Employment Status Transition Probability
Dependent Variable: I{e_{t+1}=3}

variable	parameter estimates	corrected std error	uncorrected t-statistic	corrected t-statistic
e_t=1, s_t=1	-0.52562870	0.42176812	-1.23834758	-1.24625041
e_t=2, s_t=1	-1.19341755	0.44611839	-2.66229609	-2.67511399
e_t=3, s_t=1	-0.81920118	0.46039631	-1.77001849	-1.77933914
e_t=1, s_t=2	-1.52714046	0.51255236	-3.04668919	-2.97948184
e_t=2, s_t=2	-1.87406038	0.51180484	-3.74289769	-3.66166992
e_t=3, s_t=2	-1.89114189	0.51811559	-3.73146644	-3.65003857
e_t=1, s_t=3	2.75006382	0.67224902	4.05697020	4.09084092
e_t=2, s_t=3	1.70999512	0.66075817	2.57366878	2.58792882
e_t=3, s_t=3	3.49735104	0.65323246	5.29150772	5.35391499
h_t=1, h_{t+1}=1	-0.44576119	0.08643213	-5.06447314	-5.15735518
h_t=1, h_{t+1}=2	-0.02814231	0.11184866	-0.25140546	-0.25161061
h_t=1, h_{t+1}=3	0.73477410	0.29518668	2.54913310	2.48918451
h_t=2, h_{t+1}=1	-0.38372744	0.12629923	-3.00704767	-3.03824049
h_t=2, h_{t+1}=3	0.77222603	0.34956602	2.23698564	2.20909922
h_t=3, h_{t+1}=3	0.48026267	0.18967800	2.25908999	2.53198929
$a_t \in [0,60)$	-0.12892234	0.16617958	-0.74499112	-0.77580132
$a_t \in [60,62)$	0.24293611	0.12221154	1.98347970	1.98783278
$a_t \in [62,65)$	0.41766273	0.09695995	4.22818248	4.30758012
$a_t \in [65,68)$	0.08319764	0.09619687	0.85051452	0.86486851
$y_t \in [0,4)$	0.27927137	0.21734863	1.24522584	1.28490051
$y_t \in [4,7)$	0.20449627	0.20307223	0.97685822	1.00701247
$y_t \in [7,10)$	0.32859160	0.20261126	1.57389213	1.62178349
$y_t \in [10,13)$	0.44050091	0.20823785	2.05267004	2.11537390
$y_t \in [13,21)$	0.53355789	0.21214523	2.45499095	2.51505955

$y_t \in [21,31)$	0.32767887	0.25431324	1.26370451	1.28848528
$ms_t=2,\ ms_{t+1}=2$	-0.27499860	0.33575242	-0.83147118	-0.81905171
$ms_t=1,\ ms_{t+1}=2$	-0.42346123	0.38439440	-1.12042959	-1.10163214
$ms_t=1,\ ms_{t+1}=1$	-0.41809414	0.32507531	-1.30897165	-1.28614548
$ss_t \in \{1,2\},\ s_t=1$	1.03210995	0.15168050	6.52729131	6.80449988
$ss_t \in \{1,2\},\ s_t=2$	1.08034934	0.31234593	3.47130972	3.45882320
$ss_t \in \{1,2\},\ s_t=3$	0.46282193	0.52779083	0.85695738	0.87690407

log likelihood	-9154.95723806	percent correctly predicted	82.20
grad*direc	7.13617145E-028	total observations	18778

Table 7.5: Estimated Employment Transition Probabilites, $\hat{\pi}_e$

Case I: a_t=58, ms_t=2, ms_{t+1}=2, h_t=1, h_{t+1}=1, y_t=8, ss_t=0

s_t=1			s_t=2			s_t=3		
96.1268	2.9696	0.9036	23.3057	68.9840	7.7102	33.0889	7.3921	59.5190
89.9567	8.6876	1.3557	38.6368	56.8702	4.4930	34.2108	17.1072	48.6820
93.0782	5.6418	1.2799	20.4052	73.8584	5.7363	16.4247	4.6433	78.9320

Case II: a_t=65, ms_t=2, ms_{t+1}=2, h_t=2, h_{t+1}=2, y_t=5, ss_t=2

s_t=1			s_t=2			s_t=3		
67.5569	15.6263	16.8167	6.2574	66.2566	27.4860	1.9057	5.7311	92.3632
47.1212	34.0734	18.8054	12.8050	67.4239	19.7711	2.1704	14.6103	83.2192
55.0056	24.9637	20.0307	5.6559	73.2332	21.1109	0.7446	2.8338	96.4215

Case III: a_t=75, ms_t=1, ms_{t+1}=1, h_t=1, h_{t+1}=1, y_t=20, ss_t=2

s_t=1			s_t=2			s_t=3		
69.6938	13.8217	16.4845	7.0164	63.6986	29.2850	2.0149	5.1953	92.7899
50.0203	31.0116	18.9681	14.3233	64.6629	21.0138	2.3146	13.3588	84.3266
57.6323	22.4258	19.9419	6.3905	70.9448	22.6648	0.7855	2.5631	96.6513

Case IV: a_t=85, ms_t=2, ms_{t+1}=2, h_t=1, h_{t+1}=2, y_t=12, ss_t=2

s_t=1			s_t=2			s_t=3		
69.5368	13.5160	16.9471	7.0431	62.6676	30.2893	1.9616	4.9572	93.0812
50.0408	30.4068	19.5524	14.4168	63.7895	21.7937	2.2627	12.7990	84.9383
57.5405	21.9444	20.5151	6.4371	70.0393	23.5236	0.7635	2.4417	96.7948

Case V: a_t=60, ms_t=2, ms_{t+1}=2, h_t=3, h_{t+1}=3, y_t=6, ss_t=1

s_t=1			s_t=2			s_t=3		
74.5445	11.2507	14.2048	8.8719	61.2958	29.8323	2.4960	4.8979	92.6061
56.2644	26.5466	17.1890	17.8012	61.1587	21.0402	2.8782	12.6420	84.4798
63.4964	18.8030	17.7006	8.1262	68.6548	23.2190	0.9746	2.4201	96.6054

8. Income, Wealth and Consumption

Next to employment status, the most important state variables of the DP model are income y_t and wealth w_t. The RHS records detailed information on assets and debts in each of the odd-numbered survey years, 1969-1979, as well as detailed information on the components of income in the preceding even-numbered years, 1968-1978. Consumption, c_t, is a latent control variable, essentially a time aggregration of thousands of invidividual discrete buy/no buy consumption decisions over the two year period. My strategy is to use the budget equation $w_{t+1}=w_t+y_t-c_t$ to infer consumption expenditures from measurements of w_{t+1}, w_t and y_t. There are two obstacles to this approach: 1) the RHS has no data on capital gains income, and 2) the RHS only records income in even-numbered years. Thus, capital gains and income in odd-numbered years must be imputed. A key to accurate income imputations is the use of the retrospective labor force histories used to construct the e_t state variable.

I initially tried to impute the missing income values by regressing income in even-numbered years on variables available in both even and odd-numbered years. Among the variables available in both even and odd years were the SSER earnings records (up until 1974) and the SSMBR OASDI benefit data (from 1969 to 1978). Despite the inclusion of these variables and retrospective data on total hours worked in odd-numbered years, the fits of the income regressions were not very impressive with R^2 values of 60%. Using the estimated regressions to fill in the missing income values produced intuitively unreasonable results, generating wide swings in income which occasionally turned negative or exceeded reasonable values.

An approach that turned out to work much better was something I call "full information interpolation". One can divide income into four sources: 1) wage income, 2) OASDI income, 3) unemployment insurance, and 4) other income. Since I have OASDI income in all years, that variable does not need to be imputed.[17] In addition, since other income is predominantly asset and pension income which is largely independent of labor force participation, I obtained an estimate for category 4) by simply averaging observed other income in adjacent even-numbered years. The problem thus reduced to computing wage income and unemployment compensation. Using the retrospective employment histories, I obtained an estimate of total hours worked in each year. Dividing hours worked into observed wage income I obtained a wage rate which I used to compute total wage income in odd-numbered years. If there was evidence that the worker had become involuntarily unemployed during the period, I imputed unemployment compensation as well. The resulting interpolation estimates appeared much more reasonable than the regression-based imputations. In particular there were far fewer wild swings in income, very few excessively large values, and no negative income values. Figure 8.1 plots the imputed and reported income distributions for the six year period 1973 to 1977. There is evidently little difference between the imputed and reported income distributions; both have the characteristic log-normal shape. There is a noticeable leftward shift in the distribution over time as more and more workers withdraw from the labor force. This shift is not as pronounced as it might be due to the replacement of wage income by OASDI and pension receipts. If I were to plot wage distributions only, the leftward shift would be much more pronounced.

[17] I substituted actual OASDI benefits from the SSMBR rather than reported OASDI benefits to calculate total income in even numbered years.

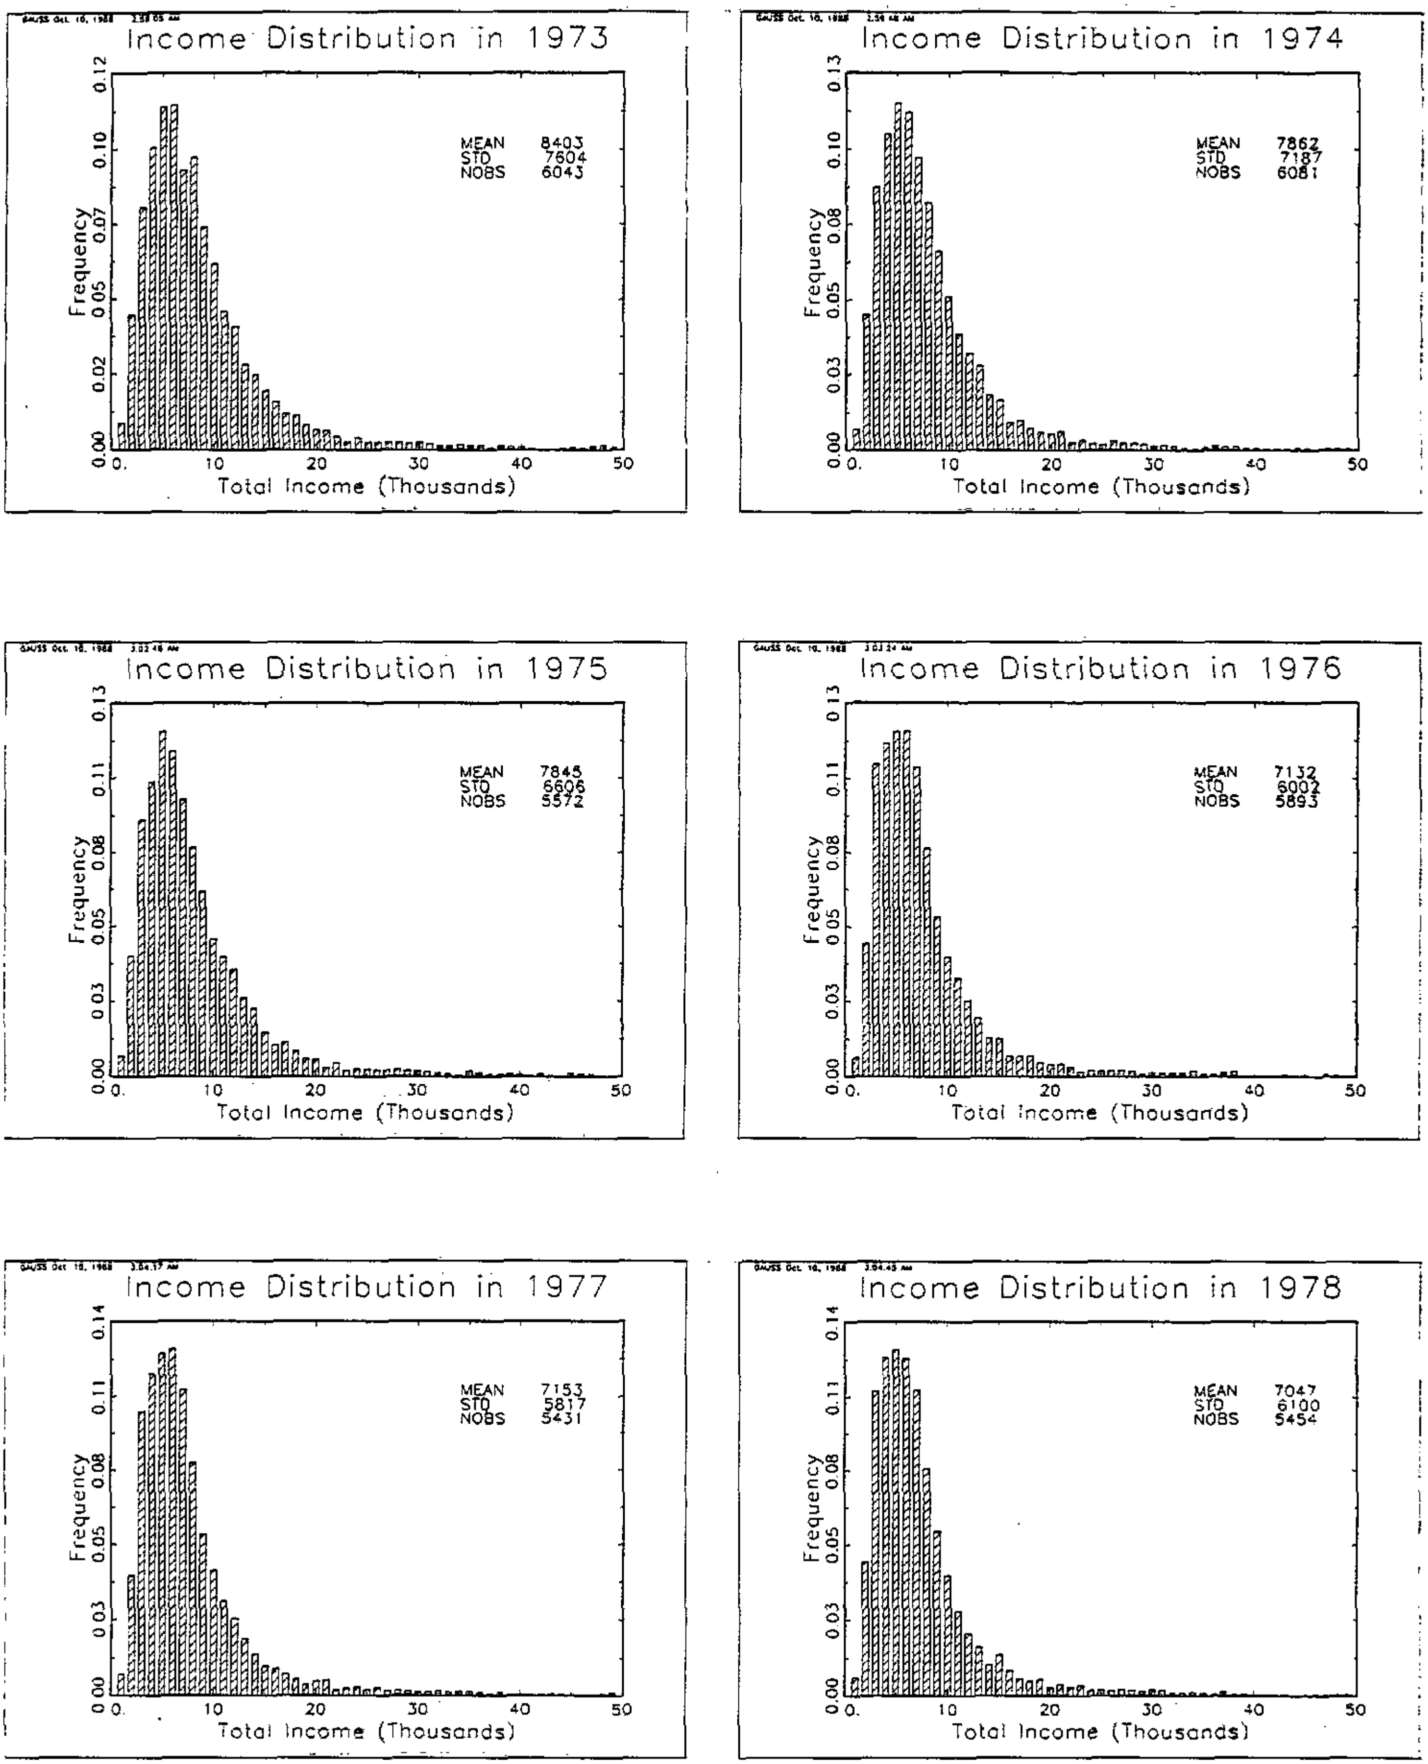
Income Distribution in 1973
Frequency
Total Income (Thousands)
MEAN 8403
STD 7604
NOBS 6043
Income Distribution in 1974
Frequency
Total Income (Thousands)
MEAN 7862
STD 7187
NOBS 6081
Income Distribution in 1975
Frequency
Total Income (Thousands)
MEAN 7845
STD 6606
NOBS 5572
Income Distribution in 1976
Frequency
Total Income (Thousands)
MEAN 7132
STD 6002
NOBS 5893
Income Distribution in 1977
Frequency
Total Income (Thousands)
MEAN 7153
STD 5817
NOBS 5431
Income Distribution in 1978
Frequency
Total Income (Thousands)
MEAN 7047
STD 6100
NOBS 5454

Figure 8.1

The existence of the seam problem in the employment data discussed in section 5 lead me to suspect the possiblity that these inconsistencies might have contaminated the imputed income data. To see whether there is any evidence of this, I plotted the distributions of income changes in figure 8.2. These distributions show no evidence of the seam problem, perhaps because wage income became an increasingly less important source of income over the survey, and because the SSER earnings records and the SSMBR OASDI benefit data allowed me to get relatively accurate measurements of the main components of income for the majority of the sample. In any event, I conclude that my income imputations appear to be fairly reliable measures of actual income.

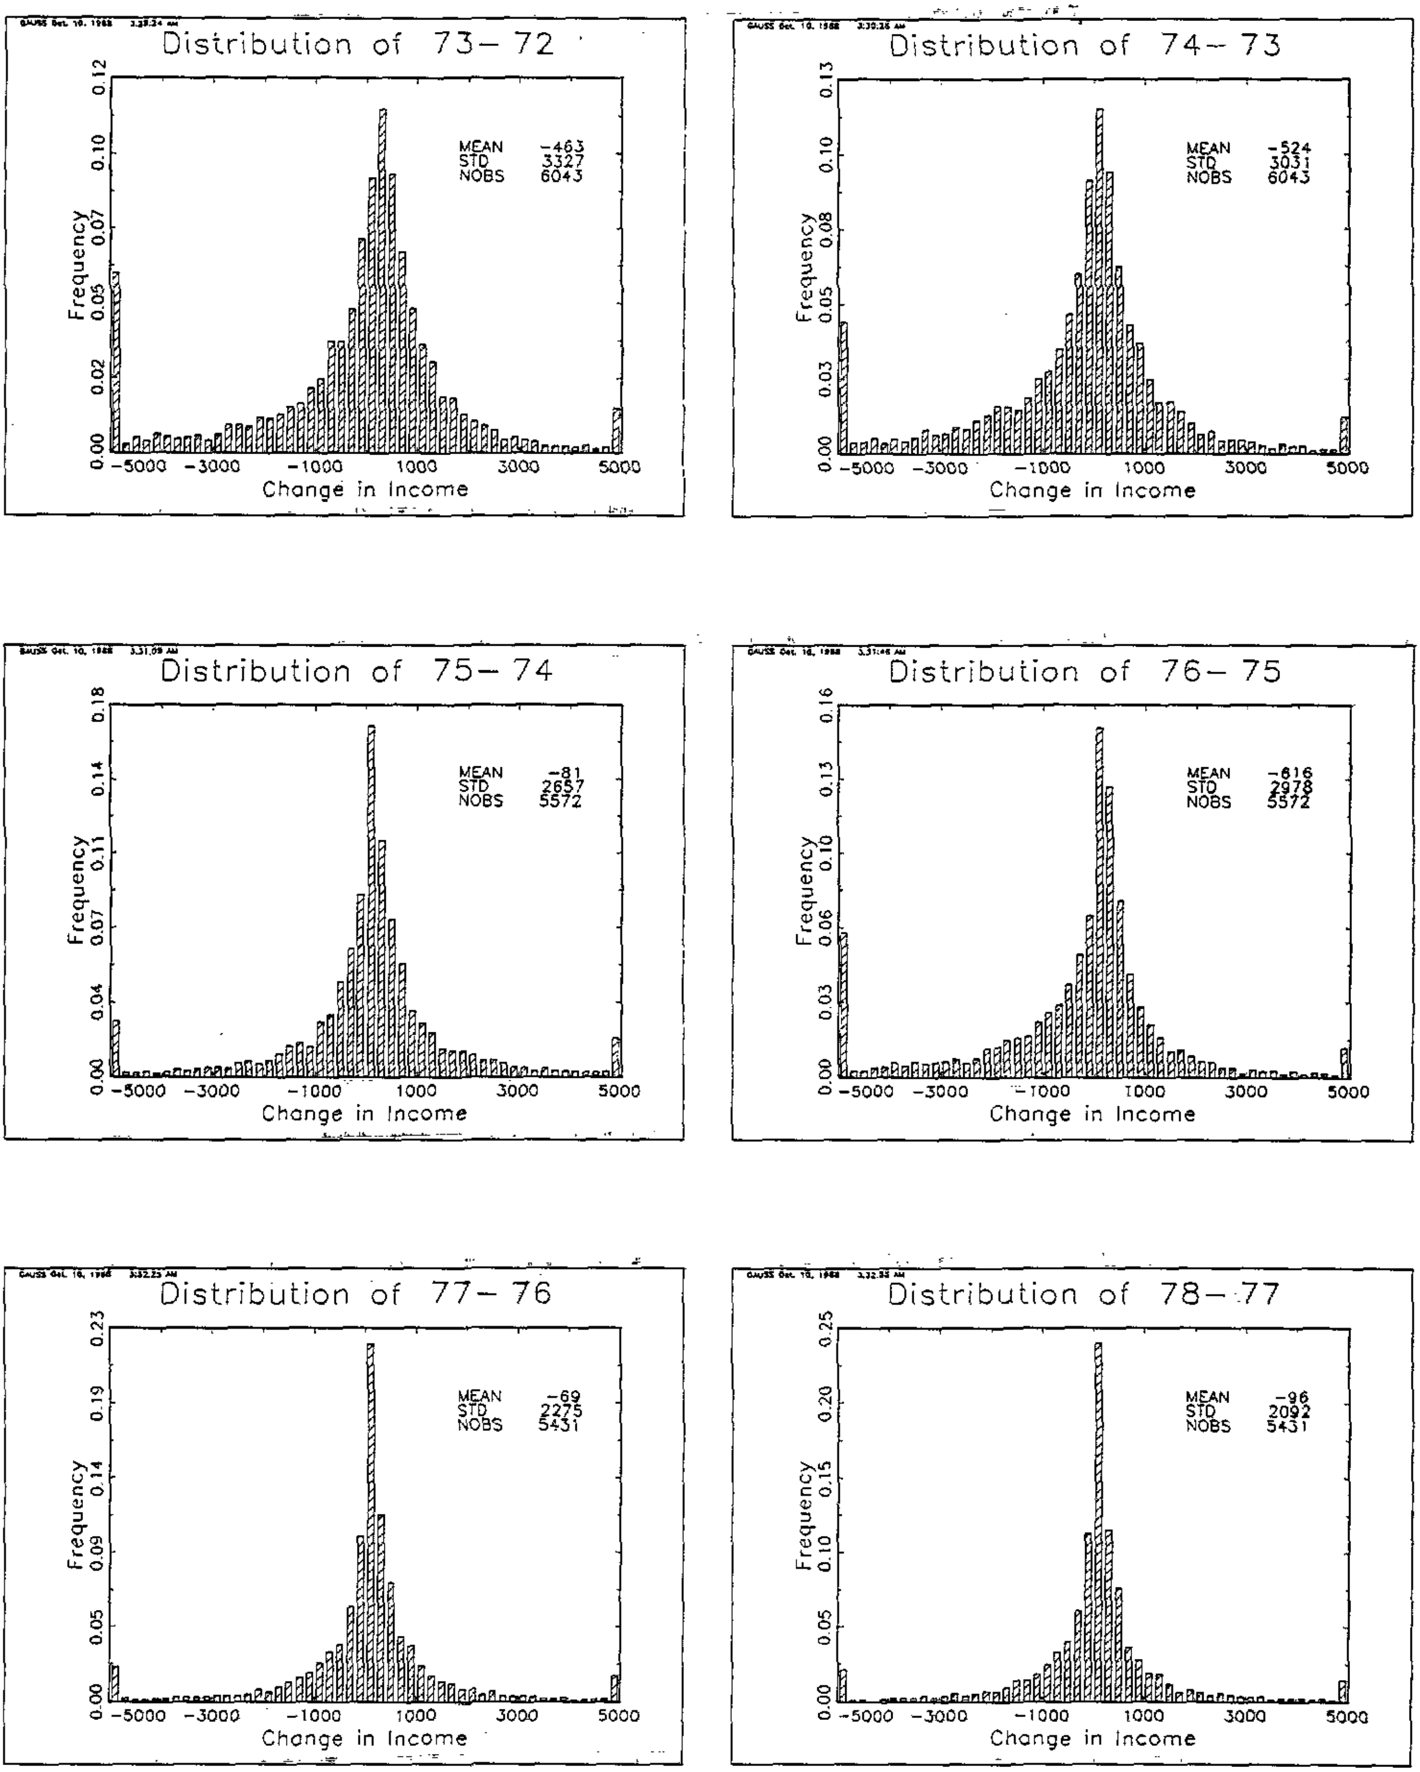
Distribution of 73- 72
MEAN -463
STD 3327
NOBS 6043
Frequency
Change in Income
Distribution of 74- 73
MEAN -524
STD 3031
NOBS 6043
Frequency
Change in Income
Distribution of 75- 74
MEAN -81
STD 2657
NOBS 5572
Frequency
Change in Income
Distribution of 76- 75
MEAN -616
STD 2978
NOBS 5572
Frequency
Change in Income
Distribution of 77- 76
MEAN -69
STD 2275
NOBS 5431
Frequency
Change in Income
Distribution of 78- 77
MEAN -96
STD 2092
NOBS 5431
Frequency
Change in Income

Figure 8.2

Having said this is not to deny the existence of systematic response errors in reported wage and OASDI benefits. For example section 4 discussed the widespread under-reporting of Social Security disability benefits. To assess how accurately respondents reported their income, I used the SSER and SSMBR datasets to compare reported and actual earnings and OASDI benefits. Because of the Social Security maximum earnings limitation, respondents had a clear incentive to under-report their wage earnings since the survey was conducted by SSA. On the other hand OASDI benefits not enter into the "earnings test", so there is no obvious incentive to under-report these receipts. Figure 8.3 presents the distribution of the percentage difference between reported wages and SSER earnings in 1970, and figure 8.4 presents the distribution of percentage response error in total OASDI benefit in 1974. The figures show no obvious evidence of systematic under-reporting, although each contains spikes at -100% indicating a non-negligible fraction of respondents falsely reporting that they had no wage or OASDI income. On the basis of these comparisons, I set flags indicating the degree of accuracy of the respondent's reports of his wage and OASDI benefits. I then used these flags in the construction of the sample boolean to screen out questionable respondents.

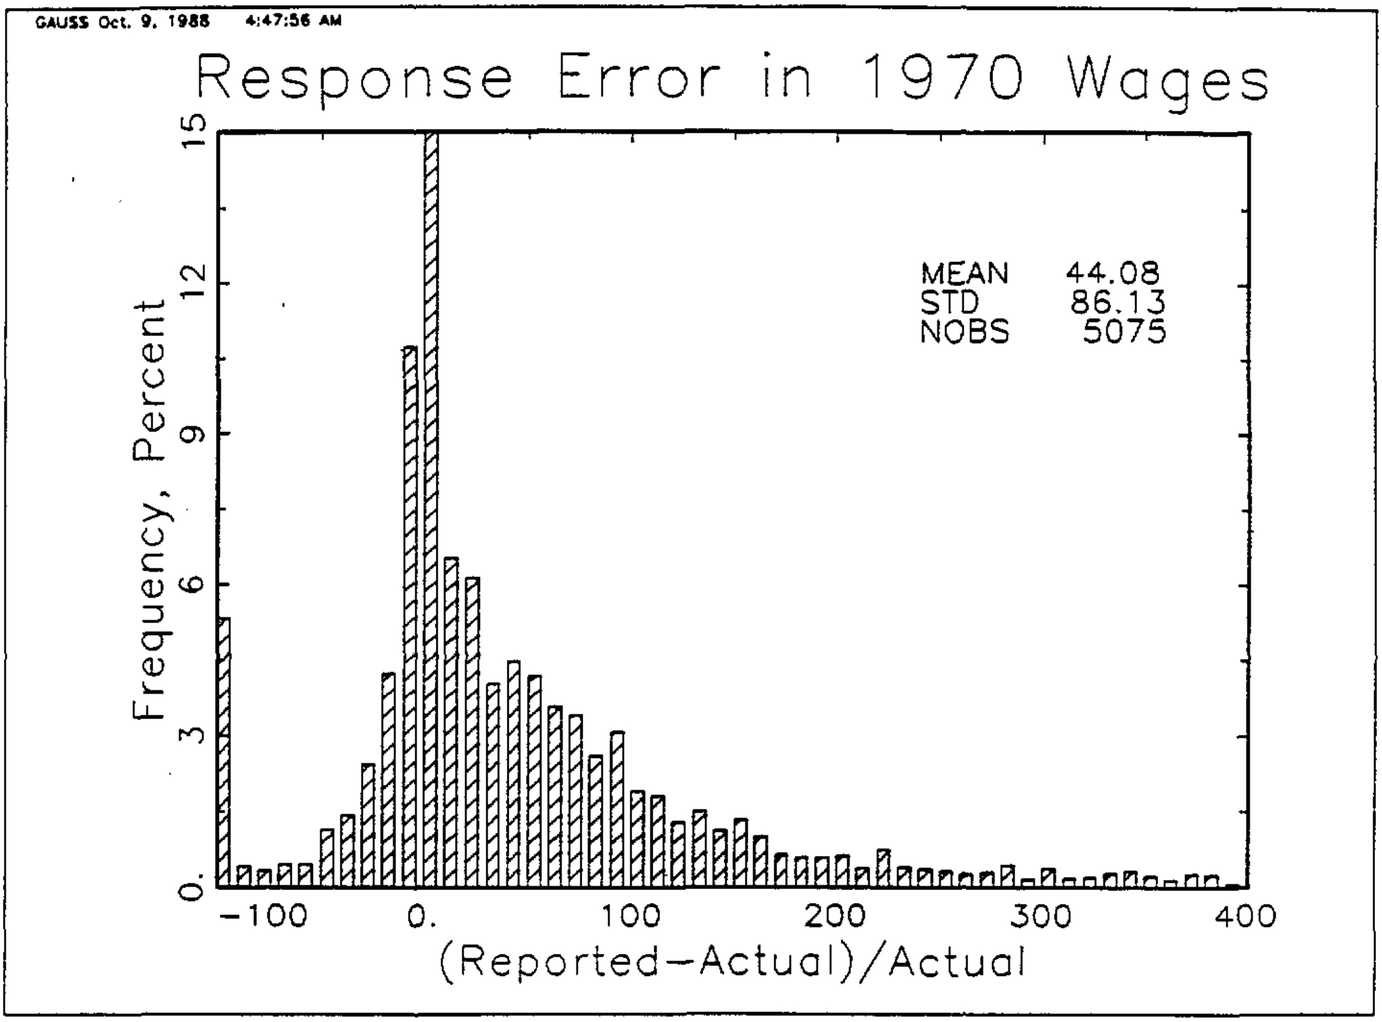
GAUSS Oct. 9, 1988 4:47:56 AM
Response Error in 1970 Wages
Frequency, Percent
0.
3
6
9
12
15
MEAN 44.08
STD 86.13
NOBS 5075
−100
0.
100
200
300
400
(Reported−Actual)/Actual

Figure 8.3

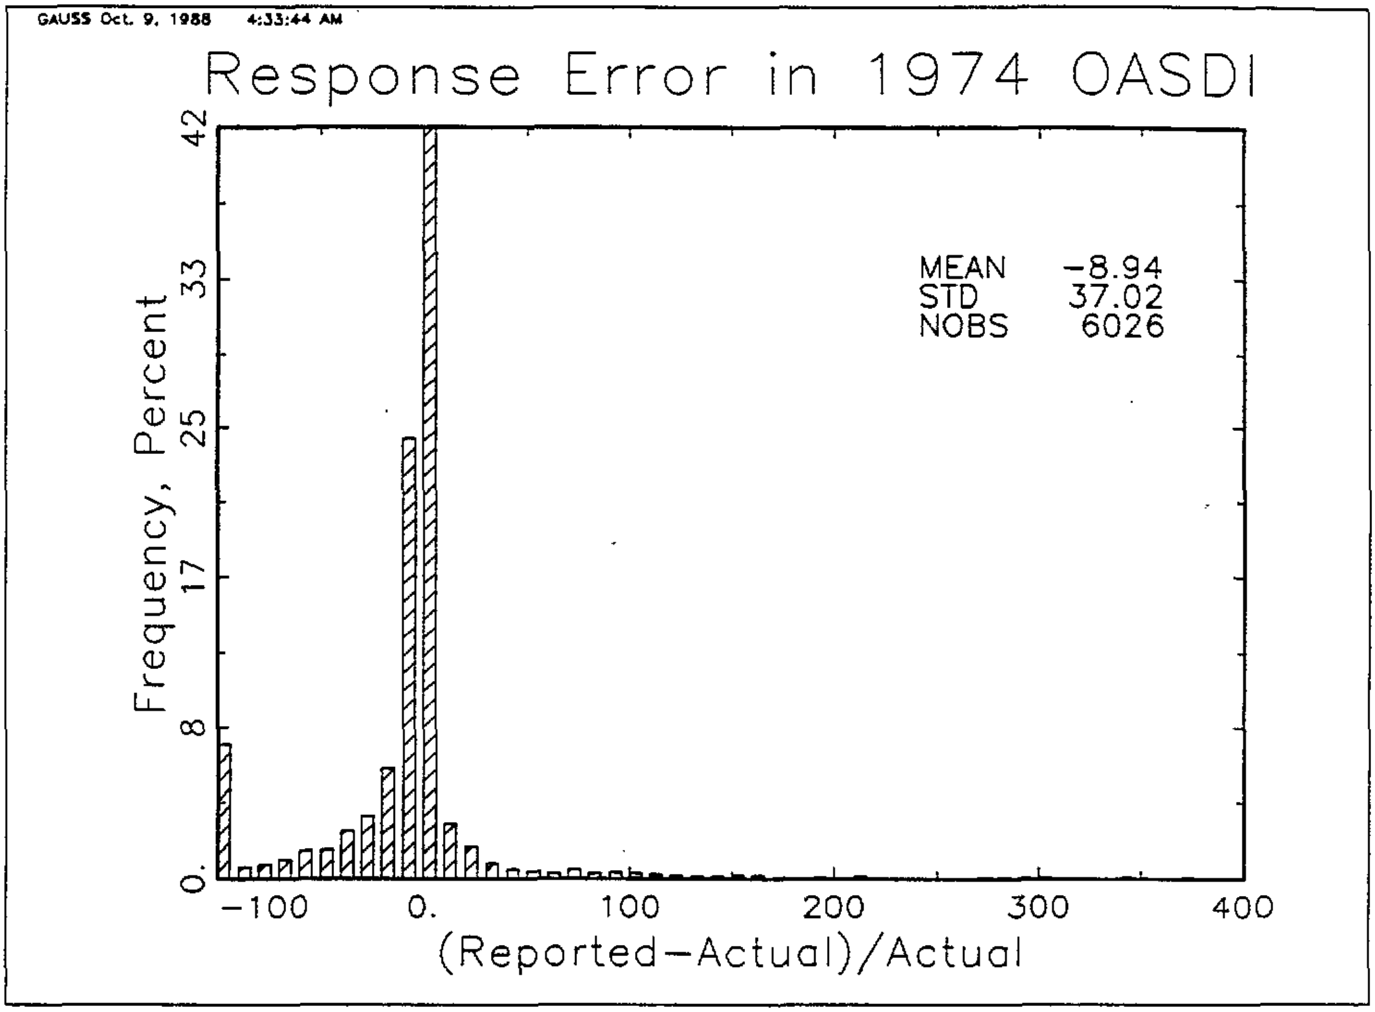
GAUSS Oct. 9, 1988 4:33:44 AM
Response Error in 1974 OASDI
Frequency, Percent
0.
8
17
25
33
42
MEAN -8.94
STD 37.02
NOBS 6026
-100
0.
100
200
300
400
(Reported-Actual)/Actual

Figure 8.4

I used the Hurd wealth data (see footnote 2) to compute respondents' net worth. Net worth consists of financial and real assets less total indebtedness, but excludes pensions, life insurance and annuities, the latter two which are fairly uncommon in the RHS anyway.[18] Wealth data are extremely hard to cross-check because major components of wealth, such as the market value of the respondent's house, are often subjective guesses. Figure 8.5 plots the distribution of wealth for the 6 survey years. Notice that there is a significant fraction of respondents, about 10%, who report that they have essentially no tangible wealth. Mean wealth levels are about $28,000 1968 dollars, equal to approximately four years of income. These distributions provide little evidence that respondents consume their wealth as they age. Figure 8.6 plots the distribution of housing value to net worth in 1969 and 1979. It shows that a large fraction of workers' wealth is tied up in housing: homeowners have an average 56% of their wealth tied up in housing in 1969 increasing to 65% in 1979. The failure of wealth to decrease over time may be partly due to the appreciation of housing in the inflationary 70's.

[18]While pensions are much more common than annuities in the RHS sample, exclusion of pension wealth is not a problem since the sample boolean already excludes workers with pensions.

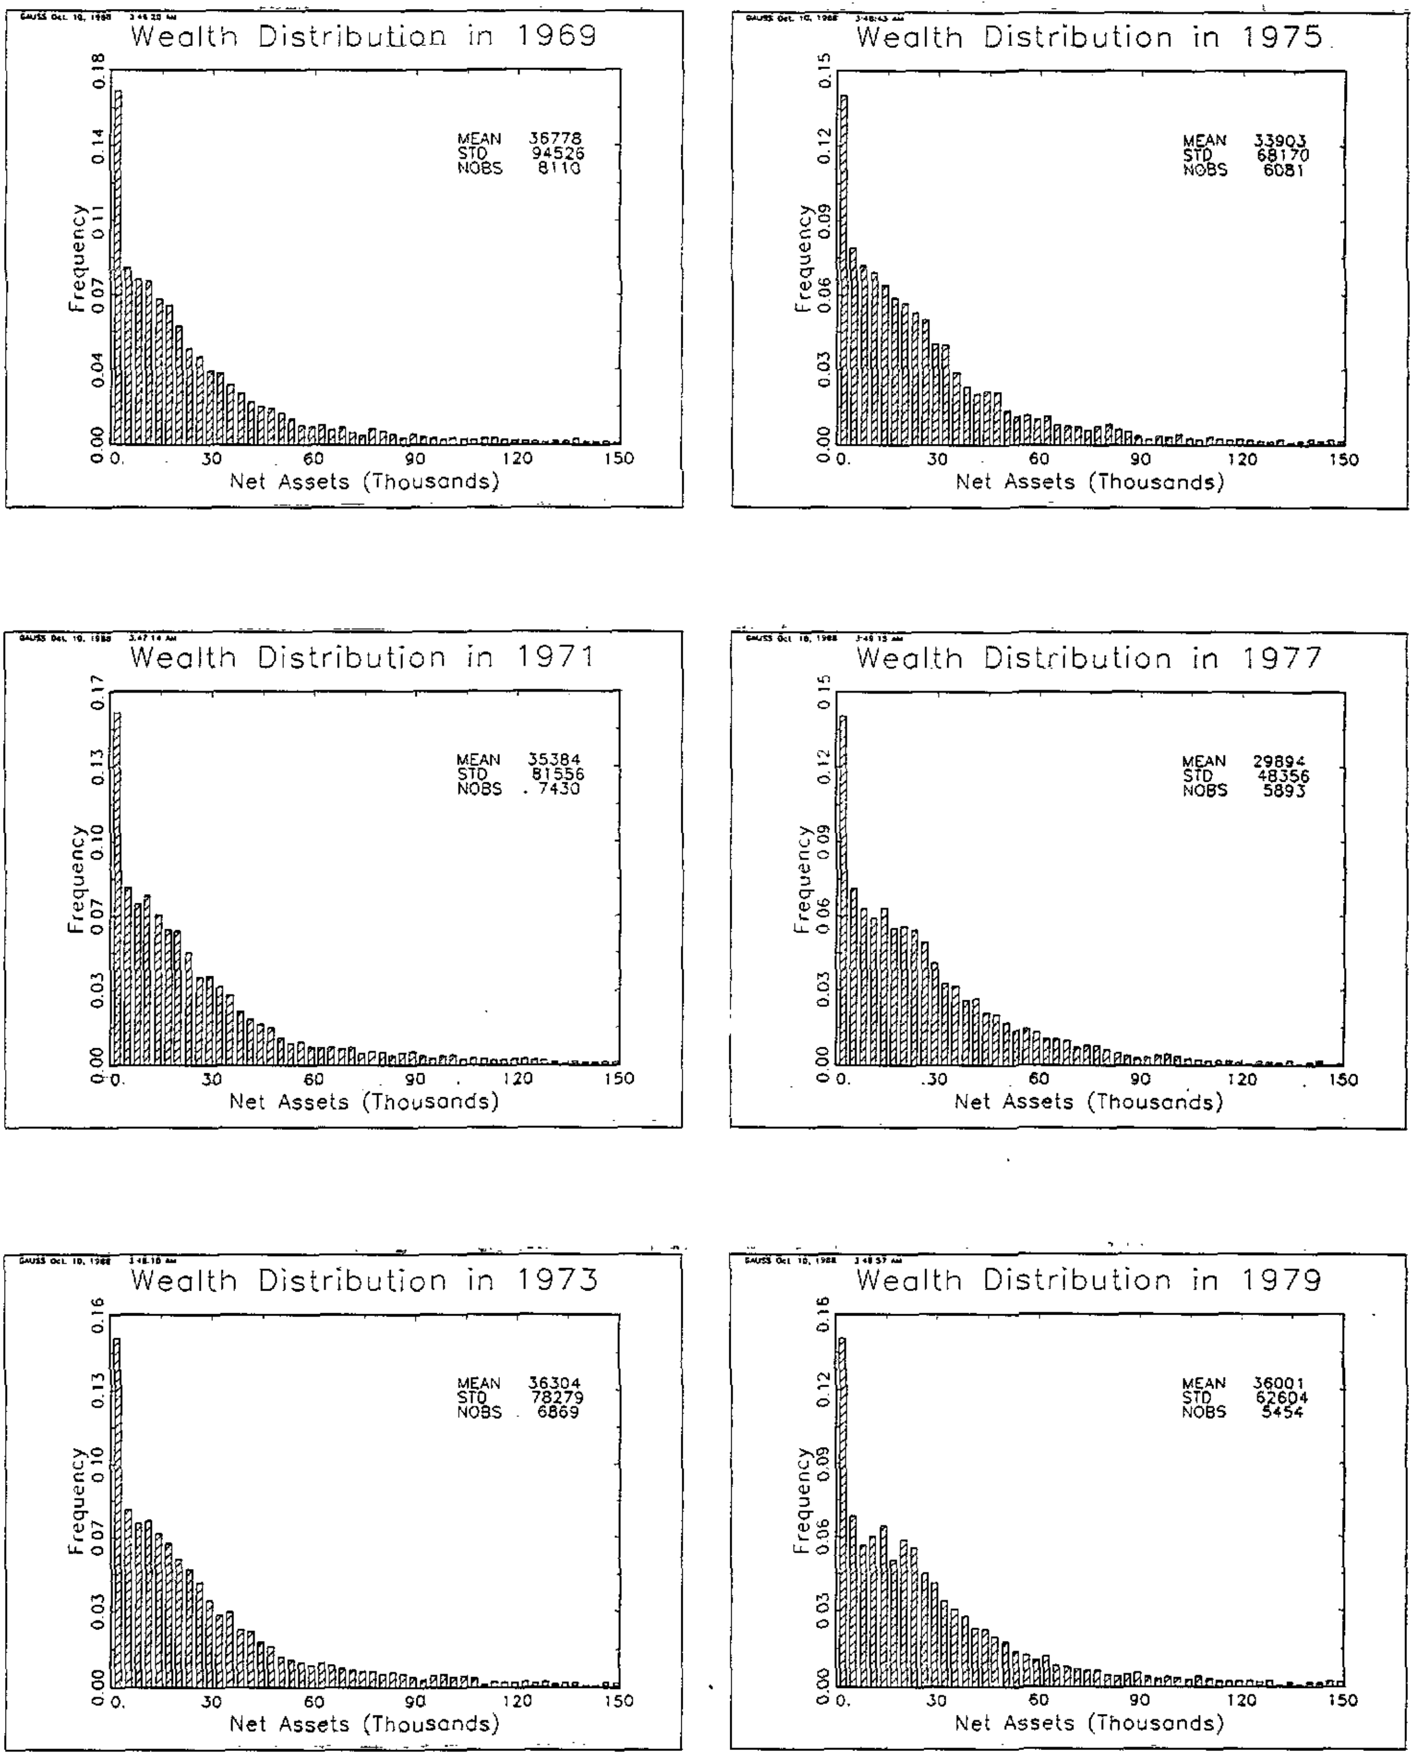
Wealth Distribution in 1969
MEAN 36778
STD 94526
NOBS 8110
Frequency
Net Assets (Thousands)
Wealth Distribution in 1975
MEAN 33903
STD 68170
NOBS 6081
Frequency
Net Assets (Thousands)
Wealth Distribution in 1971
MEAN 35384
STD 81556
NOBS 7430
Frequency
Net Assets (Thousands)
Wealth Distribution in 1977
MEAN 29894
STD 48356
NOBS 5893
Frequency
Net Assets (Thousands)
Wealth Distribution in 1973
MEAN 36304
STD 78279
NOBS 6869
Frequency
Net Assets (Thousands)
Wealth Distribution in 1979
MEAN 36001
STD 62604
NOBS 5454
Frequency
Net Assets (Thousands)

Figure 8.5

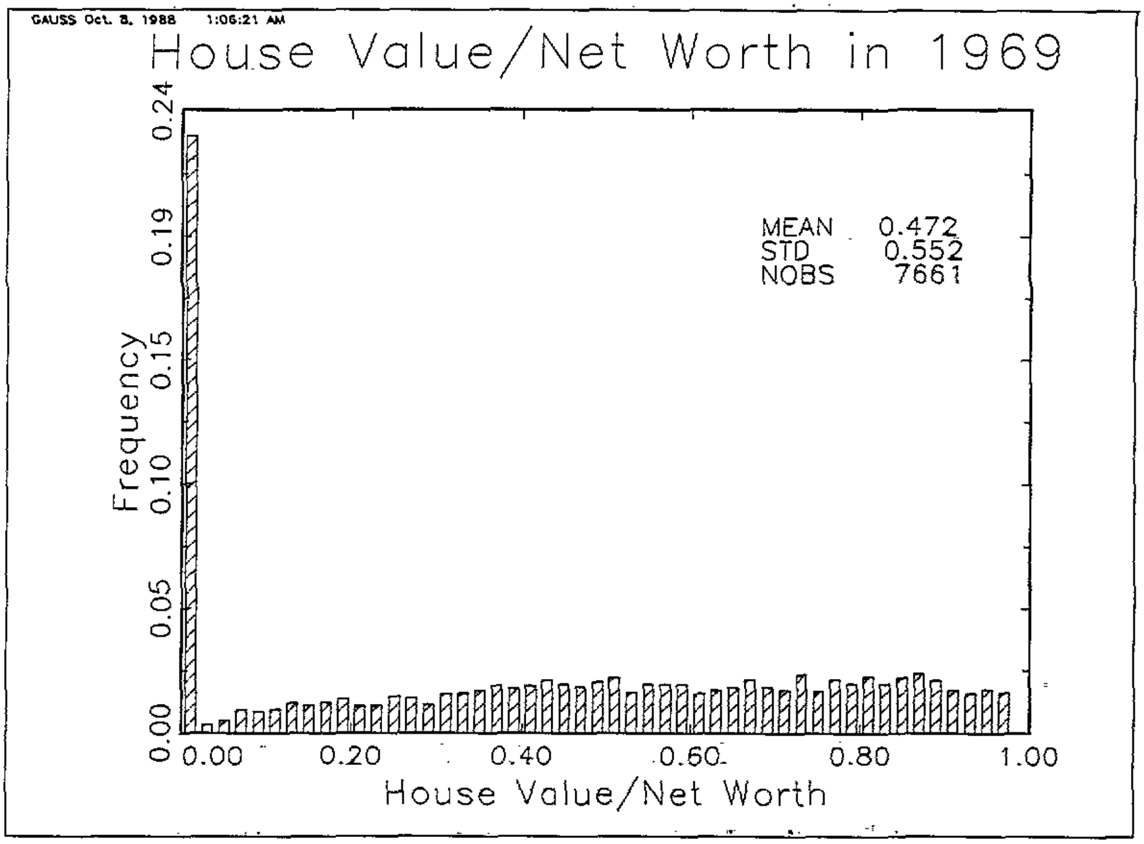
GAUSS Oct. 8, 1988 1:06:21 AM
House Value/Net Worth in 1969
MEAN 0.472
STD 0.552
NOBS 7661
Frequency
0.00
0.05
0.10
0.15
0.19
0.24
0.00
0.20
0.40
0.60
0.80
1.00
House Value/Net Worth

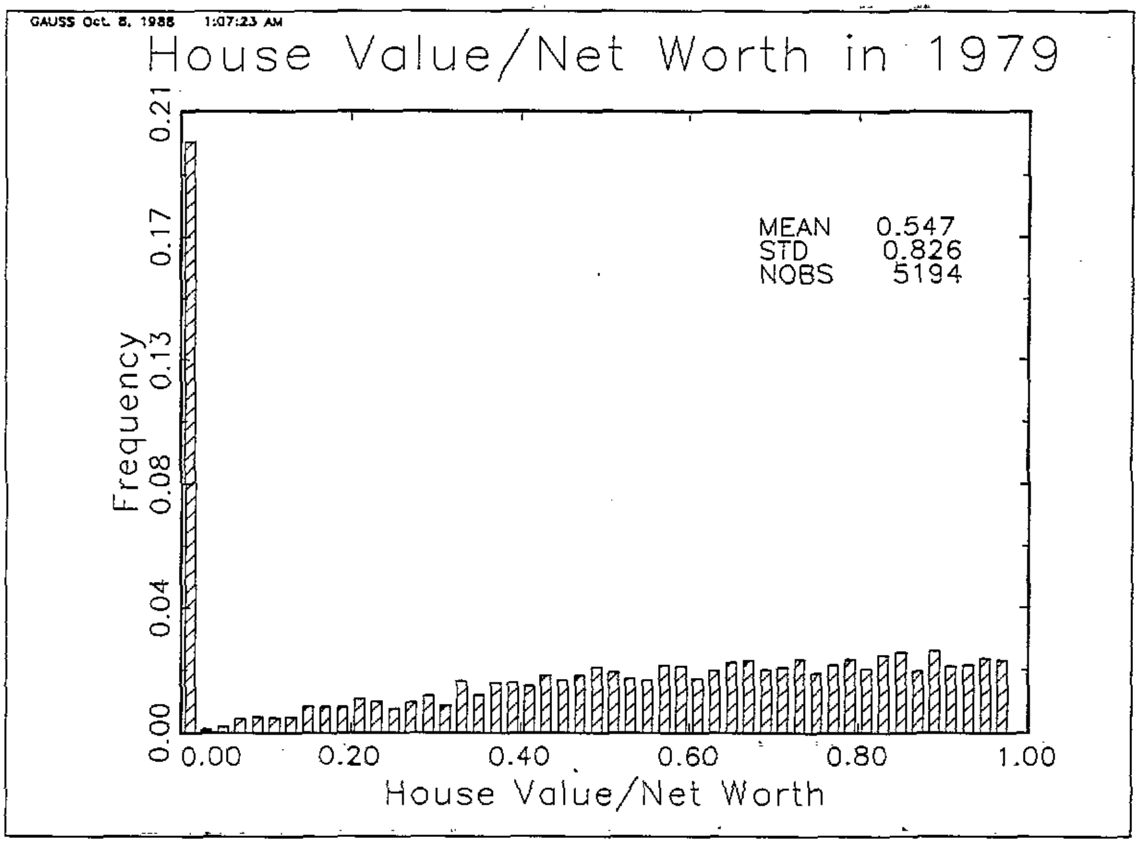
GAUSS Oct. 8, 1988 1:07:23 AM
House Value/Net Worth in 1979
MEAN 0.547
STD 0.826
NOBS 5194
Frequency
0.00
0.04
0.08
0.13
0.17
0.21
0.00
0.20
0.40
0.60
0.80
1.00
House Value/Net Worth

Figure 8.6

Using Hurd's wealth data and my imputed income series, I constructed an imputed biennial consumption series using the budget identity: $c_t=w_t-w_{t+1}+y_t$. The resulting consumption distributions are plotted in figure 8.7. Overall, the distribution of consumption looks very similar to the distribution of income plotted in figure 8.8; both income and consumption show a noticeable tendency to shift leftward over time. This fact is not an accident, since figure 8.9 shows that the distribution of wealth changes is centered about 0, suggesting that to a first approximation, $c_t=y_t$. Indeed the mean wealth change (averaged over all periods and workers) is $-658, with a standard deviation of $47,015. Given that average wealth is $28,000, it's difficult not to conclude

that most of the variation is due to measurement error. The large standard deviation suggests that it would be difficult to reject the hypothesis that $c_t=y_t$. However a simple hypothesis test of H_0: $c_t=y_t$ vs. H_A: $c_t\neq y_t$ yields a Chi-squared statistic of 6.2 with a marginal significance level of 1%: a rejection that is perhaps not surprising given that I have 31,348 observations on wealth changes.[19]

[19]The hypothesis actually tested was H_0: $w_t=w_{t+1}$ vs. H_A: $w_t\neq w_{t+1}$. It is easy to see that the budget identity implies that this is equivalent to the hypothesis test listed above. I assume that apropriate regularity conditions hold in order to justify the asymptotic Chi-square distribution for the test statistic, e.g. weak mixing conditions degree of dependence in the observations.

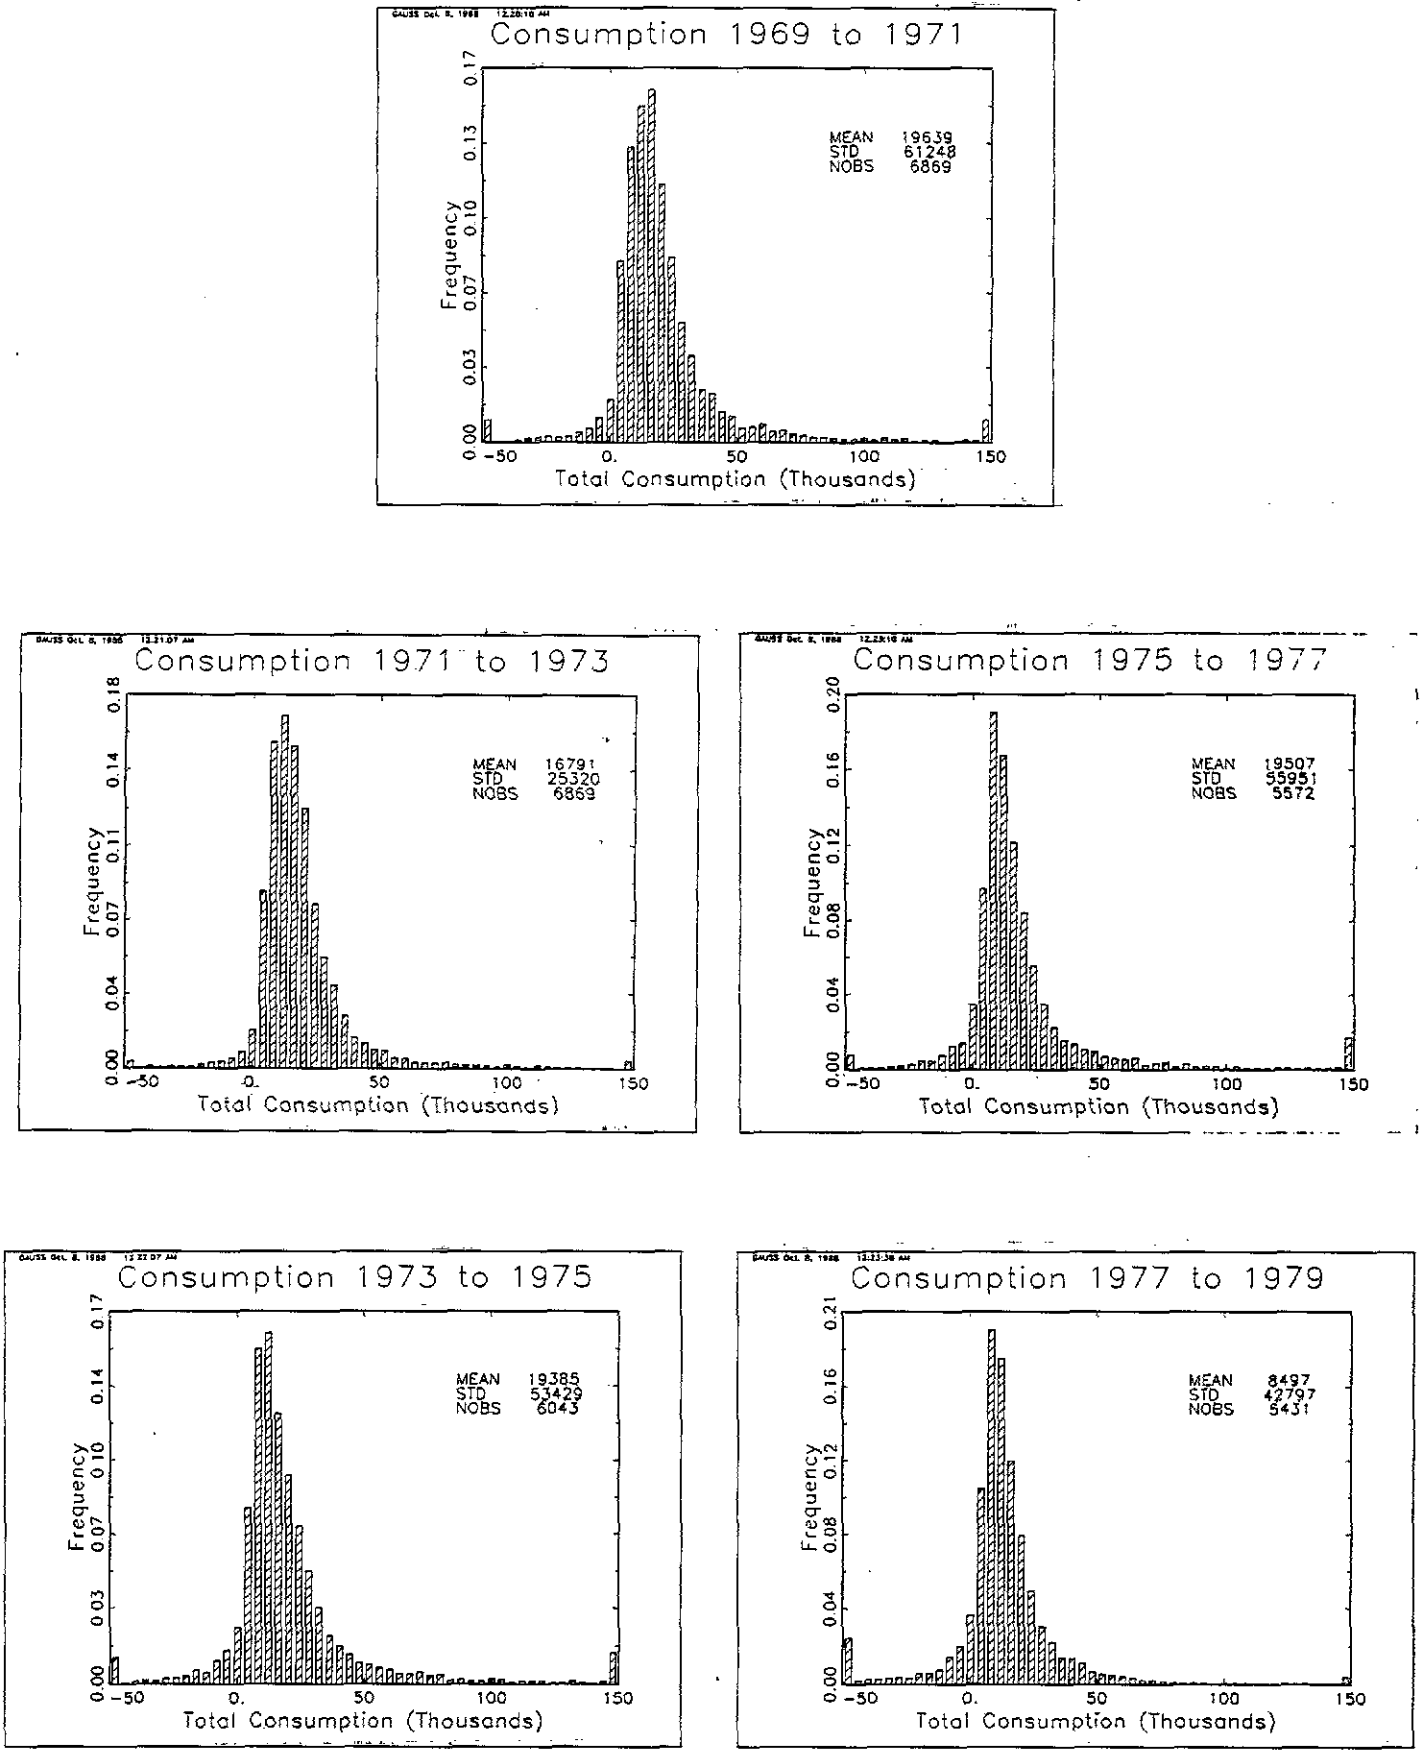
Consumption 1969 to 1971
MEAN 19639
STD 61248
NOBS 6869
Frequency
0.00 0.03 0.07 0.10 0.13 0.17
-50 0. 50 100 150
Total Consumption (Thousands)
Consumption 1971 to 1973
MEAN 16791
STD 25320
NOBS 6869
Frequency
0.00 0.04 0.07 0.11 0.14 0.18
-50 0. 50 100 150
Total Consumption (Thousands)
Consumption 1975 to 1977
MEAN 19507
STD 55951
NOBS 5572
Frequency
0.00 0.04 0.08 0.12 0.16 0.20
-50 0. 50 100 150
Total Consumption (Thousands)
Consumption 1973 to 1975
MEAN 19385
STD 53429
NOBS 6043
Frequency
0.00 0.03 0.07 0.10 0.14 0.17
-50 0. 50 100 150
Total Consumption (Thousands)
Consumption 1977 to 1979
MEAN 8497
STD 42797
NOBS 5431
Frequency
0.00 0.04 0.08 0.12 0.16 0.21
-50 0. 50 100 150
Total Consumption (Thousands)

Figure 8.7

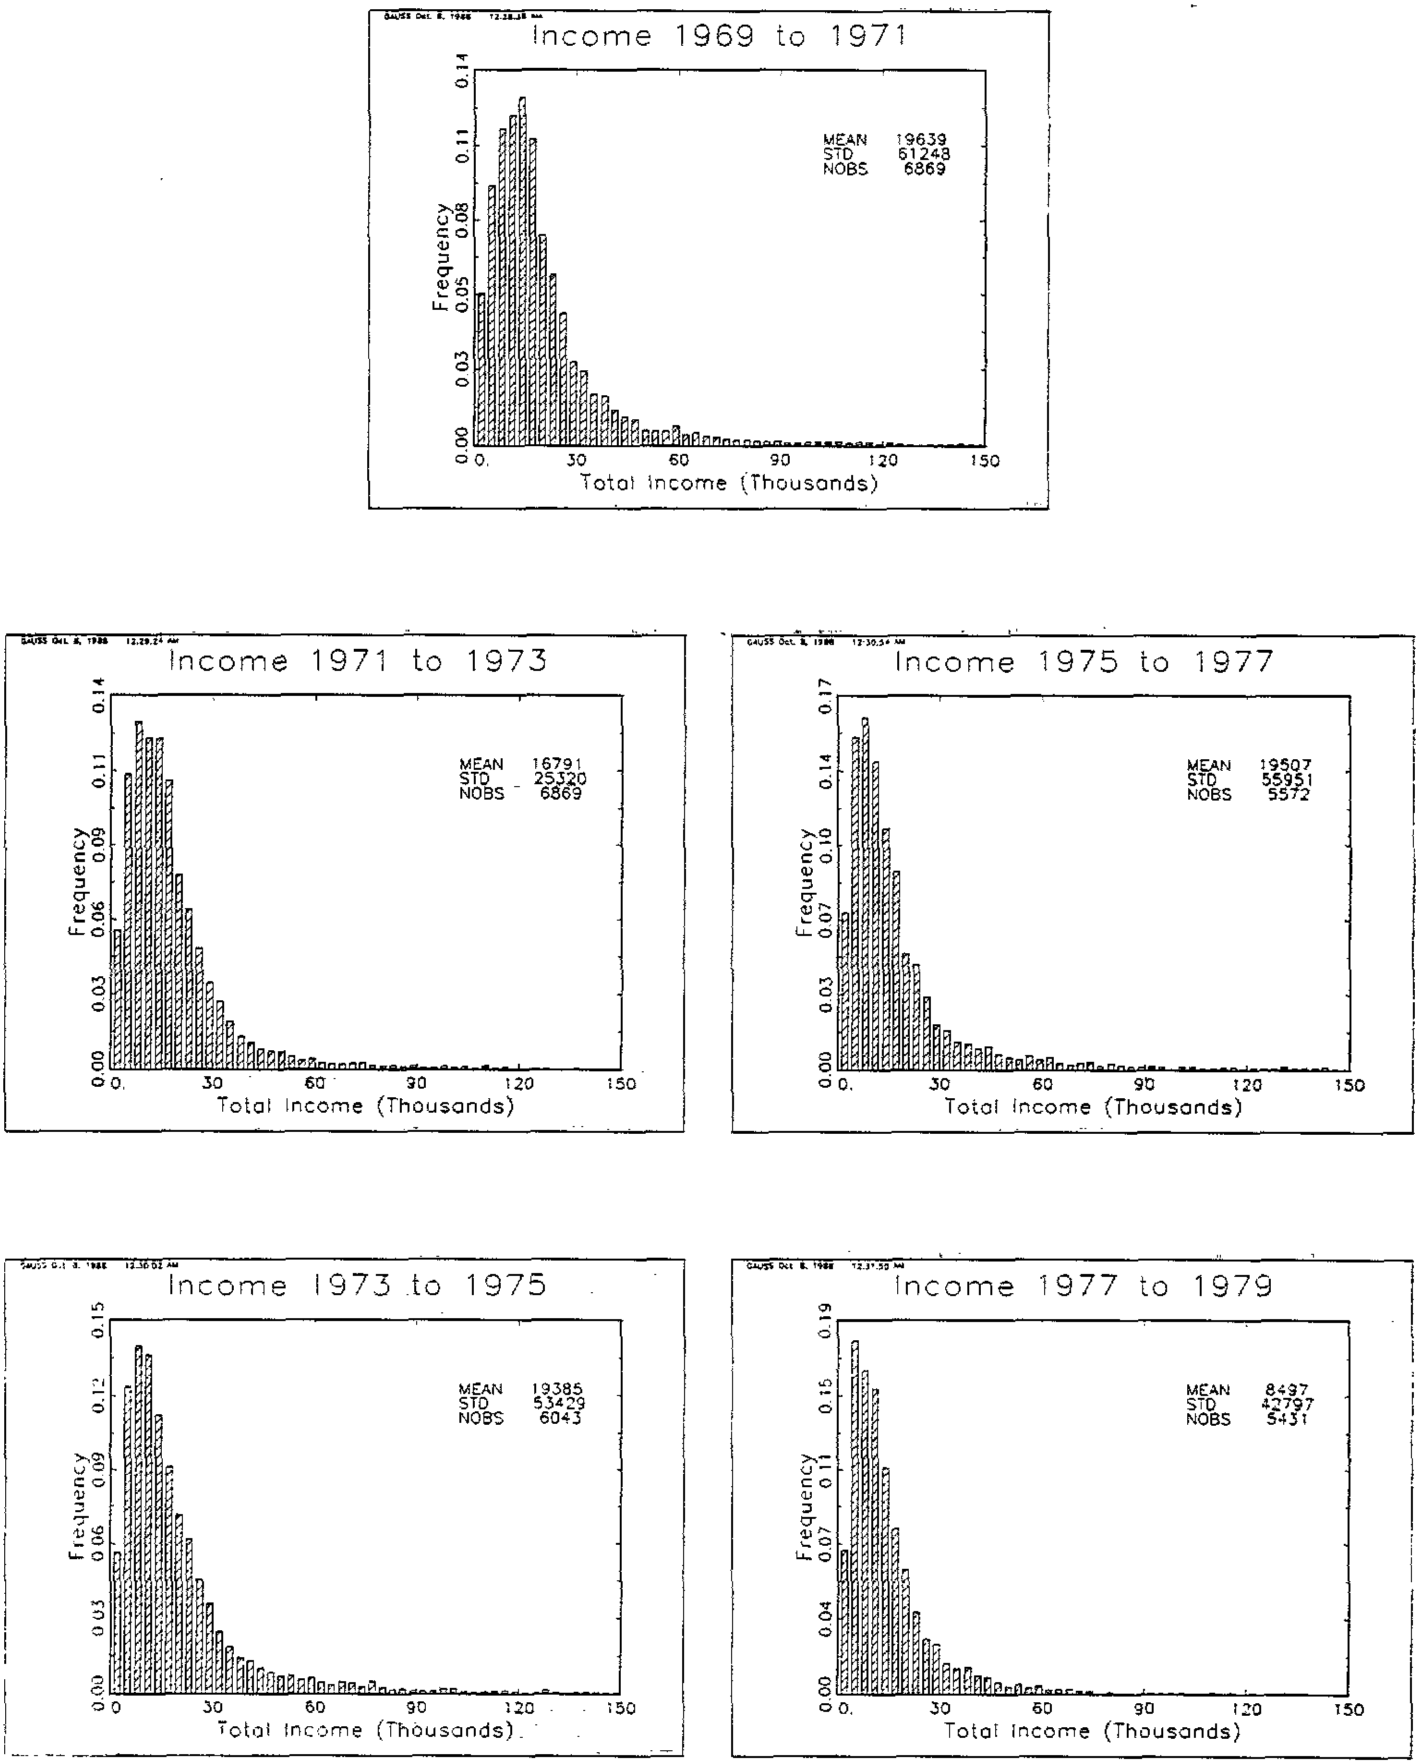

Income 1969 to 1971
MEAN 19639
STD 61248
NOBS 6869
Frequency
Total Income (Thousands)
Income 1971 to 1973
MEAN 16791
STD 25320
NOBS 6869
Frequency
Total Income (Thousands)
Income 1975 to 1977
MEAN 19507
STD 55951
NOBS 5572
Frequency
Total Income (Thousands)
Income 1973 to 1975
MEAN 19385
STD 53429
NOBS 6043
Frequency
Total Income (Thousands)
Income 1977 to 1979
MEAN 8497
STD 42797
NOBS 5431
Frequency
Total Income (Thousands)

Figure 8.8

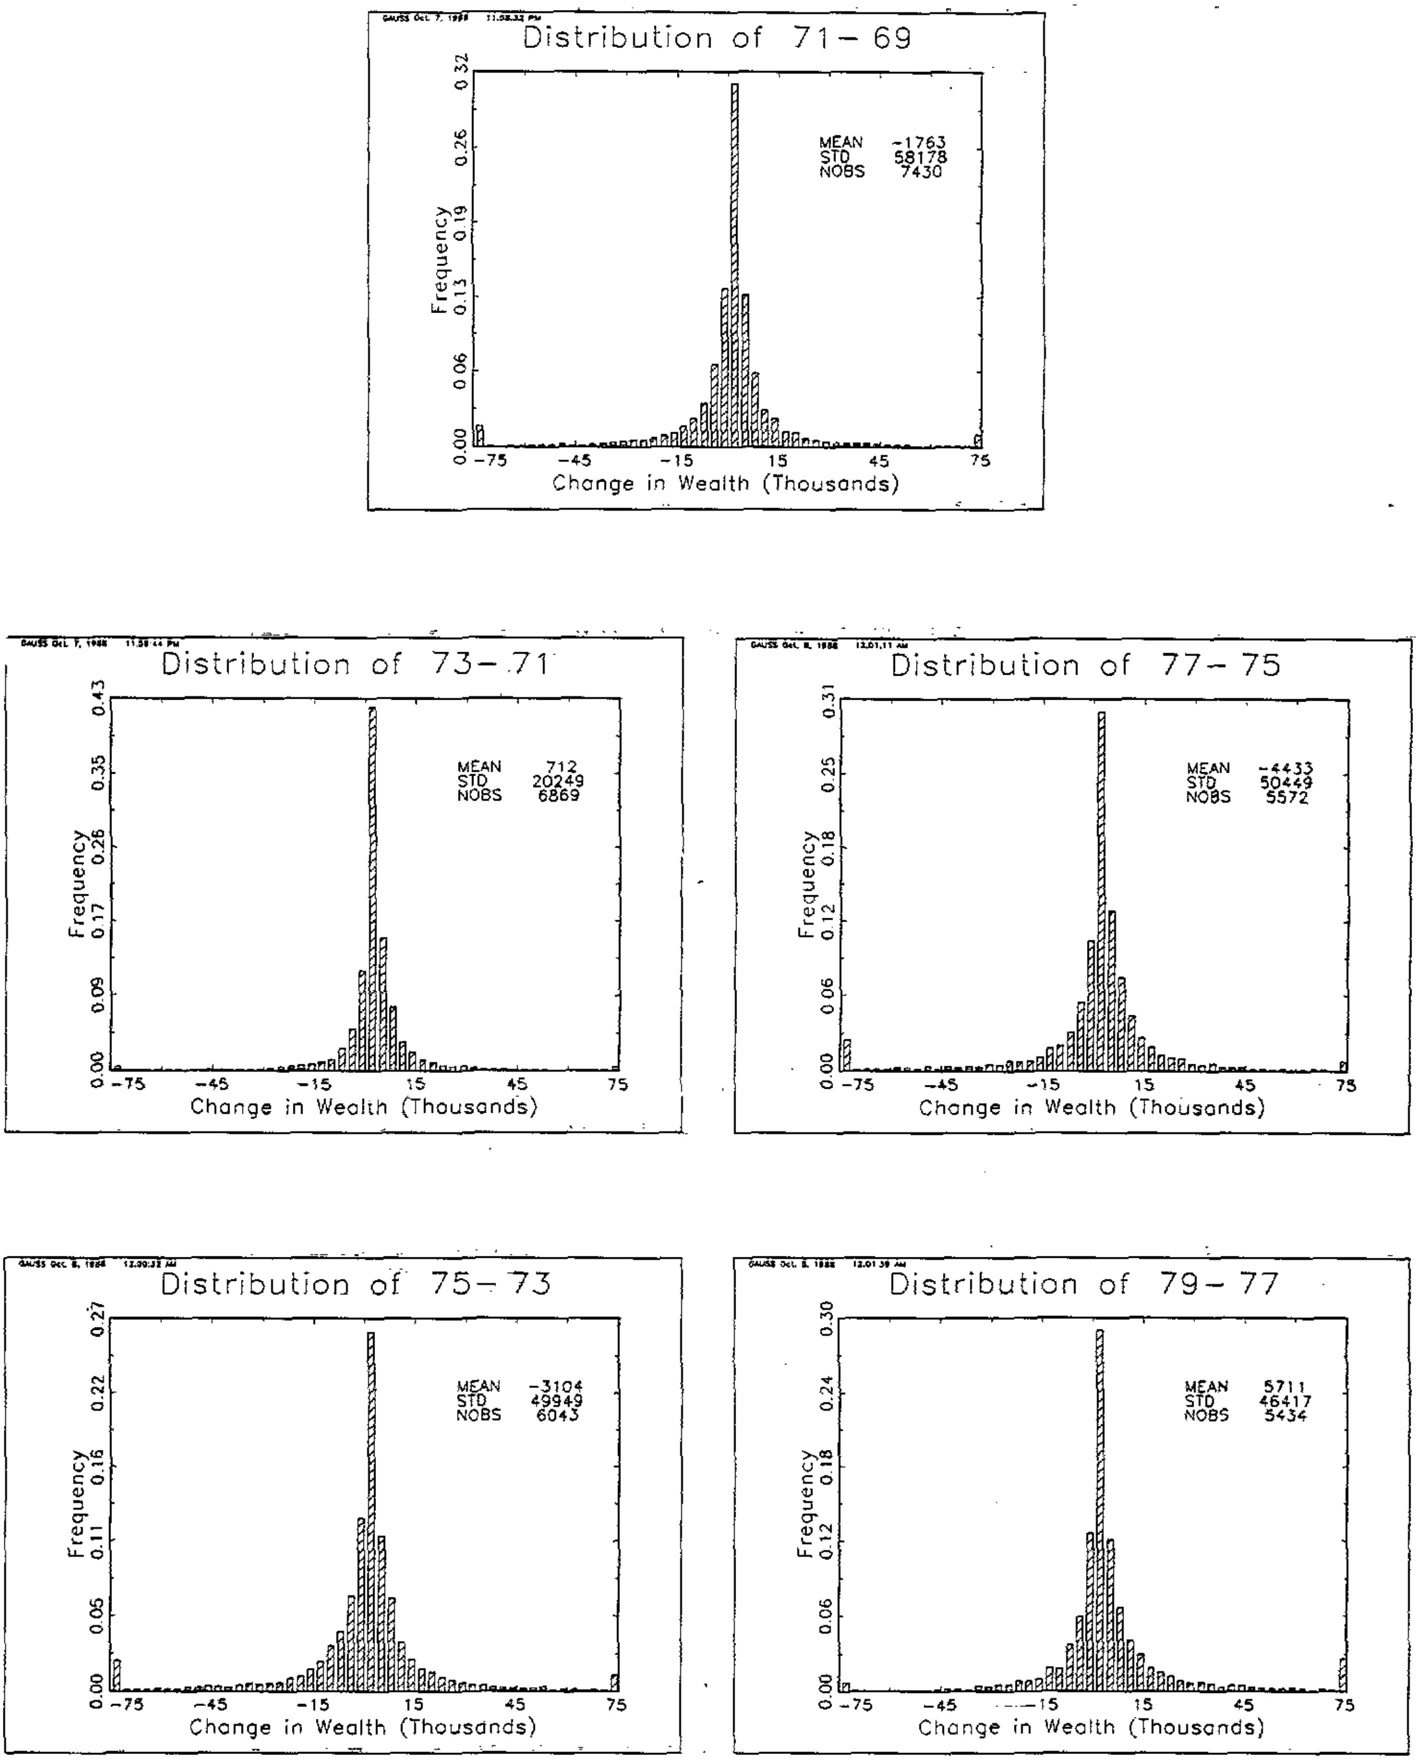

Distribution of 71− 69
MEAN -1763
STD 58178
NOBS 7430
Frequency
0.00 0.06 0.13 0.19 0.26 0.32
−75 −45 −15 15 45 75
Change in Wealth (Thousands)
Distribution of 73− 71
MEAN 712
STD 20249
NOBS 6869
Frequency
0.00 0.09 0.17 0.26 0.35 0.43
−75 −45 −15 15 45 75
Change in Wealth (Thousands)
Distribution of 77− 75
MEAN -4433
STD 50449
NOBS 5572
Frequency
0.00 0.06 0.12 0.18 0.25 0.31
−75 −45 −15 15 45 75
Change in Wealth (Thousands)
Distribution of 75− 73
MEAN -3104
STD 49949
NOBS 6043
Frequency
0.00 0.05 0.11 0.16 0.22 0.27
−75 −45 −15 15 45 75
Change in Wealth (Thousands)
Distribution of 79− 77
MEAN 5711
STD 46417
NOBS 5434
Frequency
0.00 0.06 0.12 0.18 0.24 0.30
−75 −45 −15 15 45 75
Change in Wealth (Thousands)

Figure 8.9

Whether the large variance in wealth reflects explainable differences in behavior or simple measurement error is an open question, but my initial investigations suggest the dominance of the latter. Like the employment data, aggregate consumption appears fairly smooth, slowly declining over time in apparent accord with the standard life-cycle hypothesis. However at the individual level, measured consumption is anything but smooth, making violent, unpredictable swings over time. Overall, a total of 1984 respondents have negative measured consumption in at least one of the 5 biannual survey periods. These large swings in consumption fly in the face of intuition and personal observation of the consumption behavior of the elderly, suggesting that most of the swings are due to measurement errors in wealth.

One possible reason for negative consumption is the failure to account for capital gains. Given the subjectivity of respondents' assessment of housing values and the fact that a majority of workers continue to live in the same house rather than "size down", it seemed reasonable to attribute all changes in net housing wealth to capital gains (provided the respondent did not move). Adding these housing capital gains (or losses) did reduce the number of negative consumption cases somewhat, to 1522, but overall the distribution of consumption including capital gains looked very similar to the distribution of consumption without capital gains.

The notion that response errors in wealth are driving the violent swings in consumption is confirmed by examining individual data records. Having access to a complete data record over the survey period often provides enough contextual information to enable one to intuitively identify reporting and recording errors that are responsible for negative consumption values. Table 8.1 presents relevant data for a "typical" respondent (ID number 6886) with negative measured consumption. This man, call him Bob, is coded as having the occupation of craftsman in the construction industry; most likely Bob is a carpenter. Bob responded in all six of the survey waves, and provided very complete answers; all the variable flags (with the exception of consumption) indicated very high confidence levels in his responses. Bob is married, living with his spouse and was working full-time up until 1975 when he turned 65, quit his job, and started collecting Social Security (Bob had no pensions). By all accounts, Bob is just the kind of guy I want in my sample: a typical blue collar worker who seems to provide complete, reliable answers, who has slightly above average income and most of his wealth in housing. However we can see from table 8.1 that while Bob's income declined slightly from $12,000 to $11,000 over the decade and his measured consumption was about equal to his income stream in 4 out the the 5 two-year periods, for some reason his consumption over the period 1976-77 is recorded as $-35,630. Analysis of his balance sheet reveals that between the 1975 and 1977 interviews, his house value increased from $14,000 to $100,000, increasing his overall net worth from $10,794 to $57,376.[20] This sudden increase in wealth is responsible for the recorded negative consumption in the 1976-77 biennium. A possible explanation for the increase is coding error: Bob may have reported his house value to be $10,000 in 1977 but it was mistakenly recorded as $100,000. However this explanation becomes less plausible when we realize that his house value is recorded at $150,000 in 1979: it seems very unlikely that we would get the same kind of coding error in in the same variable in two consecutive years.

If we look further into the data, we find that Bob moved between 1975 and 1977. This suggests several possibilities: 1) Bob and his wife may have moved into the house of his wealthy son and had mistakenly reported the value of his son's house as his own, or 2) Bob may have previously grossly underestimated the value of his old house, and used the capital gains on the sale of the old house to finance the purchase of his new house, or 3) Bob may have won a

[20] Note that no capital gains are imputed since Bob moved during the period, making it impossible to determine how much of the value of the new house came from capital gains on the sale of the old house.

lottery which provided an unrecorded capital gain which he used to purchase his retirement dream home, or 4) Bob may have been sitting on a nest egg which he refused to report in previous interviews and has since used it to buy his retirement home, or 5) being a carpenter, Bob may have built his own retirement home and blinded by the pride of creation, has grossly overestimated its value. Given the wealth of possible explanations, it's not easy to know what to do. One can simply exclude cases with negative measured consumption, but that still leaves the problem of hundreds of cases with implausibly large or small measured consumption, or cases where consumption changes very erratically from year to year.

Table 8.1: Selected Financial Data for "Bob"
RHS ID: 6886

		1969	1971	1973	1975	1977	1979
Personal Data							
	a_t	59	61	63	65	67	69
	ms_t	1	1	1	1	1	1
	h_t	1	2	1	1	2	2
Employment Data							
	IE	1	1	1	1	3	3
	SR	1	1	1	2	3	3
	SE	1	1	1	2	3	3
Financial Data ($1968)							
	w_t	10698	12523	13950	10794	57376	71555
	y_t	M	12033	11951	10199	10952	11429
	c_t[21]	M	10196	10524	13354	-35630	-2750
	c_t[22]	M	11154	12047	12272	-35630	11151
	capital gains	M	958	1523	-1082	0	13901

[21] This measure of c_t does not include imputed capital gains.

[22] This measure includes imputed capital gains as described on page 34.

Balance Sheet (nominal)

house value	8000	10000	13000	14000	100000	150000
mortgage	0	0	0	0	0	0
other house debt	0	0	0	0	0	0
farm value	0	0	0	0	0	0
farm mortgage	0	0	0	0	0	0
business value	0	0	0	0	0	0
business debt	0	0	0	0	0	0
real-estate value	0	0	0	0	0	0
real-estate debt	0	0	0	0	0	0
auto value	2490	2490	2495	2500	3237	4980
auto debt	0	0	0	0	0	0
savings bonds	2000	2000	2408	0	0	0
stocks	0	0	0	0	0	0
credit card debt	0	200	236	0	0	500
checking account	190	900	900	900	900	1080
savings account	0	0	0	0	0	0
face value life ins	1000	1000	913	2000	2000	2000
face value annuities	0	0	0	0	0	0
medical debts	834	0	0	0	0	0
store debts	100	0	0	0	0	0
bank debts	0	0	0	0	0	0
personal debts	0	0	0	0	0	0

In conclusion, while it is fairly easy to examine and identify reporting problems by examining observations on a case-by-case basis, it is unrealistic to think that I could screen out a sufficiently high fraction of "bad" cases to end up with a subsample for which consumption is measured accurately. Not only is case-by-case examination of 8,131 individuals impossibly time consuming, the resulting data set would be susceptible to the criticism that the sample had been "hand-picked" to support an a priori theory. If an error-identification strategy is to be successful, one should be able to write out a series of objective classification rules, say in the form of a computer program, that would allow other researchers to replicate the subsample. I have not been successful in constructing a computer program with a sufficient "intelligence" to examine the wealth data on a case by case basis, recognize the existence of a data problem, and take appropriate corrective action. As I discussed above, it is not sufficient to simply screen out cases with negative consumption because the remaining cases still suffer from reporting problems that produce unrealistically large swings in consumption. Because of these problems, I have opted against using consumption data in my first attempts at estimating the DP model. Until I see convincing evidence that changes in wealth are not dominated by measurement error, or until I am successful in constructing an "artificial intelligence" routine that discriminates accurate survey responses from innacurate responses, I will adopt the null hypothesis that $c_t=y_t$ and focus on "explaining" the joint dynamics of $x_t=(y_t,e_t,a_t,ms_t,h_t)$ and $d_t=(s_t,ss_t)$, excluding w_t and c_t from the model.

9. Estimating the Stochastic Process of Income

All that remains is to specify and estimate the final component of workers' beliefs, the transition density for income π_y. The lognormal shapes of the income distributions plotted in section 8 suggest that the transition density π_y should have a lognormal distribution with parameters $(\mu\ \sigma)$ which are linear-in-parameters functions of the state and control variables listed in the decomposition (3.2). As is well-known, if a random variable $\tilde{y}$ has a lognormal distribution, then its mean and variance are given by

$$E[\tilde{y}] = \exp\{\mu+\sigma^2/2\}$$

(9.1)

$$\mathrm{var}[\tilde{y}] = \exp\{2\mu+2\sigma^2\} - \exp\{2\mu+\sigma^2\}$$

It's extremely important to allow both μ and σ to depend on the state variables, since if σ is fixed, then (9.1) and the autoregressive properties of the income process will imply that the variance of y_{t+1} is an exponentially increasing function of current income y_t. Thus, by failing to specify σ properly, one is making an implicit assumption about the form of heteroscedasticity that may grossly misrepresent workers' actual beliefs. Once we have decided on the appropriate functional forms for μ and σ, the lognormal model is fairly easy to estimate: one obtains initial estimates of (μ,σ) by a log-linear regression, and uses these as starting values for computing the final parameter estimates by maximum likelihood.[23] There is a minor problem concerning the fact that the DP model requires y_t and its transition density π_y to be discretized. My approach was to discretize y_t as an independent variable entering (μ,σ), but to do the estimation treating the dependent variable y_{t+1} as a continuous variable. After estimating the relevant parameters, it's easy generate a discrete transition probability matrix $\hat{\pi}_y$: simply compute the area under the lognormal density corresponding to each of the discrete income cells for y_{t+1}.

The hard part is to specify how the parameters (μ,σ) depend on the underlying state and control variables. The specification is crucial here, because not only must π_y embody workers' expectations about how future income depends on current their employment, health and marital status, it must also embody the relevant rules and actuarial structure of the Social Security OASDI system, including the regressive nature of the payout schedule, the extra payments to spouse, the penalty for early retirement, and the "earnings test" for workers under 70. As Idiscussed in my earlier paper, by estimating π_y using income data over the decade of the 70's (during which Social Security benefits increased more than 50% in real terms) I have implicitly assumed that workers have "semi-rational" expectations: i.e. they correctly anticipated the increase in benefits over the 70's, but did not expect any benefit changes

[23] Although the likelihood function is concave, using the regression starting values (as opposed to zero starting values) substantially reduces the number of iterations needed to converge.

thereafter.[24]

My initial attempts to estimate π_y yielded disappointing results. Although the coefficient estimates for the marital status, employment status and search variables had reasonable signs and magnitudes, the variables representing the structure of OASDI benefits either had small, insignificant coefficients or else had the wrong sign. The estimated model looked as if workers were unaware of key features of the OASDI benefit plan, and the few provisions they did know about seemed to be regarded as taxes instead of benefits. Apparently, the Social Security benefit structure was "drowned out" by sample selection bias. A simple explanation of the problem goes as follows. High income workers typically continue working beyond retirement age, whereas low income workers retire as soon as they can. A regression model attempts to fit this data by generating negative returns to retirement. My solution to the problem was to augment the data set with "artificial" data on the incomes retired workers would have received in the absence of OASDI payments. Thus, corresponding to each data record for a retired worker receiving OASDI ($ss_t \in \{1,2\}$), I created a duplicate record deducting all OASDI benefits from the worker's income y_t and setting $ss_t=0$. This procedure, which nearly doubled the number of observations, produced dramatically improved results. In particular, nearly all the Social Security variables had significant coefficients with correct signs and magnitudes. In effect, the augmented data "drowned out" the sample selection bias, allowing me to capture the true underlying OASDI benefit structure. The existence of the SSMBR data set was absolutely crucial to the success of this procedure since as I have shown, the magnitude of response error in the self-reported values of certain Social

[24] In fact, the large benefit increases in the 70's put severe strain on the Social Security trust fund, necessitating substantial tax increases to fund the system. Fully rational workers might expect real benefit decreases in the future.

Security benefits such as SSDI is so large as to render them useless.

A final problem I encountered concerned the estimation of age-income effects. In my initial specifications I included the polynomial terms in the age variable a_t to capture the independent effects of aging on income. Just looking at the estimated coefficients, the estimated model seemed quite reasonable, with age terms all entering with highly significant coefficients. However when I plotted out the age-income profiles, the results were clearly far from reasonable. In models that included only a linear term in a_t the age-income profile sloped upward, whereas in models with quadratic and cubic age terms the age-income profile was hump-shaped: rising until age 70 and then sharply falling thereafter. The incomes predicted by the hump-shaped profiles were completely unreasonable: at the top of the hump a 70 year old worker who was currently earning $10,000 could expect to earn nearly double that amount two years later if he continued working. On the other hand, on the downward sloping part of the profile, say at age 80, the worker would only expect to make half as much even if he continued working. The reason behind these strange results is lack of data on earnings for very old men. As I have discussed before in section 5, the RHS has no data on workers older than 73. Thus, estimation of age-income profiles until age 98 requires pure extrapolation over a region where there are no observations to guide us. Including polynomial age terms in the regression produced unreasonable forecasts because the estimation procedure chose the coefficients to get a good fit in the region where there are a lot of observations, namely, for ages 58 to 68. Since there are no observations beyond age 73, the regression doesn't "care" what its predictions are in that range, and thus the crazy results. In order to avoid the extrapolation problems inherent in the use of polynomial terms, I tried specifications using age dummies, which entail the implicit extrapolation that age-income profiles are constant after age 73. In spite of my hopes, the age dummies also yielded somewhat disappointing results: the estimated age-income profile fluctuated up and down with no clear pattern. Since I have little a priori knowledge of the correct shape of the age-income profile, I decided simply not to include a_t in the estimation of π_y.

Table 9.1 presents the specification for π_y that I finally settled upon. The main implications of table 9.1 have already been discussed in conclusion 6 of section 3, and will not be repeated here. However to convince the reader that the estimated model really does endow workers with sensible income expectations, I present a graphical summary of the predictions of the model in figures 9.1 to 9.6.

Table 9.1: Estimates of Income Transition Probability

Dependent variable $\ln(y_{t+1})$

variable	parameter estimates	corrected std error	uncorrected t-statistic	corrected t-statistic
		σ parameters		
constant	-0.25126219	0.02855069	-21.34212254	-8.80056532
$\ln(y_t)$	-0.51962543	0.01628730	-86.07562382	-31.90371387
		μ parameters		
constant	-0.12229749	0.04289815	-3.28057767	-2.85088040
$\ln(y_t)$	0.94987850	0.00587781	189.14920514	161.60425480
$h_t=1,\ h_{t+1}=1$	0.02940221	0.00766624	3.95473986	3.83528475
$h_t=1,\ h_{t+1}=3$	-0.27613498	0.06620089	-7.13456992	-4.17116714
$h_t=2,\ h_{t+1}=2$	0.00755871	0.01090634	0.74628145	0.69305704
$h_t=2,\ h_{t+1}=3$	-0.24915939	0.07111086	-5.63796860	-3.50381631
$h_t=3,\ h_{t+1}=3$	0.05060739	0.01724244	3.34964834	2.93504796
$s_t=1,\ e_{t+1}=1$	0.22207861	0.01554447	13.04453198	14.28666204
$s_t=1,\ e_{t+1}=3$	-0.24745222	0.02112119	-12.23124090	-11.71582949
$s_t=2,\ e_{t+1}=1$	0.19803935	0.02666134	6.40437111	7.42796085
$s_t=2,\ e_{t+1}=2$	0.00128486	0.01516328	0.07500851	0.08473524
$s_t=2,\ e_{t+1}=3$	-0.31740913	0.02646962	-13.52888152	-11.99144848
$s_t=3,\ e_{t+1}=1$	0.02553642	0.06605507	0.45591327	0.38659293
$s_t=3,\ e_{t+1}=2$	-0.03441756	0.02664083	-1.09440228	-1.29191001
$s_t=3,\ e_{t+1}=3$	-0.18673725	0.01732835	-10.63689745	-10.77640100
$s_t=3,\ e_{t+1}=3,\ y_t<4$	-0.08231817	0.01151675	-7.26648962	-7.14769260
$s_t=3,\ e_{t+1}=3,\ y_t>15$	-0.11454721	0.02976604	-6.99269653	-3.84825189
$ms_t=2,\ ms_{t+1}=2$	-0.19961059	0.04176342	-5.83765826	-4.77955591
$ms_t=1,\ ms_{t+1}=2$	-0.30755954	0.04933132	-8.25716818	-6.23456987
$ms_t=1,\ ms_{t+1}=1$	-0.04921835	0.03813186	-1.56720305	-1.29074065
$ss_t\neq0, ms_t=2, ms_{t+1}=2^{*}$	-0.04721417	0.04996894	-1.19327620	-0.94487044

$ss_t \neq 0,\ e_{t+1}=1^*$	0.04971474	0.02732560	1.40854924	1.81934668
$ss_t \neq 0,\ e_{t+1}=2^*$	0.47145927	0.04392680	9.16446610	10.73284009
$ss_t \neq 0,\ e_{t+1}=3^*$	0.52746853	0.02280896	28.06081154	23.12550001
$ss_t = 2,\ e_{t+1}=1^*$	-0.01470458	0.03635928	-0.29965741	-0.40442436
$ss_t = 2,\ e_{t+1}=2^*$	0.00994914	0.04475291	0.15380495	0.22231272
$ss_t = 2,\ e_{t+1}=3^*$	0.19631143	0.02480131	8.38115753	7.91536464
$ss_t \neq 0,\ a_t \geq 70,\ e_{t+1}=1^*$	0.36644945	0.15467874	2.16755253	2.36910033
$ss_t \neq 0,\ a_t \geq 70,\ e_{t+1}=2^*$	0.33556980	0.08699671	2.38721172	3.85727015

log likelihood -2.91266614E+005 initial log likelihood -2.94748686E+5
grad*direc 2.03125379E-026 number of observations 39494

*These variables are all multiplied by $1/\ln(y_t)$.

(Table 9.1 about here)

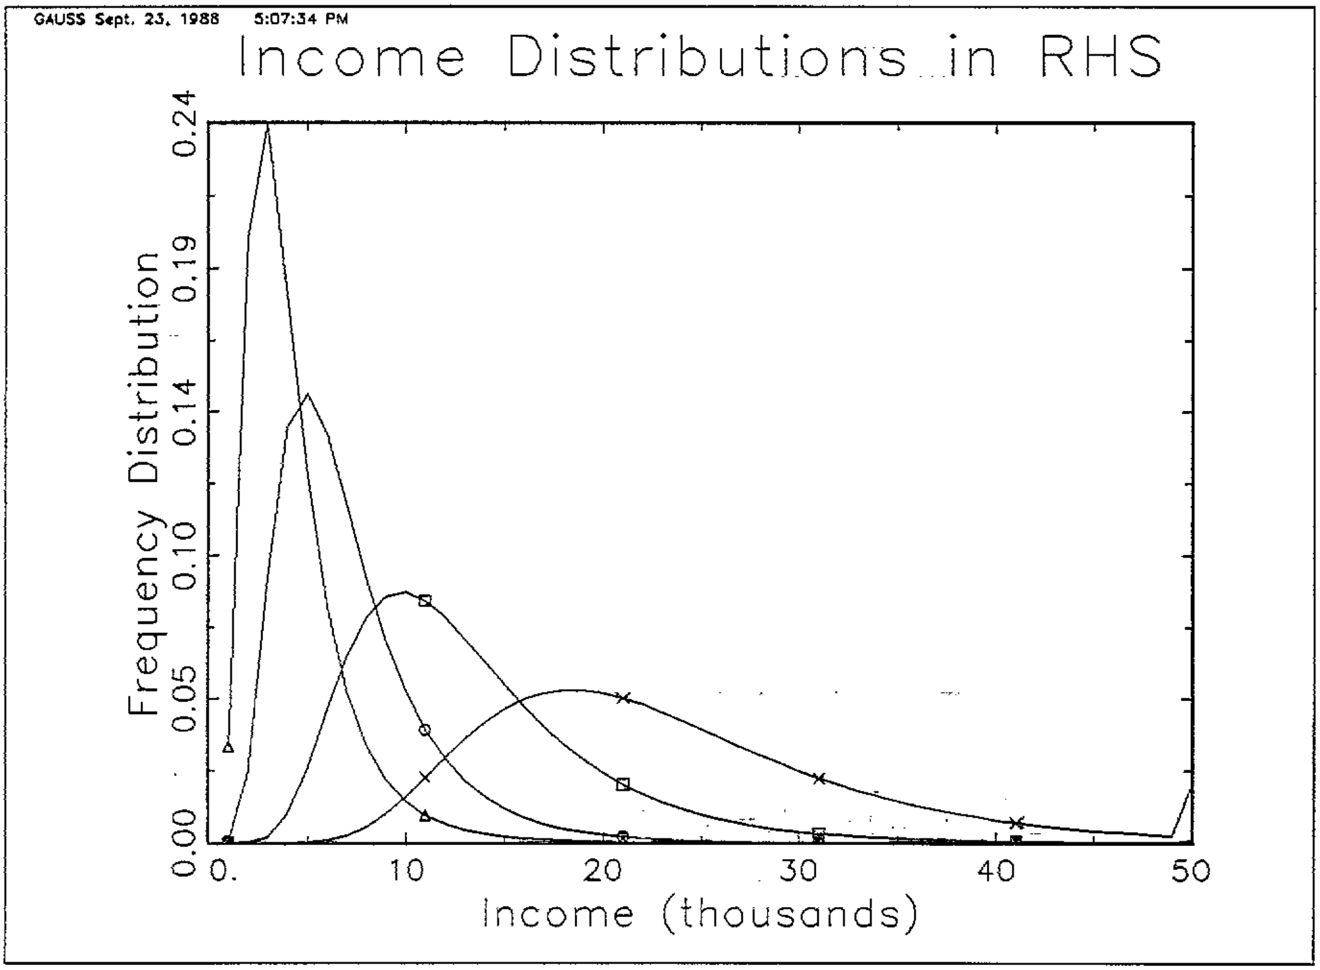
GAUSS Sept. 23, 1988 5:07:34 PM
Income Distributions in RHS
Frequency Distribution
0.00
0.05
0.10
0.14
0.19
0.24
0.
10
20
30
40
50
Income (thousands)

Figure 9.1

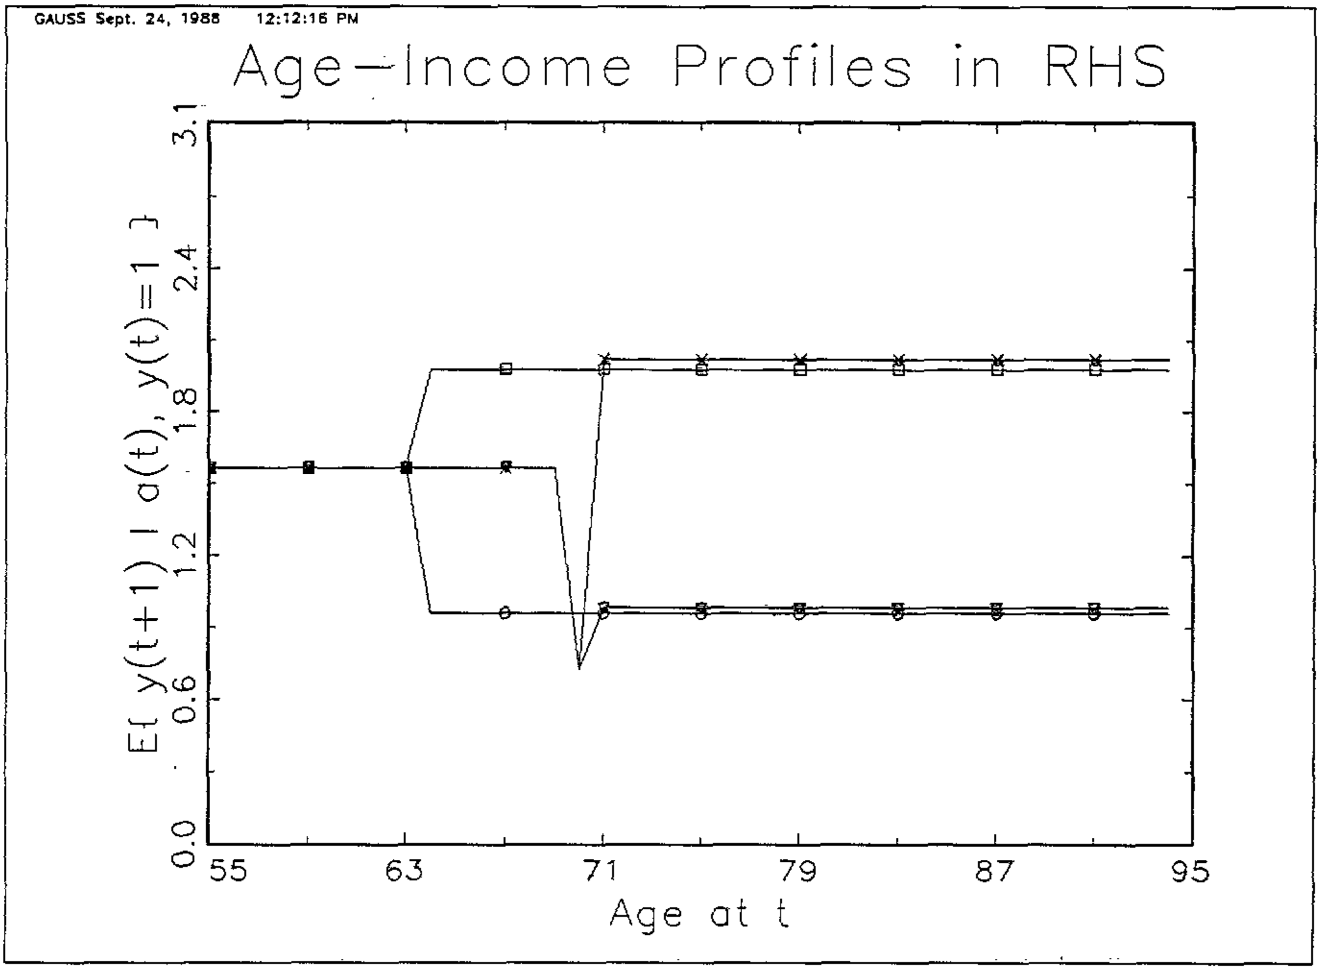

GAUSS Sept. 24, 1988 12:12:16 PM
Age-Income Profiles in RHS
E{ y(t+1) | a(t), y(t)=1 }
0.0
0.6
1.2
1.8
2.4
3.1
55
63
71
79
87
95
Age at t

Figure 9.2


GAUSS Sept. 24, 1988 12:17:57 PM
Age-Income Profiles in RHS
E{ y(t+1) | a(t), y(t)=3 }
0.
1
2
3
4
5
55
63
71
79
87
95
Age at t

Figure 9.3

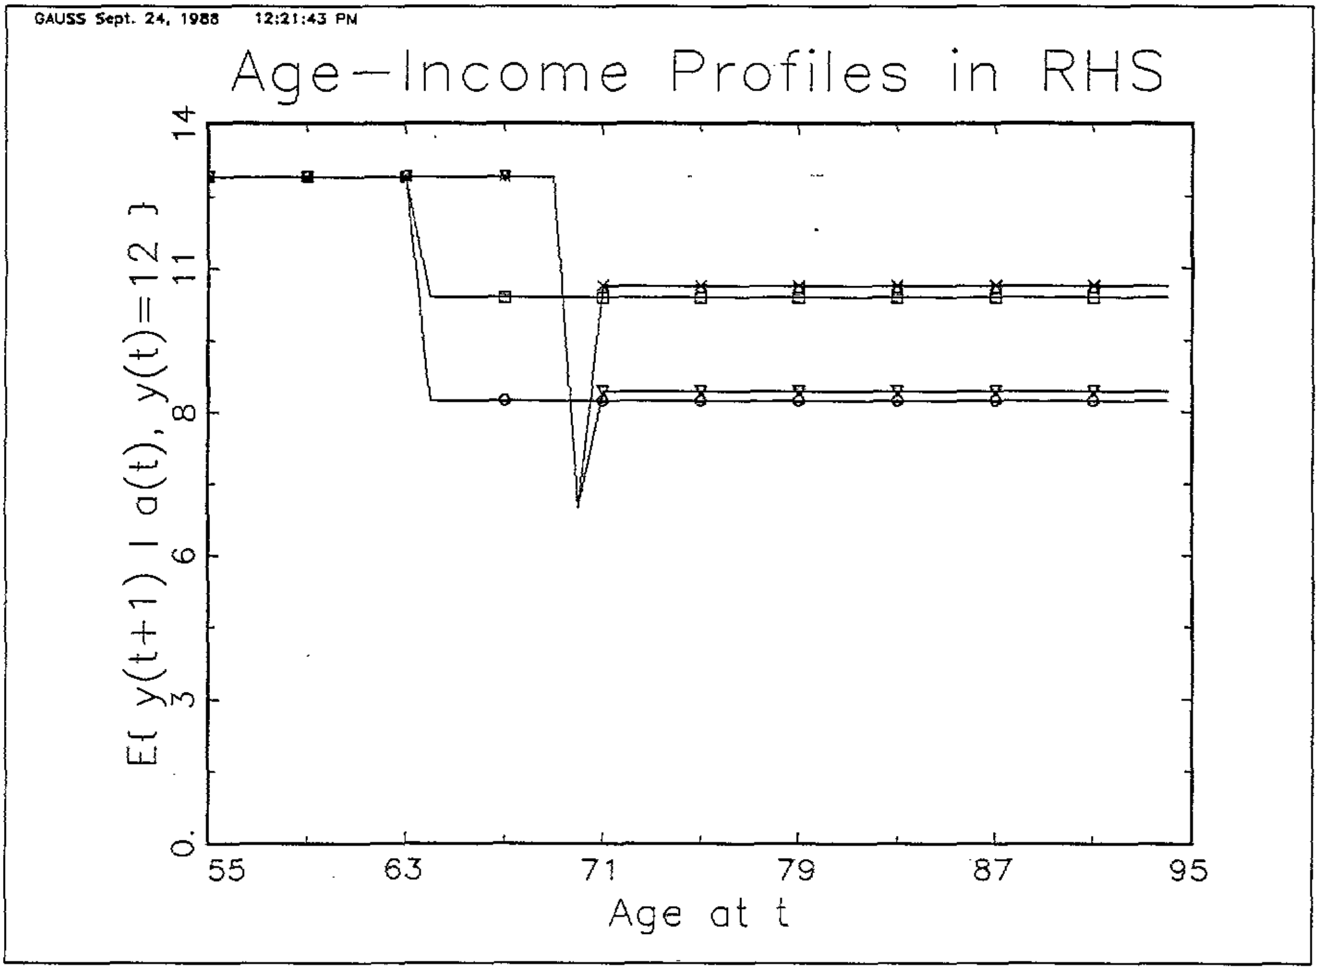
GAUSS Sept. 24, 1988 12:21:43 PM
Age-Income Profiles in RHS
E{ y(t+1) | a(t), y(t)=12 }
0.
3
6
8
11
14
55
63
71
79
87
95
Age at t

Figure 9.4

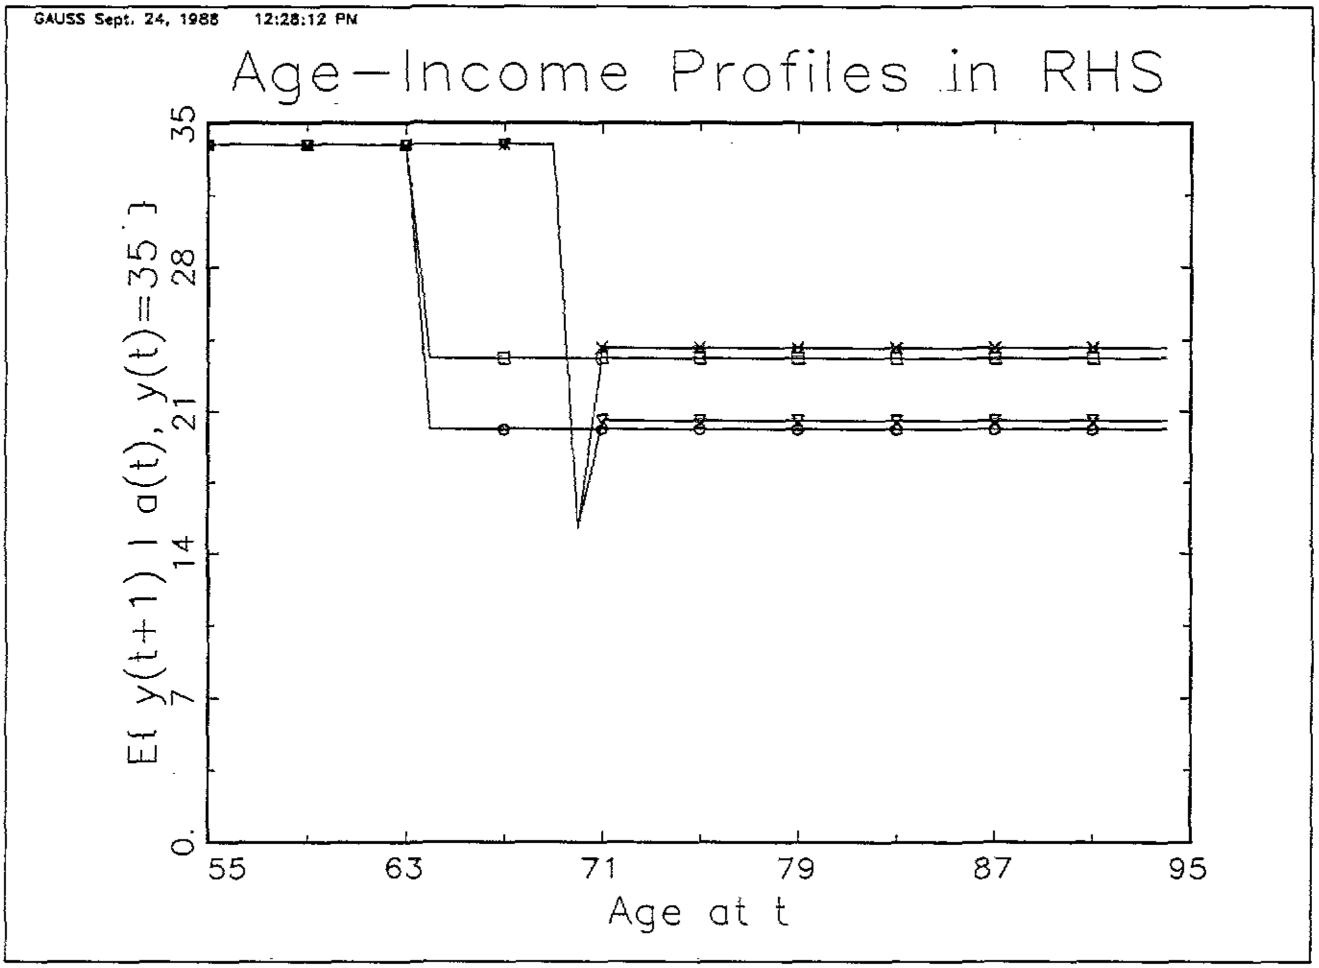
GAUSS Sept. 24, 1988 12:28:12 PM
Age-Income Profiles in RHS
E[y(t+1) | a(t), y(t)=35]
35
28
21
14
7
0.
55
63
71
79
87
95
Age at t

Figure 9.5

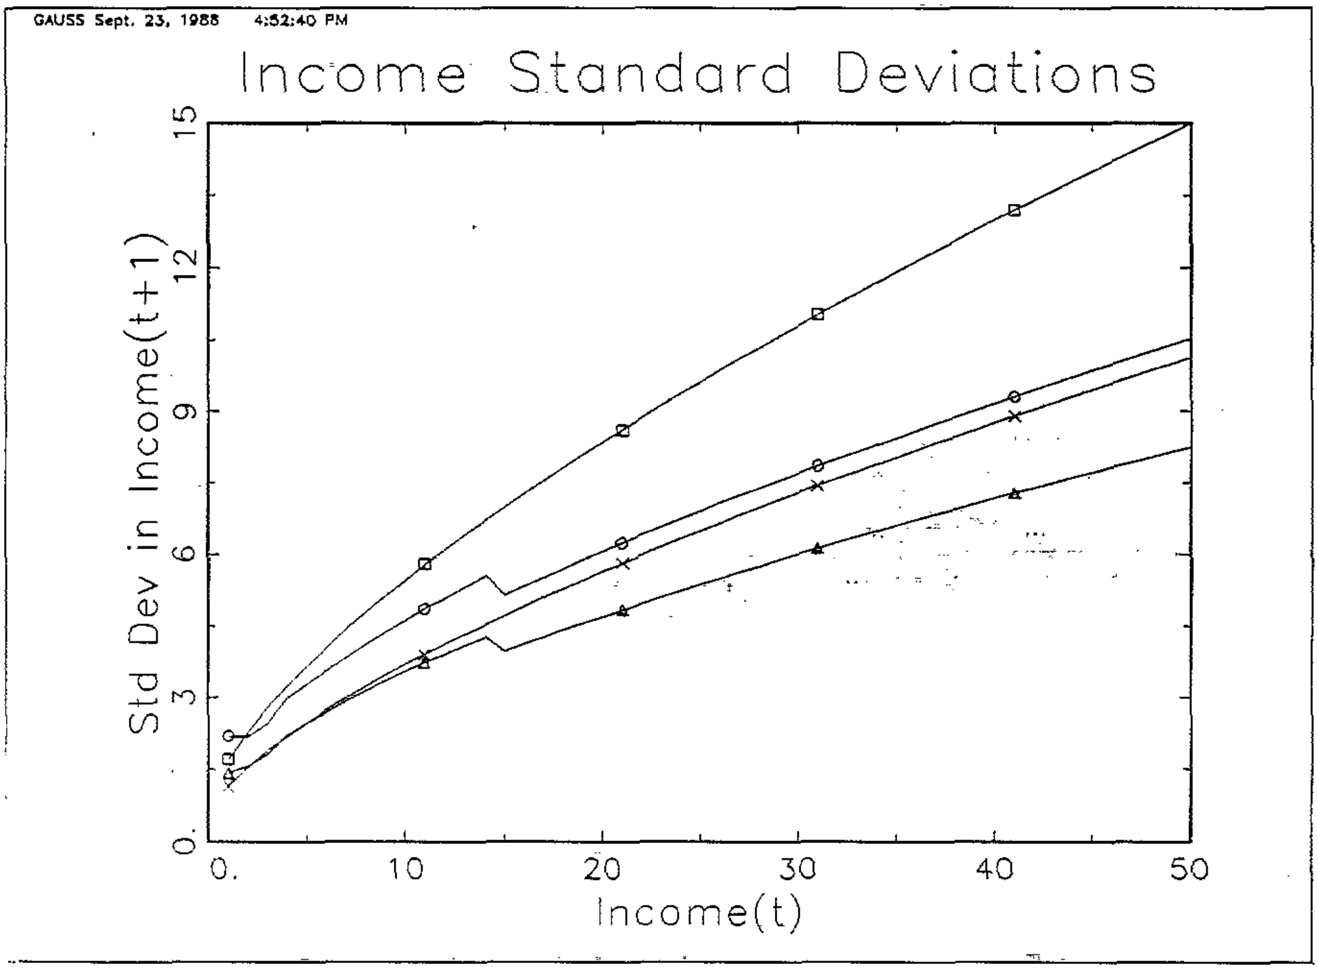
GAUSS Sept. 23, 1988 4:52:40 PM
Income Standard Deviations
Std Dev in Income(t+1)
0.
3
6
9
12
15
0.
10
20
30
40
50
Income(t)

Figure 9.6

Figure 9.1 presents the estimated transition densities $\hat{\pi}_y$ for four configurations of the conditioning variables listed in (3.2), or in plainer language, the subjective distributions of $\tilde{y}_{t+1}$ for four different workers. The sharply peaked density marked with triangles represent the expectations of single man, aged 75, who is disabled, out of the labor force, and receiving a total income of y_t=\$4,000. The density marked with the circles corresponds to a 65 year old retired man who is married, in health state $h_t=2$, and receiving an income of y_t=\$7,000. The density marked with boxes corresponds to a married 58 year old man who is in good health, working full-time with a total income of y_t=\$12,000. Finally the curve marked with x's corresponds to a wealthier 80 year man who is not retired, married, in good health and continues to work full-time earning an income of y_t=\$20,000.

Figures 9.2 to 9.5 present workers' income expectations, $E\{y_{t+1}|e_{t+1},ms_{t+1},h_{t+1},x_t,d_t,a_t\}$, plotted as a function of the age of the worker. Although the profiles are flat by construction (the estimated model excluded a_t), the figures provide an indication of the dynamics of income as workers retire. Each of the figures contain four curves, corresponding to four different retirement paths. The curves marked with boxes correspond to working full-time up until age 65 and then collecting OASI. The curve marked with circles shows what the worker would expect if he quit working but did not start collect OASI. The other two curves represent the expectations of a worker who works full-time up until his early 70's, but then becomes disabled and has to quit work. The lower curve represents what the worker would expect if there were no OASDI program to cover him, the higher curve represents what the worker would expect if he applied for OASDI. Note carefully that the curves in figures 9.2 to 9.5 represent conditional expectation functions: they are not the same as the sample paths of the income process. Given the strong autocorrelation in income, actual sample paths of income may look quite a bit different. The figures clearly show the progressive nature of the Social Security system. In figure 9.2, a very low income worker actually expects to do better by retiring and collecting OASDI than continuing to work at his low paying full-time job. However for a very high income worker figure 9.5 shows that the percentage replacement rate of OASDI benefits is much smaller: Social Security is not such a good deal for these workers. Figures 9.3 and 9.4 present intermediate cases for workers earning $3,000 and $12,000, respectively.

Finally, figure 9.6 plots out the standard deviation of $\tilde{y}_{t+1}$ as a function of current income, y_t. The four curves are all upward sloping, representing the fact that the higher worker's current income is, the more uncertain he is about his future income. Note that while uncertainty does increase with y_t, the increase is far from proportional: this is a direct consequence of the fact that $\ln(y_t)$ enters the σ parameter with a large, significant negative coefficient as you can see from table 9.1. The four curves in figure 9.6 correspond to four classes of workers. The curve marked with boxes corresponds to a 60 year old worker who is married, in good health and working full-time. The curve marked with circles corresponds to a worker who is 88 and disabled and out of the labor force. The curve marked with triangles corresponds to a worker who is 68, single, in health state h_t=2, and is retired and receiving Social Security. The final curve, marked with x's, corresponds to a 55 year man who is single, in health state h_t=2, and working part-time.

10. Modelling the retirement decision

The SSMBR data allow me to determine exactly when a worker applys for and receives OASI benefits. In my opinion, the only sensible and precise definition of the concept of "retirement" is to define it in terms of collection of OASI benefits. I used the SSMBR data set to construct the control variable ss_t defined in (2.2). Figure 10.1 summarizes this variable in terms of the implied distribution of age of first receipt of OASDI. The distribution has a pronounced bimodal shape, with peaks at the early retirement age of 62 and at the normal retirement age of 65. Overall, over 60% of the sample retires between the ages of 62 to 65. The distribution differs significantly from a frequency distribution tabulated by Burtless and Moffitt (1984) who used the RHS data, the instantaneous measure of labor force participation IE, and a definition of "retirement" to be a sudden, discontinous drop in labor supply to under 30 hours per week. Burtless and Moffitt's frequency distribution for the age of retirement is presented in figure 10.2. My estimate of the retirement frequency distribution also differs signficantly from the estimate of Sueyoshi (1986), who also used the RHS but adopted an hours-of-work based definition of retirement and partial retirement.Sueyoshi's estimate of the retirement age distribution is presented in figure 10.3.The main difference between the "age of first receipt of OASDI" definition of retirement and the hours of work definitions is the pronounced peak of early retirees that shows up under the former definition. Apparently, there a substantial fraction of workers apply for OASI at age 62, but continue working until age 65. Only by separating the concepts of "retirement" and "labor supply" will we be able to understand this behavior.

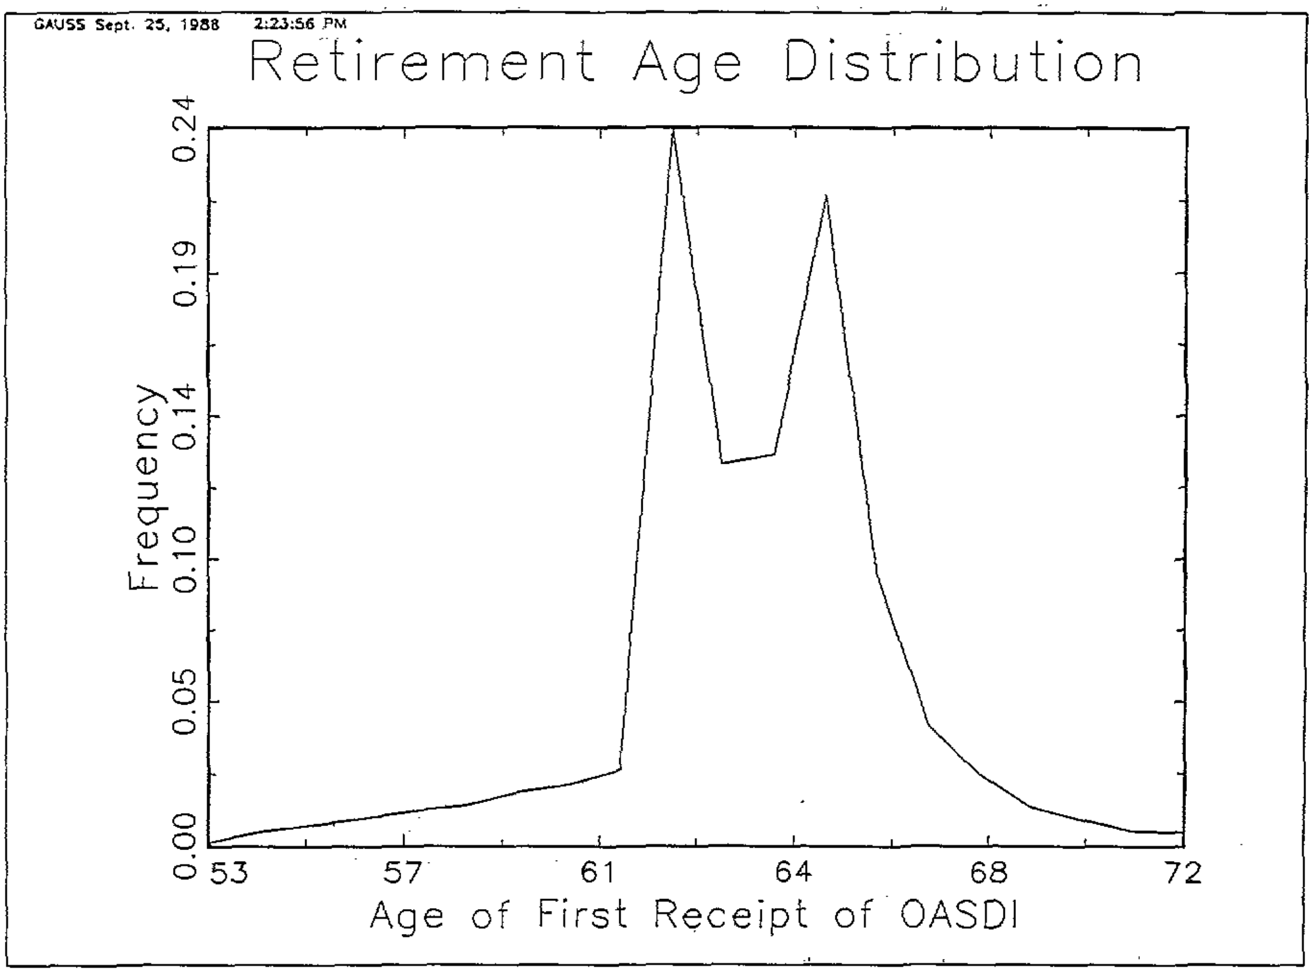

GAUSS Sept. 25, 1988 2:23:56 PM
Retirement Age Distribution
Frequency
0.00
0.05
0.10
0.14
0.19
0.24
53
57
61
64
68
72
Age of First Receipt of OASDI

Figure 10.1

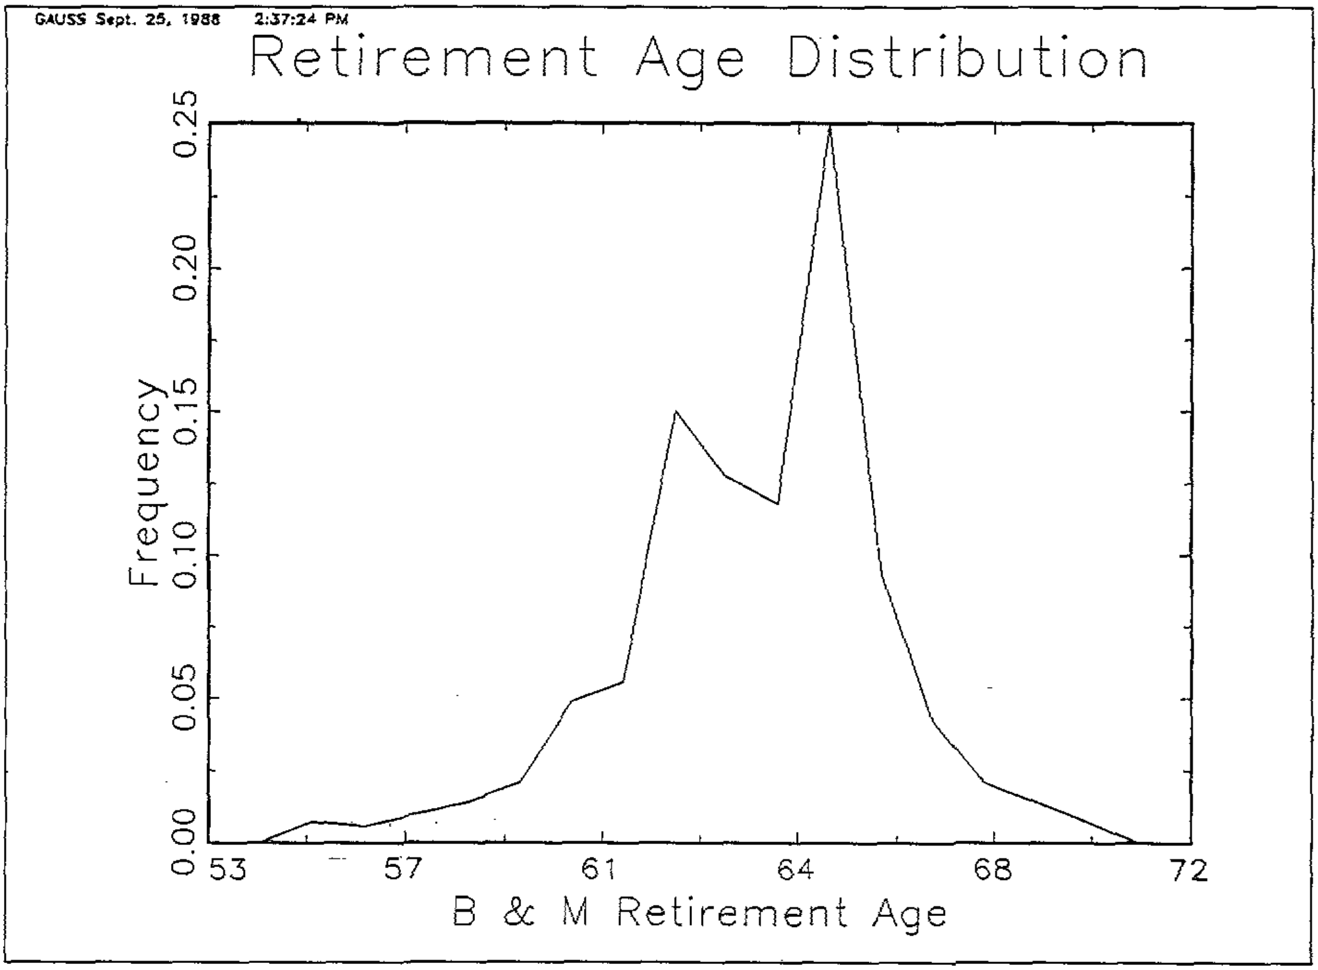
GAUSS Sept. 25, 1988 2:37:24 PM
Retirement Age Distribution
Frequency
0.00
0.05
0.10
0.15
0.20
0.25
53
57
61
64
68
72
B & M Retirement Age

Figure 10.2

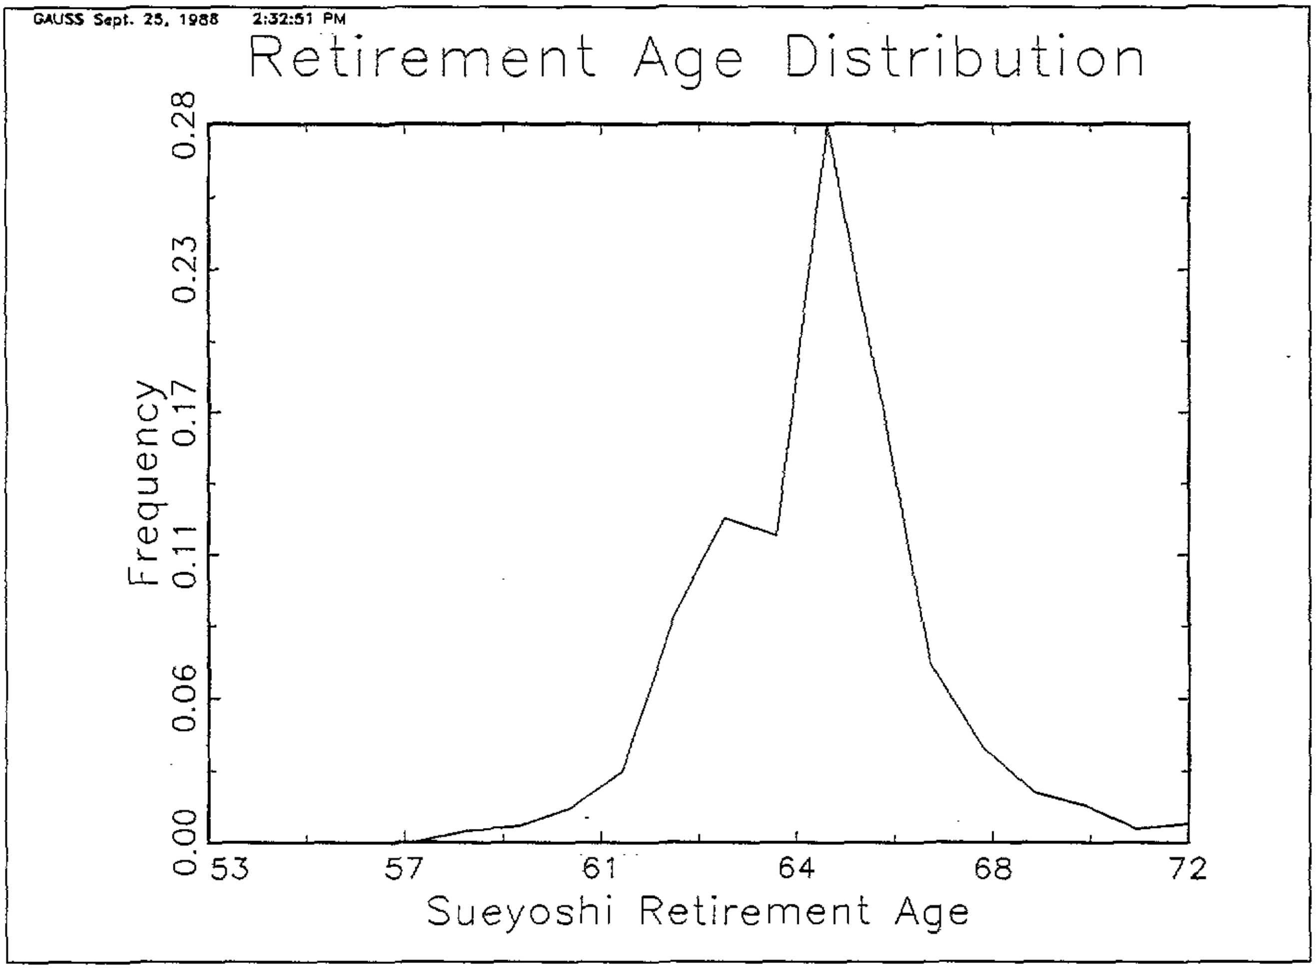
GAUSS Sept. 25, 1988 2:32:51 PM
Retirement Age Distribution
Frequency
0.00
0.06
0.11
0.17
0.23
0.28
53
57
61
64
68
72
Sueyoshi Retirement Age

Figure 10.3

The estimation results in the previous sections of the paper suggest two conclusions about the decision to retire and collect OASI: 1) workers who retire and receive OASI appear to be less healthy than their counterparts who continue to work, and 2) once a worker starts collecting OASI, he is significantly less likely to return to work on either a full or part-time basis. Analysis of the RHS data provide evidence of the role of self-selection in the decision to collect OASDI; less healthy workers are more likely to quit their jobs and retire early, even given the permanent 20% penalty for early retirement. This may be rational behavior given the reduced life-expectancy of unhealthy workers and the well-known fact that the OASI benefit structure is not actuarially fair. A more complete analysis of these issues must await the estimation of the DP model in the third part of this series.

References

Berkovec, J. and Stern, S. (1987) Job Exit Behavior of Older Men. University of Virginia, typescript.

Bertsekas, D. and Shreve, S. (1978) *Stochastic Optimal Control: The Discrete Time Case* New York, Academic Press.

Bound, J. (1986) The Disincentive Effects of the Social Security Disability Program. University of Michigan, typescript.

Burtless, G. and Moffitt, R. (1984) The Effect of Social Security Benefits on the Labor Supply of the Aged in H. J. Aaron and G.Burtless (eds.) *Retirement and Economic Behavior*. Washington, D.C. Brookings Institution.

Gustman, A.L. and T.L. Steinmeier (1983) Minimum Hours Constraints and Retirement Behavior *Economic Inquiry* 3, 77-91.

Gustman, A.L. and T.L. Steinmeier (1984) Partial Retirement and the Analysis of Retirement Behavior *Industrial and Labor Relations Review* 37, 403-15.

Gustman, A.L. and T.L. Steinmeier (1986) A Structural Retirement Model *Econometrica* 54-3, 555-84.

Hamermesh, D.(1984) Consumption During Retirement: The Missing Link in the Life Cycle *Review of Economics and Statistics* 66, 1-7.

Hansen, L.P. (1982) Large Sample Properties of Generalized Method of Moments Estimators *Econometrica* 50, 1029-1054.

Hill, D. (1988) Response Error Around the Seam: An Analysis of Change in a Panel with Overlapping Reference Periods. University of Michigan, typescript.

MaCurdy, T.E. (1983) A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty *International Economic Review* 24-2, 265-89.

Mott, F.L. and Haurin, R.J. (1985) Factors Affecting Mortality in the Years Surrounding Retirement in H.Parnes (ed.) *Retirement Among American Men* Lexington Press.

Parnes, H.S. (ed.) (1981) *Work and Retirement* Cambridge, M.I.T. Press.

Rust, J. (1988a) A Dynamic Programming Model of Retirement Behavior in D.Wise (ed.) *The Economics of Aging.* Chicago, University of Chicago Press, 359-398.

Rust, J. (1988b) Maximum Likelihood Estimation of Discrete Control Processes *SIAM Journal on Control and Optimization* 26-5, 1-19.

Stock, J.H. and Wise, D.A. (1988) Pensions, The Option Value of Work, and Retirement, J.F.K. School of Government, Harvard University, typescript.

Sueyoshi, G. (1986) Social Security and the Determinants of Full and Partial Retirement: A Competing Risks Analysis. University of California-San Diego, typescript.

U.S. Census (1979) *Statistical Abstract of the United States*, Washington, U.S. Government Printing Office.

van Dijk, N.M. (1984) *Controlled Markov Processes: Time Discretization* CWI Tract II, Amsterdam, Mathematische Centrum.

White, H. (1982) Maximium Likelihood Estimation of Misspecified Models *Econometrica* 50, 1-26.