

Retirement Income Security and Well-Being in Canada

Michael Baker, University of Toronto and NBER
Jonathan Gruber, MIT and NBER
Kevin Milligan, UBC and NBER

May, 2004

Preliminary Draft – Not for Citation – Comments Welcome

Expenditures on social security in Canada are projected to increase substantially over the next decades. For example, expenditures on the primary transfer program for seniors, which totaled over \$23 billion in 1999/2000, are projected to rise from \$25 billion in 2001 to \$109 billion by 2030, or from 2.3 percent of GDP to 3.2 percent (Office of the Superintendent of Financial Institutions 2002a). Benefit expenditures in the largest of the two public pension plans, which totaled over \$20 billion in 2000/01 are projected to rise to \$74 billion in 2025 (Office of the Superintendent of Financial Institutions 2002b). As in many other developed nations, the forecasted increase in expenditures has led many to question the long term health of these programs.

There are a variety of solutions possible to this long-term problem. Many of them involve reducing, in one way or another, the benefits available to retirees in Canada. For example, the government could cut the OAS amount, change the translation of past earnings into CPP/QPP benefits, or raise program entitlement ages. As discussed in Baker, Gruber and Milligan (2003, 2004), these changes could have significant impacts on both retirement behavior and program finances. For example, we find that raising the age of eligibility for retirement programs in Canada would lower program expenditures by over 22%.

While improving program finances, such a change could also have pernicious effects on the well-being of the elderly in Canada. The income from retirement programs made up 45.5% of the average elderly family's gross income in Canada in 2000. If benefits from these programs are reduced, it could lead to significant reductions in the standard of living of elderly families.

But would it? This depends on the reaction of other sources of support to the elderly. For example, a large literature has investigated the question of whether social security benefits

simply serve to crowd out savings by the elderly. Other sources of income support for the elderly include their own labor supply or transfers from other family members. If these other sources of support are increased as Social Security benefits are reduced, then there may be little implications of benefits reductions for elderly well-being.

In this paper, we investigate directly the “other side” of Social Security reform: the implications for elderly well-being. We do so by using the large variation across birth cohorts in income security entitlements in Canada that arise from reforms to the program over the past 35 years. This variation allows us to explore the effects of benefits on elderly well-being while controlling for other factors that affect well-being over time and by age.

The paper proceeds as follows. In Section I, we provide a brief background on the institutional features of the income security programs for Canadian elders. In section II, we describe our data sources. Section III discusses our empirical strategy for using variation in benefits across birth cohorts to identify the effect of benefits on well-being. Section IV presents time series evidence on the role of Social Security, and Section V presents regression evidence. Section VI concludes with the implications of our findings for Social Security policy.

Section I: Background on the Canadian Income Security System

Public transfer/pension plans for seniors in Canada fall into two groups: 1) the Old Age Security program, and 2) the Canada and Quebec Pension Plans. The Old Age Security program, which encompasses the Old Age Security (OAS) pension, the Guaranteed Income Supplement

(GIS), and the Allowance, are transfer programs financed out of general tax revenues. The Canada and Quebec Pension Plans (CPP/QPP) are contributory programs that pay a benefit which is proportional to an individual's earnings over his/her working life.

The Current Parameters

OAS is the older of the two programs, dating back to 1952. Its current rules, which have been relatively stable since the 1970s, provide for a pension payable to all individuals aged 65 or older who satisfy a residency requirement.^a In June 2004 the monthly benefit paid to individuals who fully satisfied the residency requirement was \$463.39. This benefit is clawed back from "high income" pensioners at a 15 percent rate, starting at incomes of \$59,790 (2004). Benefits are full indexed to the CPI and fully taxable under the Income Tax Act.

Available since 1967, the GIS component of OAS is an income tested supplement that is also payable to those aged 65 or older. The income test is applied annually based on taxable income (as defined by the Income Tax Act, excluding any OAS pension) at the family level. Benefits are taxed back at a 50 percent rate, except in families where only one partner is 65 or older, in which case the tax rate is 25 percent. There are separate monthly benefits for married and single individuals, fully indexed to the CPI, which amounted to \$550.73 and \$358.73, respectively, in June 2004. These benefits are not subject to income taxes.

Finally, the Allowance (formerly the Spouse's Allowance) is an income tested benefit payable to 60-64 year old partners of individuals who are aged 65 or older. It is also available to

^a Individuals must have been a Canadian citizen or legal resident of Canada at some point before application, and have resided in Canada for at least 10 years (if currently in Canada) or 20 years (if currently outside Canada). The benefit is prorated for pensioners with less than 40 years of Canadian residence, unless they are "grandfathered"

widow/widowers who are aged 60-64. This part of OAS was introduced in 1975. For the partners of 65+ year olds the benefit equals the sum of an OAS pension plus a GIS at the married rate. Benefits are reduced by 75 cents for each dollar of income to the OAS part of the benefit is eliminated. At this point the benefit (as well as the partner's GIS) is reduced according to standard GIS rules. Widow/widowers receive a slightly higher benefit the benefit, but the income test is applied equivalently. Again benefits are fully indexed and not subject to income taxes.^b

CPP and QPP pensions, available since 1966, are financed by employer and employee contributions. In 2004, these are (each) payable at a rate of 4.95 percent on earnings between \$3500 and \$40,500 (the Year's Maximum Pensionable Earnings, YMPE). Pensions are available to individuals aged 60 or older who have made contributions in at least one calendar year in the ("contributory") period stretching from age 18 to the year of application (or age 70).^c Benefits are actuarially adjusted for applications at different ages in the interval 60 through 70. The adjustment is a reduction (premium) of 0.5 percent for each month the application precedes (follows) the 65th birthday. Also, applications prior to age 65 are subject to a retirement test, which is that the individual's annual rate of earnings cannot exceed the maximum retirement pension payable at age 65, for the year in which the pension is claimed.^d

The amount of benefits received is based on the individual's "average" earnings over the contributory period. Any months (a) receiving a disability pension, (b) spent rearing small

under rules that apply to the persons who were over age 25 and had established attachment to Canada prior to July 1977.

^b More details on the Allowance are available in Baker (2002).

^c For individuals who were 18 prior to 1966, the contributory period starts on January 1, 1966.

children,^e (c) between age 65 and the commencement of the pension^f, and (d) 15 percent of the remaining months, are excluded from the calculation.^g In each of the remaining months, earnings are expressed as a fraction of 1/12 of the current YMPE, to a maximum of one.^h The average of these fractions is then multiplied by the average value of the YMPE over the previous five years (including the year of application), and then by 0.25, to arrive at the benefit level. Benefits are fully indexed and subject to income taxes.

Both the CPP and QPP pay survivor pensions when a contributor dies and has made contributions for the lesser of 10 years or one third of the number of years in the contributory period.ⁱ For survivors under age 65, benefits are equal to a flat rate benefit plus 37.5 percent of the earnings-related pension of the deceased spouse. Benefits are reduced if the claimant is younger than age 45 and is not disabled and has no dependents. For survivors aged 65 and above, the pension is equal to the greater of a) 37.5 percent of the deceased's retirement pension plus 100 percent of the survivor's own retirement pension, or b) 60 percent of the deceased's retirement pension plus 60 percent of the survivor's own retirement pension. There is an upper

^d There are no restrictions on returning to work after the benefit is being paid.

^e This is defined as months where there was a child less than 7 years of age and the worker had zero or below average annual earnings.

^f Periods after age 65 to age 70 can be substituted for periods prior to age 65 if this will increase their future retirement pension.

^g The last three of these exclusions cannot be used, however, to reduce the contributory period below 120 months after taking into account the offset for months of disability pension receipt.

^h Excess earnings in one month above 1/12 of the YMPE may be applied to months in the same calendar year in which earnings are below 1/12 of the YMPE.

ⁱ There is also a lump sum death benefit, which is generally equal to one-half of the annual CPP/QPP pension amount up to a maximum \$2,500.

cap on total payments equal to the maximum retirement pension payable in that year.^j There are also orphan benefits payable to the children of the deceased.^k

The CPP and QPP also pay disability pensions. The benefit is equal to a flat-rate portion plus an earnings-related portion equal to 75 percent of the applicable CPP/QPP retirement pension, calculated with the contributory period ending at the date of disability. More information can be found in Gruber (2000).

Program Changes over Time

Changes to the OAS and CPP/QPP programs over time are the basis of identifying any impact of social security on well being. These changes have altered both the level of benefits available to claimants, as well as the eligibility criteria.

The age of eligibility for OAS pensions was changed over the period 1965 to 1970. Initially these pensions were available starting at age 70 but starting in 1965 the age was lowered annually in one year increments until it reached 65 in 1970. Full indexation of benefits did not start until 1973. Previously, adjustments to benefits were ad hoc. Finally, the residency requirements for benefits were changed in 1977. Originally individuals who did not fully satisfy the requirement received no benefit, but starting in 1977 pro-rated benefits were granted for partial fulfillment.

The GIS was introduced in 1967. Benefits were indexed starting in 1973. There were also one time real increases in benefits of 69 percent in 1971, 20 in 1972, 44 percent in 1979/80

^j If the surviving spouse is receiving his or her own CPP disability pension, the sum of the earnings-related portion of the two pensions cannot exceed the maximum retirement pension available in the year.

^k More details on survivor pensions are available in Baker et al (2004).

and 16 percent in 1984. The introduction of the program, as well as these benefit changes, should have heterogeneous impacts across birth cohorts. For example, older cohorts who missed out on CPP/QPP pensions (introduced in 1966) are more reliant on GIS than younger cohorts whose GIS is taxed back because of these pensions. Also, the introduction of the Allowance in 1975 led to a dramatic temporal change in the public support available to 60-64 year olds. This support was extended to widows aged 60-64 in 1985.

The introduction of the CPP/QPP in 1966 had a differential impact across birth cohorts. The parameters of the system were phased in over a ten year “transition” period. First, pensions were initially (as of January 1967) payable starting at age 68. The minimum age was lowered in subsequent years until it settled at age 65 in 1970. Second, pensions were prorated. The benefit was calculated as the standard pension entitlement multiplied by the fraction of the transition period that the individual had made contributions to the plans. For example, if the benefit application was made in January 1972 after 6 years of contributions, the standard pension entitlement would be multiplied by 0.6, which is just the number of months of contributions (72) divided by 120 months, the length of the transition period. This means that those born before 1900 got no pensions, those born between 1901 and 1910 got partial benefits, and those born in/after 1911 were 65 in 1976 and so they got full benefits. The details of the phase-in were announced in 1965, so the 1901 cohort got a windfall, while later cohorts could have anticipated the benefits they would receive.

Starting in 1970, CPP/QPP benefits were available starting at age 65. This changed in 1984 for the QPP and 1987 for the CPP when more flexible rules were introduced. As described

above, these rules allowed application anytime between ages 60 and 70 subject to an actuarial adjustment.

The real value of pensions payable by the CPP/QPP changed quite dramatically over the period 1975 and 1986. Recall that the benefit calculation includes a moving average of the YMPE. While the YMPE was initially set to equal average earnings, it was not indexed. By the mid 1970s it had fallen to roughly 70 percent of the average. In 1975 both plans were amended to allow the YMPE to rise at a rate of 12 1/2 per annum until the average was once again attained in 1986. This change was unanticipated and probably not widely appreciated.

Finally, both the CPP and QPP initially had earnings tests on the pensions of beneficiaries aged 65 to 69. Benefits were recovered at a rate of 50 cents for every dollar of earnings in excess of 18 percent of the YMPE on an annual basis. An additional 50 cents was recovered for earnings in excess of 30 percent of the YMPE. Therefore, at higher levels the tax back was 100%. The earnings test was eliminated from the CPP in 1975. The second tier of the test was eliminated from the QPP in 1973, while the remaining tier was dropped in 1977.

In figure 1 we graph the real value of the different pensions and benefits starting in 1961. Many of the reforms and changes reviewed can be clearly seen in this picture. In particular, the real increases in the GIS in 1971/72, 1978/79 and 1984, the indexing of the OAS pension, the effects of the CPP transition period and the recovery of the YMPE to the level of average wages.

Section II: Data Sources

Our analysis encompasses three aspects of elderly wellbeing: incomes, consumption, and happiness. In all cases, we use survey microdata provided by Statistics Canada. Below, we describe the details for each of our data sources in detail. In addition, we explain exactly how we form the wellbeing variables we use in the analysis from the raw data sources.

There are several issues common to all of the data sets that we use. First, we use 2001 Euros for all of the analysis. We update dollar values to 2001 using the Canadian Consumer Price Index, then translate to Euros using the December 31, 2001 Euro-Canadian dollar exchange rate. Second, we scale all income and consumption data by an equivalence scale to account for the size of the household / family. The first adult is counted as one, each subsequent adult counts as 0.7, and each child under age 18 counts as 0.5. A third issue is the definition and weighting of elderly households. We define an elderly household as one in which there is at least one member age 60 or over. All other households are therefore classified as non-elderly. Furthermore, we weight the results for the elderly households by how many members are age 60 or over, effectively placing our results on an individual basis for the elderly.

Income data

We draw our income data from two different sources. From 1971 to 1997, we have data from the Survey of Consumer Finances. This survey reports information on the education, age, family structure, labour market activity, and incomes of respondents. In particular for income,

the data is disaggregated to a degree that serves our purposes well. In particular, we use the variables for after-tax income, Canada / Quebec Pension Plan income, and OAS/GIS/SPA income to measure the actual public pension benefits received by elderly Canadians.

The Survey of Consumer Finances is conducted using the sampling frame of the monthly Labour Force Survey, which is a stratified random sample of Canadians with some oversampling of smaller provinces to ensure sufficient sample size. With the survey weights, nationally representative results can be obtained.¹ The survey is available in odd years from 1971 to 1981, then 1982, and finally on an annual basis from 1984 to 1997. We discard the 1996 and 1997 surveys because we have other data for those years, described below. The survey is available for different family definitions and for individuals. We use the census family sample, which ranges in size from 25,927 up to 47,840.^m

The Survey of Consumer Finances was replaced in the late 1990s by the Survey of Labour and Income Dynamics. The new survey is available for 1993 to 2000 and for our purposes is similar to the previous survey. One interesting difference is the source for income data. For the Survey of Consumer Finances, income data was asked on a recall basis. For the Survey of Labour and Income Dynamics, however, respondents may check a box indicating consent to attach data directly from tax records. A high percentage of respondents take this option (over 80 percent in 2000).

¹ The sampling frame excluded residents of the three northern territories, those living in institutions, and inhabitants of native reserves. These groups typically account for less than three percent of the population.

^m A census family is comprised of parent(s) with their children, childless couples, or grandparents living with their grandchildren. Individuals on their own who are not in a census family are classified as “non-family individuals”, but are included in the census family sample.

We form five income variables for our analysis. The first is total income from income security pensions, which is the sum of reported CPP/QPP and OAS-GIS-SPA income. Second, we take the reported total after-tax income for the family. Third, we form a measure of relative income poverty. For each year, we find the median family income for non-elderly families. We set a poverty line at 40 percent of this median, then form a dummy variable for each elderly family indicating whether it is over or under the poverty line. Fourth, we measure absolute poverty by repeating the same exercise but using the poverty line for 1971 in all years. Finally, we form a measure of inequality by taking the difference between the 90th and the 10th percentile of family income. This difference is then normalized by dividing by the mean non-elderly income for the year. Combined, these five measures will richly describe the income patterns of elderly Canadian families over the three decades we study.

Consumption data

The consumption data we use comes from two different surveys, each conducted over a number of years. The first survey is the Family Expenditure Survey. It was first conducted in 1969, followed irregularly by surveys in 1974, 1978, 1982, 1984, 1986, 1990, 1992, and 1996. The sampling frame in several of the years included only residents of large cities, so we restrict our sample in all years to residents of large cities to maintain comparability. This survey is also based on the Labour Force Survey sampling frame, so survey weights can produce nationally representative results.

The unit of observation changed for the Family Expenditure survey in 1990. Prior to 1990, the concept used was the spending unit, defined as “a group of people living in the same dwelling who depend on a common or pooled income for major expenses or one financially

independent individual living alone.” This is similar to a definition of an economic family, one or more of which can live in any one dwelling. From 1990 on, the unit of observation was the household, defined simply as those in the same dwelling unit. The sample size varies considerably, with between 4,569 and 15,140 observations.

The second set of consumption data is the annual Survey of Household Spending, available for the years 1997 through 2000. This survey has larger sample sizes than the Family Expenditure Surveys did, with around 14 and 18 thousand households per year.

The key variable we extract from these surveys is the current consumption of the household. Included in this measure is spending on goods and services throughout the year. It is equivalent to the total expenditure of the household less personal taxes, gifts and donations, and life insurance. In addition to the current consumption variable, we construct three other measures of elderly wellbeing similar to those we made for income. There is a relative consumption poverty measure constructed as being under 40 percent of the median non-elderly household’s consumption level in each year, an absolute measure based on the median non-elderly household in 1969, and a 90-10 inequality measure normalized by mean non-elderly consumption.

Happiness data

To undertake our analysis of the happiness of elderly Canadian households, we put together several waves of the General Social Survey. The sample is formed using a random digit dialing methodology, with supplemental samples drawn from the Labour Force Survey sampling frame in some years. Sample sizes are around ten thousand observations per year. The

happiness question that we require is asked in the years 1985, 1986, 1989-1991, 1996, and 1998. Unfortunately, age is only reported in 5-year ranges, so we cannot identify the exact year of age of the respondents. As the happiness question is individually based, we use the individual as the unit of observation for the happiness data.

The exact wording of the question is “presently, would you describe yourself as . . .” Possible responses are very happy, somewhat happy, somewhat unhappy, and very unhappy. We form two variables to use for our analysis. First is an indicator for having responded “very happy.” The second is an indicator for having responded either somewhat or very unhappy.

Section III: Empirical Strategy

For the regression analysis that forms the core of this paper, we want to regress well-being outcomes on retirement income. Because observed retirement income may be determined by the same factors that determine the outcomes we wish to study, we create *simulated benefits* that are exogenous to the outcomes. We do this by attempting to capture legislative variation in the system over time. In general, benefits are composed of two components: laws and the characteristics of recipients. Laws come in two forms: the euro amount of the benefits and the age at which benefits are available. The idea of the simulated benefit approach, in its ideal form, is to abstract from all differences in characteristics of recipients and focus solely on variation in benefits that arises from law changes. Ideally, we would take exactly the same person, put him in every single birth cohort, and then compute his benefits. In this way, we would hold

characteristics constant, and any benefits variation that we saw over time or across birth cohorts would have to be due to law changes.

Of course, this is impossible in reality. Two types of factors may differ across birth cohorts that can affect benefit determination. The first type includes factors which are unlikely to be endogenous to Income Security rules (at least to a first approximation), but which are potentially important omitted determinants of well-being: differences in earnings histories; differences in capital income (which can matter for GIS eligibility and benefits); and differences in spousal labor supply. The second type is factors which are likely endogenous to Income Security rules, such as differences in actual observed retirement ages. The approximately ideal simulated benefits approach would hold the first type constant, and compute benefits for each birth cohort based on that constant comparison. A more difficult question is whether one wants to hold the second set of factors constant, as this may be part of the effect that we want to capture in our results.

Thus, we consider three approaches. First, we show results using actual benefits received, as measured by reported IS income at each age. Second, we employ a “mixed simulation” approach, in which the first set of factors is held constant, but the second set of factors is allowed to vary. In particular, we base the benefit amount on a fixed earnings history across all birth cohorts, not the actual earnings history by cohort. The motivation for this mixed simulation approach is to take account of changes in retirement behaviour that may be responses to changes in program structure. At the same time, by using the fixed earnings history the mixed approach allows for the fact that there could be correlated changes over time in tastes for leisure (as manifested in early retirement) and measures of well-being that could bias the results if each

cohort's actual earnings were used. Finally, we also consider a “pure simulation” approach, whereby both the first and second sets of factors are held constant. For that approach, we use the base cohort earnings history *and* the base cohort retirement patterns.

Methodology

In order to impute benefits to a retired family we must know their earnings history, their capital income, at what age they retired, their marital status, and the rules that were in place that determined their benefits. Combined, this information allows us to calculate ‘typical’ benefits for families in our data. We describe how we approach each of these factors in turn.

We form the earnings profile by taking deciles of earnings in each year and age in the Survey of Consumer Finance data over all individuals, separately for men and women. We then average these deciles to form an earnings value for each age and year. The motivation for this decile-based approach is to capture non-work in the best way possible. Median earnings of women in earlier years were zero, for example, so using the decile approach allows us to capture the fact that at least some of the women in those years had earnings. Since the oldest income microdata available to us is 1971, we impute earnings to each year previous to 1971 by deflating the 1971 age-earnings profile by the growth in the industrial composite wage. The cohort we chose to use for the simulations is the 1920 cohort, which is aged 65 in 1985, the middle year of the period our data covers.

For capital income, we use data from the Survey of Consumer Finances. We define capital income as the sum of income from investments and income from retirement pensions, which includes both withdrawals from Registered Retirement Savings Plans and Registered

Retirement Income Funds. We form deciles of capital income for each age and year by sex. The decile approach allows us to capture nonlinear effects of the income security system that would be missed if we simply used the median or mean for assigning capital income. We assign the capital income deciles to the earnings deciles by age, year and sex, assuming that the lowest earning decile has the lowest capital income, the 2nd lowest earners with the 2nd lowest capital income, and so on. At the end of the calculation, we average over the deciles to arrive at the simulated benefit.

We do not observe the retirement data for elderly Canadians in our data. Because CPP/QPP benefits depend on the date of retirement through the actuarial adjustment of early retirement benefits, we must account for the possibility of different retirement dates. We do so by calculating the retirement benefits that would be received at each potential age of retirement between the ages 55 and 70. We then average these benefits using the observed retirement probabilities at each age for the cohort as weights. We form the cohort retirement probabilities with the Survey of Consumer Finances data. We assume that none are retired at age 55 and that all are retired by age 70. The retirement rate is the change in the proportion of respondents who are employed in one year over the previous year. Using these observed rates, we form a cohort- and sex-specific set of probabilities for retirement that sum to one. With these probabilities, we then determine a retirement age-weighted average of the benefits the family is currently receiving.

Finally, we must consider different family types. Someone we observe at an older age may have previously been married but is now widowed. For such a person, we must assign spousal benefits in order to get a correct measure of total Income Security benefits. We take

observed family types by age and sex in the 1990 General Social Survey and develop a set of probabilities.ⁿ For each family in our simulation, we average the benefits over each of the seven family types using the 1990 probabilities as weights.

Our goal is to have simulated benefits for each someone of each age in each year. To form the two types of simulated benefits we use for the analysis, we use the described methodology in slightly different ways. For both simulations, we use the 1920 cohort earnings profile, the 1990 cross-sectional family types, and the 1985 cross-sectional capital income profile. For the partially simulated benefit, we use each cohort's observed retirement pattern to find the simulated benefits. In contrast, the fully simulated benefit applies the 1920 cohort's retirement pattern to all cohorts in our data. So, the key difference between the partially and fully simulated benefits is whether cohort-specific or fixed retirement rates were used in the construction of the benefits.

Description of Variation

The result of this simulation exercise is a benefits measure that varies across birth cohorts only. In the regression analysis below, we will control for both age and year effects, allowing us to identify the impact of IS programs on well-being solely from variation across birth cohorts. In this section, we show that such variation is sizeable, and explain its legislative origins.

Figure 2 opens our discussion by showing the evolution of IS income over time. The first line in this figure shows the average benefits paid to elders under IS programs, in 2001 Euros per person. The line is derived from aggregate spending on the programs divided by the number of

ⁿ The seven family types we consider are married, male single never married, male widowed, male

Canadians age 60 and over.^o The rise has been substantial, from around 2,000 Euros in the early 60s up to 4,000 in the 1970s and then 7,500 by the 1990s. The increase from 2000 over 1965 is 263 percent. The second line shows the ratio of those average benefits to average worker income, with the scale on the right-hand side axis. Over the same 1965 to 2000 period, this ratio increased by 179 percent, indicating that benefit growth has far-outpaced earnings growth over the last generation. This gives an indication of the very large extent to which Canada's programs have been expanded.

Figure 3 then shows our partially and fully simulated benefits measures, graphed over time. These also grow steadily over time, in parallel fashion to actual benefits income, and move very closely together. While they follow the same trends, our simulated benefits appear lower than actual benefits all along the time period under consideration.

Figure 4 shows the evolution of these benefits measures over birth cohorts more clearly. We graph the simulated benefits levels for each cohort in the year they reach age 65. The first upward increment in the late 1960s corresponds to the extension of OAS benefits down to age 65. The great rise in the 1970s is caused by the phasing in of the CPP/QPP over 1970 to 1976, the introduction and expansion of the GIS in 1967, 1971, 1979, and 1980, and finally the SPA in 1976. Through the 1980s and 1990s, benefits grow much more slowly, increasing only with higher earnings through the YMPE. The partially simulated benefits line is more volatile than the fully simulated line, reflecting the additional variation introduced by the cohort-specific retirement rates.

separated/divorced, female single never married, female widowed, female separated/divorced.

^o The source is Human Resources Development (1999).

The final picture of the variation in Canada's recent history is in Figure 5. Here, we expand on the age 65 variation from the previous figure to show benefits at several ages. Benefits at age 60 are zero until the 1980s when early retirement is introduced in Quebec (in 1984) and the rest of Canada (in 1987). Benefits at ages 70, 80, and 90 are similar until the mid 70s, when those who are 70 begin to show CPP/QPP income. These same cohorts a decade later as 80 year olds begin showing CPP/QPP income in the 1980s, so they pull higher than the 90 year olds who do not receive any CPP/QPP.

Regression Analysis

Having created these various measures of benefits (actual, partially simulated, full simulated), we can then use them in regression analysis to assess how IS generosity affects our various measures of well-being. We will in each case run regressions of the mean outcome for an age-year cell (e.g. mean after-tax income) on each of our three benefits measures. In doing so, we will consider four alternative regression models.

The first is a model that controls for age and year in a linear fashion. That is, we estimate models of the form:

$$Outcome_{ay} = \beta_0 + \beta_1 Benefits_{ay} + \beta_2 X_{ay} + \beta_3 AGE_a + \beta_4 YEAR_y + e_{ay},$$

where $Outcome_{ay}$ is one of the well-being outcomes, $Benefits_{ay}$ is the measure of benefits, X_{ay} is a vector of cohort characteristics, Age_a are controls for age, $Year_y$ are controls for year, and e_{ay} is an error. The cohort characteristics include the average earnings, marital status, and education levels as measured before retirement when the cohort was age 50. Also included in X is a control

for the average value of the dependent variable among the non-elderly families for the given year to capture any sharp national deviations in outcomes that are reflected in elderly and non-elderly alike.

Using this framework, we estimate four models. The first uses linear terms for age and year, and no cohort characteristics. The second adds the cohort and non-elderly controls to the linear specification. The third regression model replaces the linear age and year controls with explicit age and year dummies. In this model, we fully control for any variation in benefits by age or over time, so that all that identifies the model is variation across cohorts. Finally, we augment the age/year dummy regressions with controls for cohort characteristics. This is our richest model, imposing as an identifying assumption only that there are no differences across birth cohorts, conditional on observable characteristics, associated with both program generosity and outcomes. All of our regressions are weighted by the cell size from the micro-data to replicate the full survey outcomes.

Section IV: Evidence

Time Series Evidence

We begin our presentation of the effects of IS programs on well-being by examining the evolution of our various well-being indicators over time. We use the equivalence-scale adjusted measures, so the data represent the amount per effective person in 2001 Euros. In each case, we show two graphs, each with two lines. The two lines in each graph represent the well-being

measures for the elderly and the non-elderly; we include the latter group to show overall trends that are unlikely related to IS program generosity. The first graph shows these two series rescaled to fit on the same graph; the second graph shows the series in index form, so that they start at the same point.

Figure 6 shows the graph for actual benefits received, as reported in our microdata sources. Similar to the aggregate and simulated benefits shows in earlier figures, there is a strong increase from 1971 through 2000, with a flattening out over the 1990s.

Figure 7 shows the results for average after-tax family income. Average incomes rise rapidly from 1970 through 1980, and very slowly thereafter, for both the elderly and the non-elderly. The incomes of the elderly households are higher than for working families after 1980, with gaps of around 2,000 Euros per year. So this figure does not appear to offer much support for an effect of IS on well-being, in that there are no large rises in well-being of the elderly relative to the non-elderly during this period of rapid IS expenditure growth. The increase in 1998 to 2000 in incomes reflects the change from the Survey of Consumer Finances over to the Survey of Labour and Income Dynamics. The difference is not in reported family incomes, but instead in the average size of census families across the two data sources – families in the new survey are smaller on average than in the older survey, leaving the income to be spread less thinly.

Figures 8 and 9 show the evolution of poverty over this time period, with the relative and absolute income poverty measures respectively. By both measures, there is a steep reduction in elderly poverty relative to the non-elderly over the 1970s and 1980s, which is consistent with IS program growth. For relative poverty, however, this reduction does not begin until the late

1970s, after the major CPP/QPP expansion. In the indexed version, we can see clearly the drastic drop in elderly poverty, from 100 in 1971 down to under 40 by the early 1980s, for a drop of 60 percent. For absolute poverty, the reduction is very steep from 1971 onwards for the elderly, but it is also steep for the non-elderly as well. Once again, the major improvement for the elderly relative to the non-elderly is in the late 1970s and early 1980s. The timing of these movements corresponds to the expansions of the GIS and also the slow of extension of CPP benefits into an ever-higher proportion of the elderly cohorts included in the elderly sample.

Figure 10 shows the evolution of income inequality over this time period. After moving together, the series for the elderly and non-elderly diverge in the early 1980s, with the inequality of the elderly falling much more rapidly in the 1980s and early 1990s than for the non-elderly. Through the 1990s, inequality among working families increased slightly while it decreased among elderly families.

We next turn to our consumption-based measures of well-being. Here the time series are less complete due to more infrequent surveys. The volatility reflects in part the differences in survey coverage and sample size across different waves of the consumption surveys. Nonetheless, it is clear from Figure 11 that the time series for the elderly and non-elderly consumption largely move together, suggesting relatively little effect of IS programs.

Figure 12 and 13 examine relative and absolute consumption poverty measures. In both cases, the elderly and non-elderly move closely together until the late 1980s and early 1990s, at which point there is a larger fall in elderly consumption poverty than non-elderly consumption poverty. Once again, given the more or less steady rise in IS benefits over this period, it is hard to attribute this to IS effects.

Figure 14 measures consumption inequality. Interestingly, there is a sharp fall in consumption inequality during the 1970s, when the CPP/QPP are being phased in and the GIS expanded. This suggests that these programs may have bolstered the consumption of the bottom in a way which lowered overall inequality.

Finally, figures 15 and 16 show the results for our happiness measures. There is clearly variation across different samples, reflecting either true differences in happiness across time or some kind of differences in the sampling methodology. The measures for the elderly and non-elderly move very closely, suggesting little effect of the IS expansions on elderly happiness.

Regression Results

We next move from time series graphical analysis to regression analysis, which allows us to exploit the variation in benefits across birth cohorts documented in Figure 5. As is clear from that figure, a major source of variation is the evolution of benefits at age 60 over time, relative to other ages. This allows us to, in principle, include both age and year dummies in our model, controlling for general differences in well-being by age and over time. At the same time, we have also found that our models using linear age and year measures do not work well when we include 60-64 year olds in the sample. The reason for this is illustrated in Figure 17, which shows our measure of simulated benefits graphed against income, by age/year cell. It is clear that the set of points to the left, for 60-64 year olds, is described by a different process than is the set of points to the right, for 65 and over. Thus, imposing a linear age and year trend on this combined sample may be doing large violence to the data. However, once we include both age and year trends, the data line up better. Figure 18 shows a scatter-plot of the residuals from a

regression of income on a full set of age and year dummies plotted against benefits. The distinct clouds visible in Figure 17 are now removed in the data, allowing us to reasonably fit a line through the variation that remains. In the regressions below, we will show three sets of results: using the linear and age dummy models for age 65 and over, and using the age dummy model only for age 60 and over. We will rely on the last of these as our most reliable model.

The results of our analysis are presented in Tables 1-3, for the linear specification for age 65, the dummies specification for age 65, and the dummies specification for age 60, respectively. Each table has three sets of two columns: the first excludes controls for cohort characteristics (and young worker characteristics for the linear model), while the second includes them. While the controls make some difference in the linear models, there is little effect in the models that include age and year dummies, so we focus our discussion on the results including controls. The three sets of columns correspond to the use of actual benefits, partially simulated benefits, and fully simulated benefits.

The first panel of each table focuses on actual benefits received. In the second and third panels, using simulated benefits, we first run regressions of IS income on simulated benefits. This forms an implicit “first stage” for translating our simulated benefits calculation into an effect on actual benefits received. By comparing the remaining rows to this row, we can form an implicit instrumental variables calculation of the impact of each dollar of Income Security benefits on well-being. The magnitude of this first stage coefficient varies substantially across the tables, but for the age 60 dummies regression in Table 3 that we prefer the magnitude is very close to 1.

The second row shows the effect on average family income. There is in general an effect

on total family income which mirrors the effect on Income Security benefits, suggesting little “crowding out” of other income by these benefits increases. Correspondingly, in the first and third tables, there are sizeable reductions in poverty rates associated with benefits increases. In the third table, which we prefer, this reduction occurs only for relative and not absolute poverty measures. This is consistent with the time-series graphs in figures 8 and 9, in which only the relative measure differed from the trend in the non-elderly poverty rate. There is mixed evidence on the impacts of benefits on income inequality, but for the specification we prefer higher benefits raise income inequality among the elderly.

The next panel considers the impact on consumption, so that we can assess whether IS payments are crowding out other forms of consumption smoothing. In general, this appears not to be the case: the impact on consumption is very similar to the impact on both IS income and total income. This is particularly true in Table 3. So this finding suggests that each dollar of IS income is translated directly into consumption, with coefficients around 0.6.

Once again, focusing on the third column, we see effects on relative consumption poverty, but not on absolute consumption poverty. And we also see a rise in consumption inequality among the elderly (although it is not significant), mirroring the rise in income inequality.

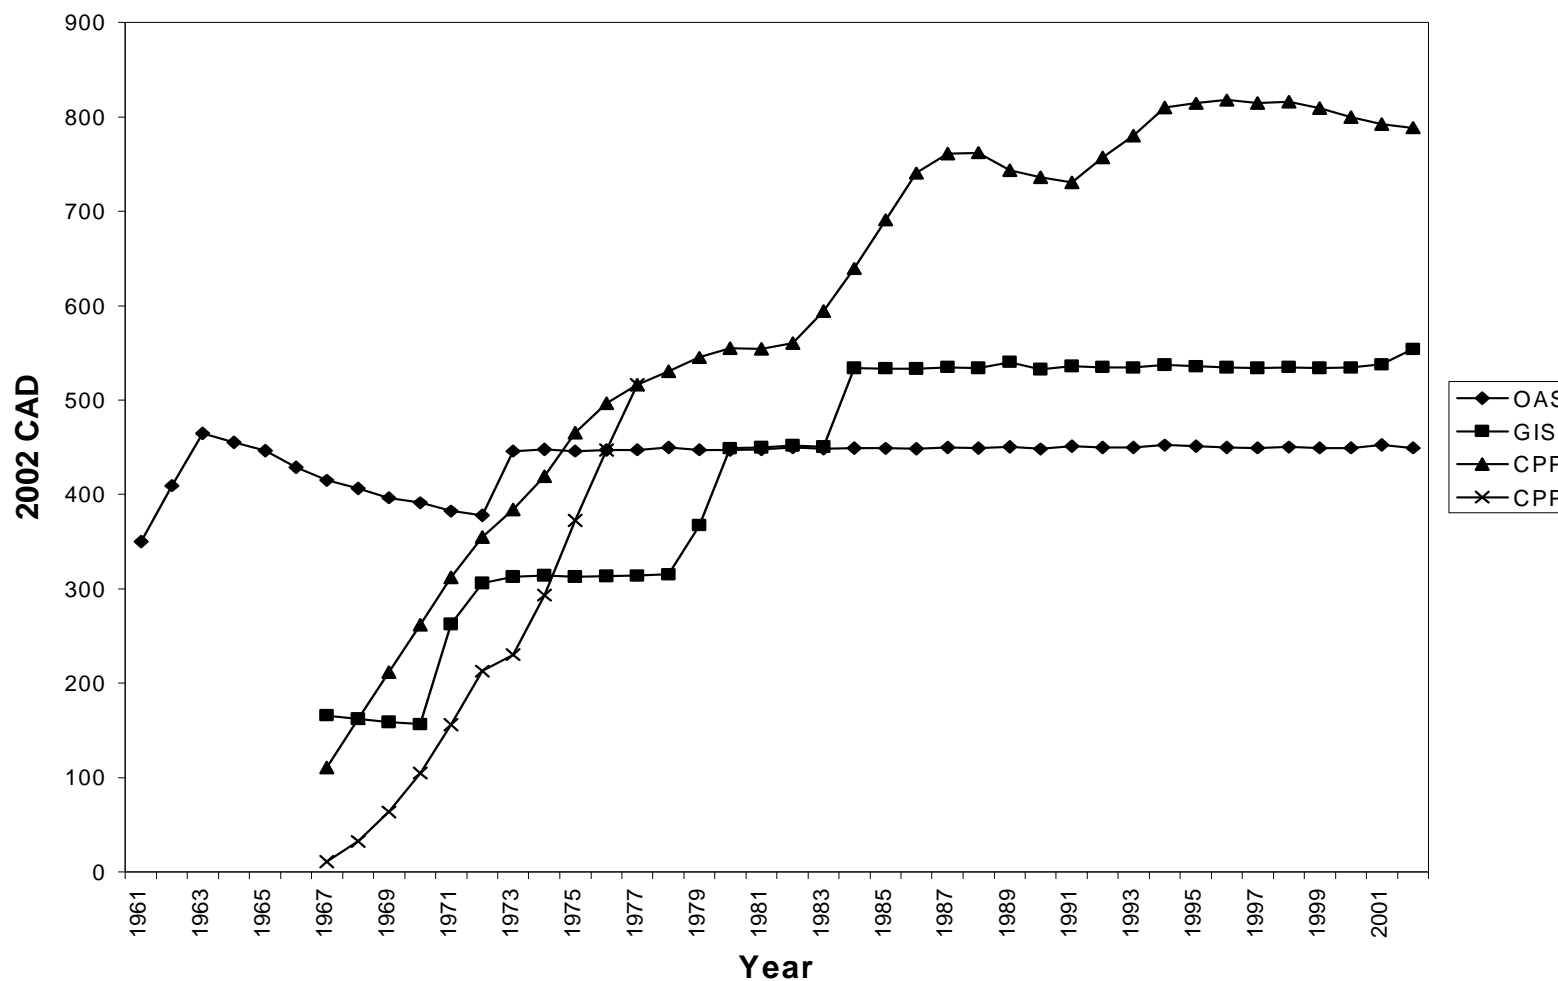
The final panel considers our measures of happiness. Unfortunately, we find little consistent pattern of effects of Social Security benefits on self-reported happiness. This may be due to the lack of variation in the data we have. The happiness data is much more limited than the consumption and income data, so this may not be surprising.

Conclusion

References

- Baker, Michael, "The Retirement Behaviour of Married Couples: Evidence From the Spouse's Allowance", *Journal of Human Resources*, 37, Winter 2002, 1-34.
- Baker, Michael, Jonathan Gruber, and Kevin Milligan (2003), "Simulating the Response to Reform of Canada's Income Security System," NBER Working Paper No. 9455.
- Baker, Michael, Jonathan Gruber, and Kevin Milligan (2004), "Income Security Programs and Retirement in Canada," in Jonathan Gruber and David Wise (eds.) *Social Security Programs and Retirement Around the World: Micro Estimation*. Chicago: University of Chicago Press.
- Human Resources Development (1999) HRDC Redbook. Statistics on income security.
- Office of the Superintendent of Financial Institutions (2002a), *Actuarial Report on the Old Age Security Program as at December 31, 2000*. Ottawa.
- Office of the Superintendent of Financial Institutions (2002b), *Canada Pension Plan: 19th Actuarial Report as at December 31, 2000, Supplement*. Ottawa.

Figure 1: Real Maximum CPP, OAS and GIS Benefits 1961-2002



Notes: GIS benefits are for a single individual. The CPP phase-in calculation is for December of the indicated year.

Figure 2: Average Income Security Benefits

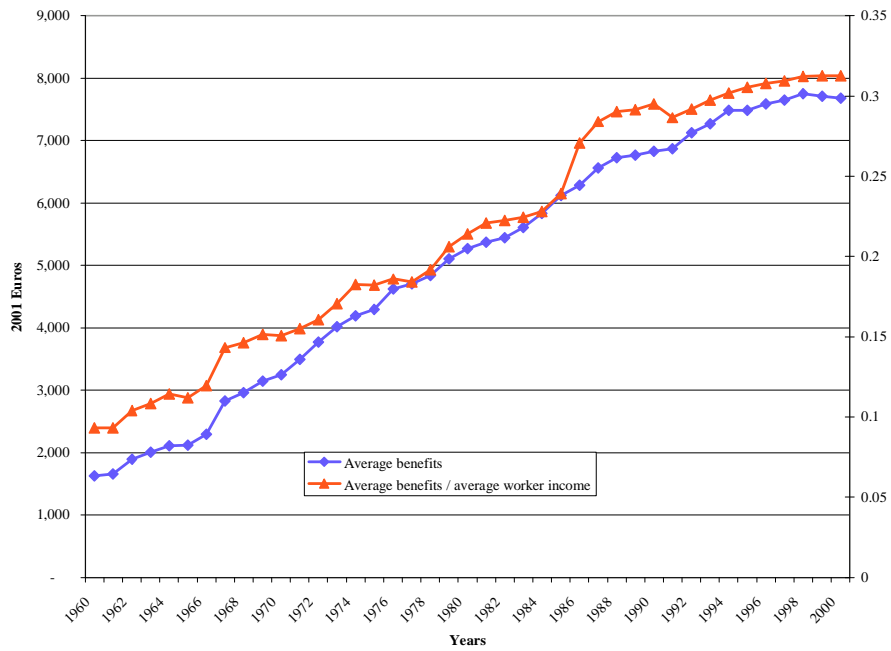


Figure 3: Average Simulated Benefits

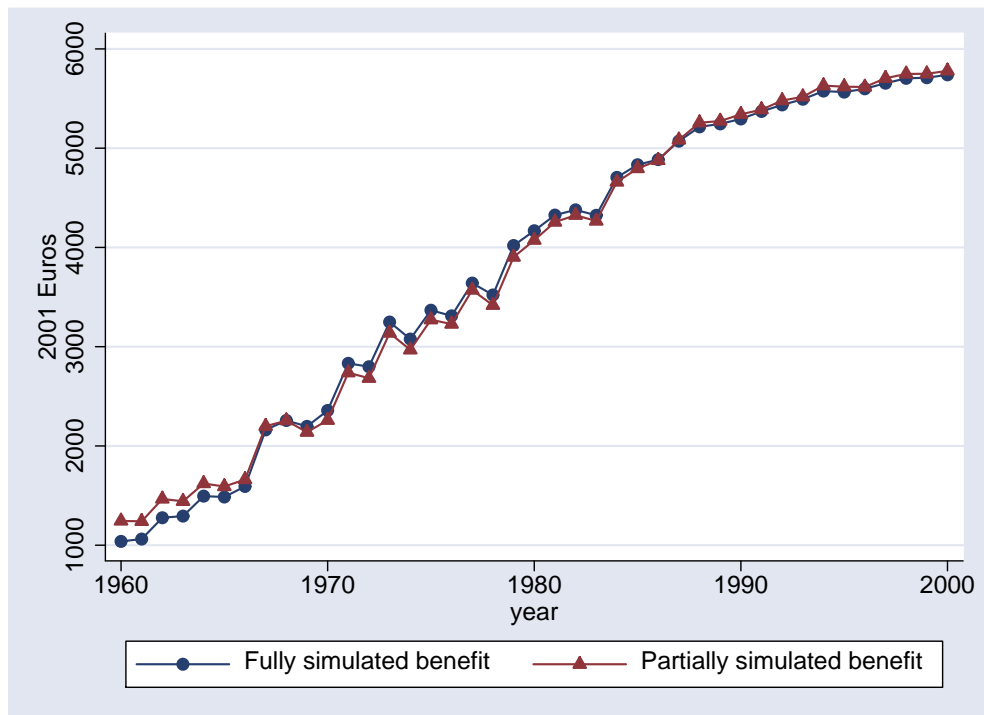


Figure 4: Benefits at age 65 in different years

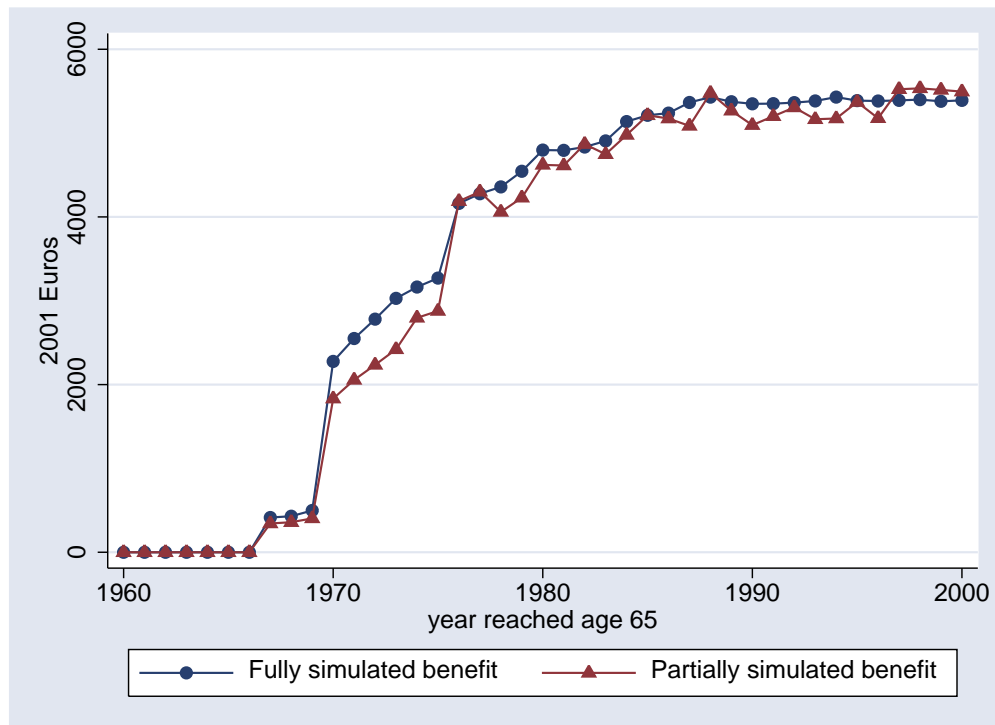
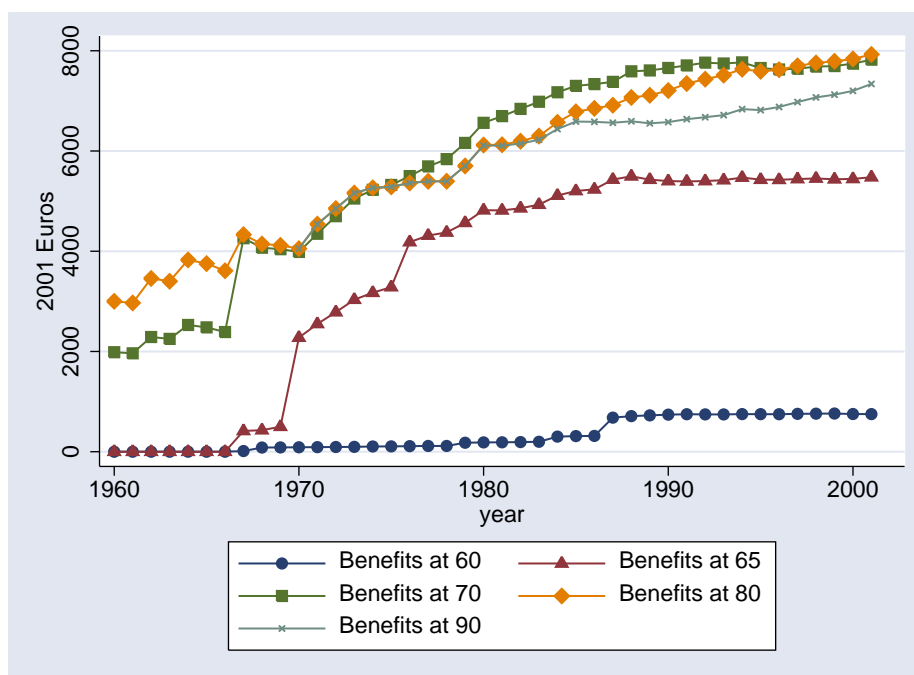


Figure 5: Benefits at different ages through time



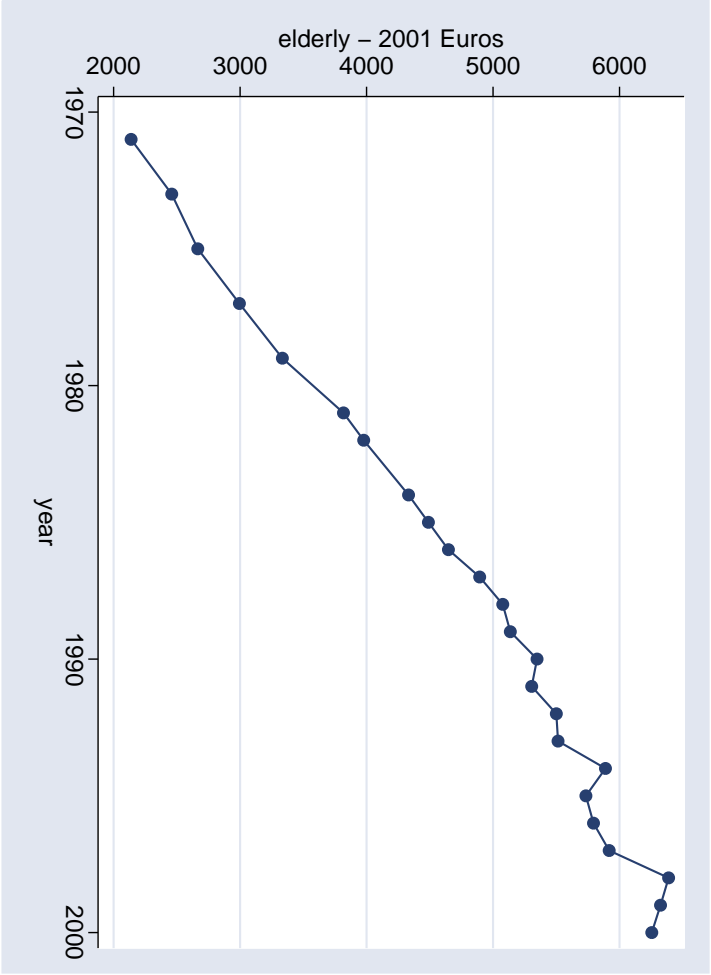
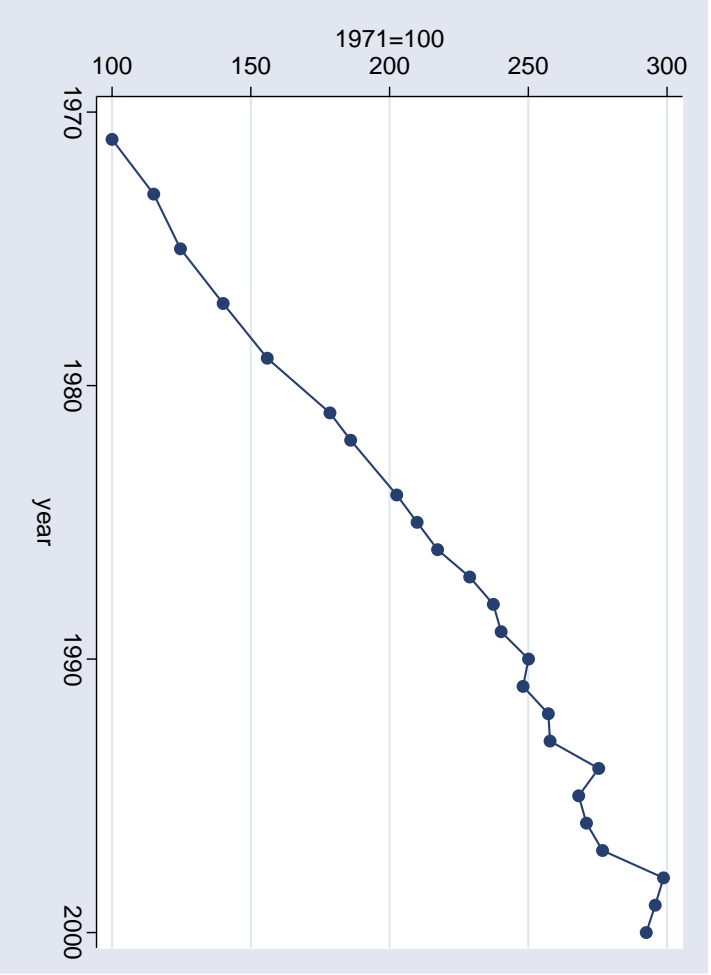


Figure 6: Benefits

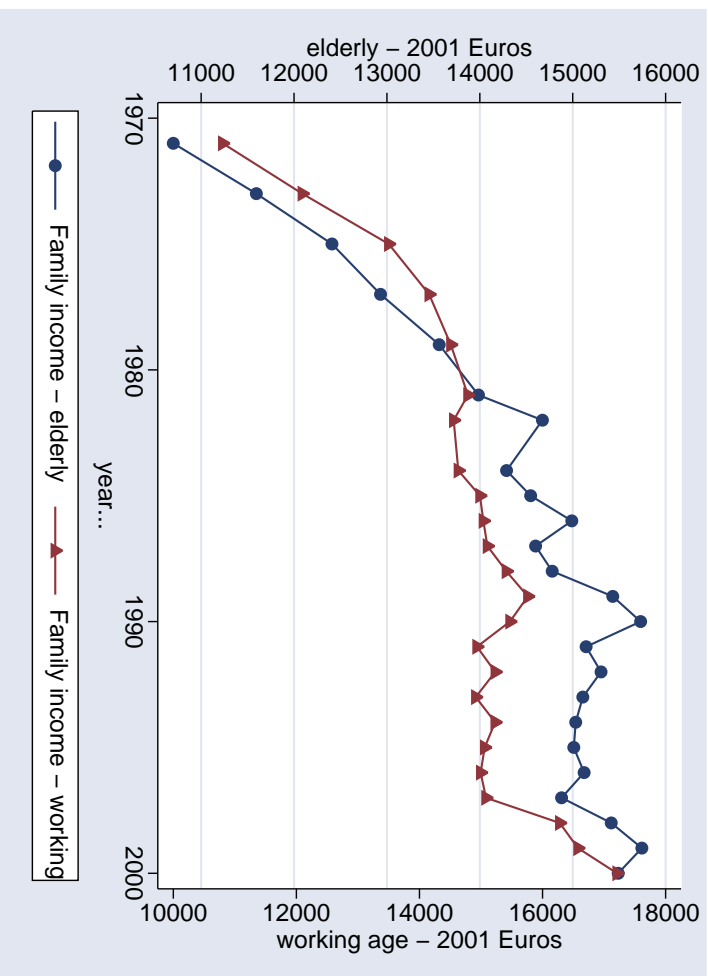


Figure 7: After-tax income

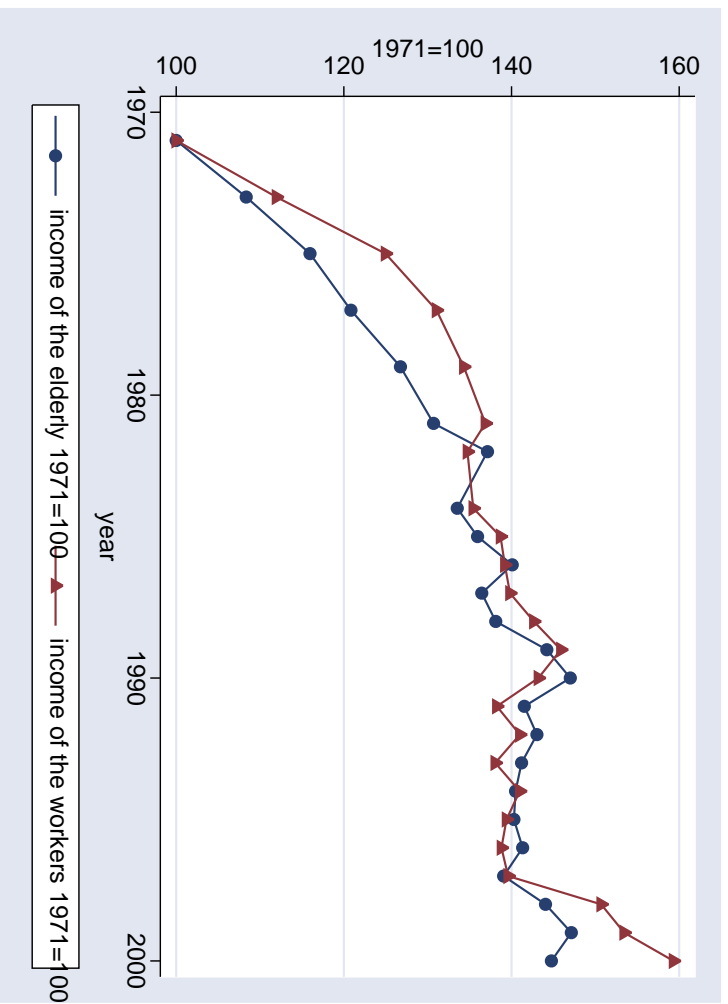


Figure 8: Relative income poverty

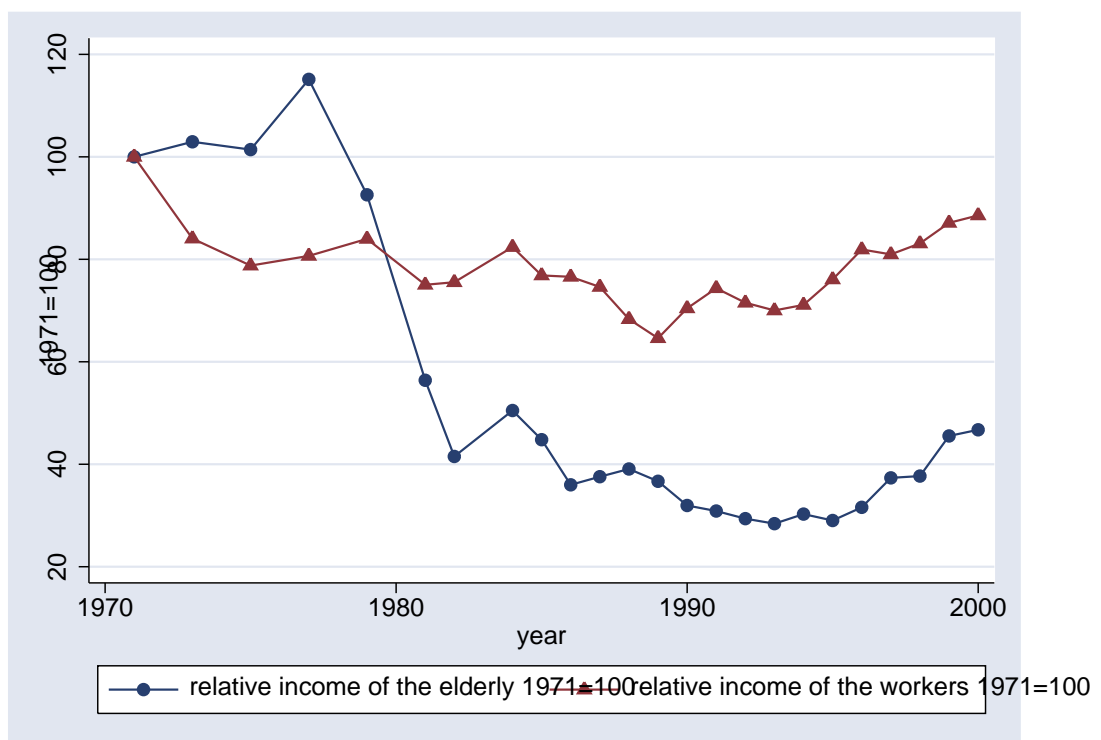
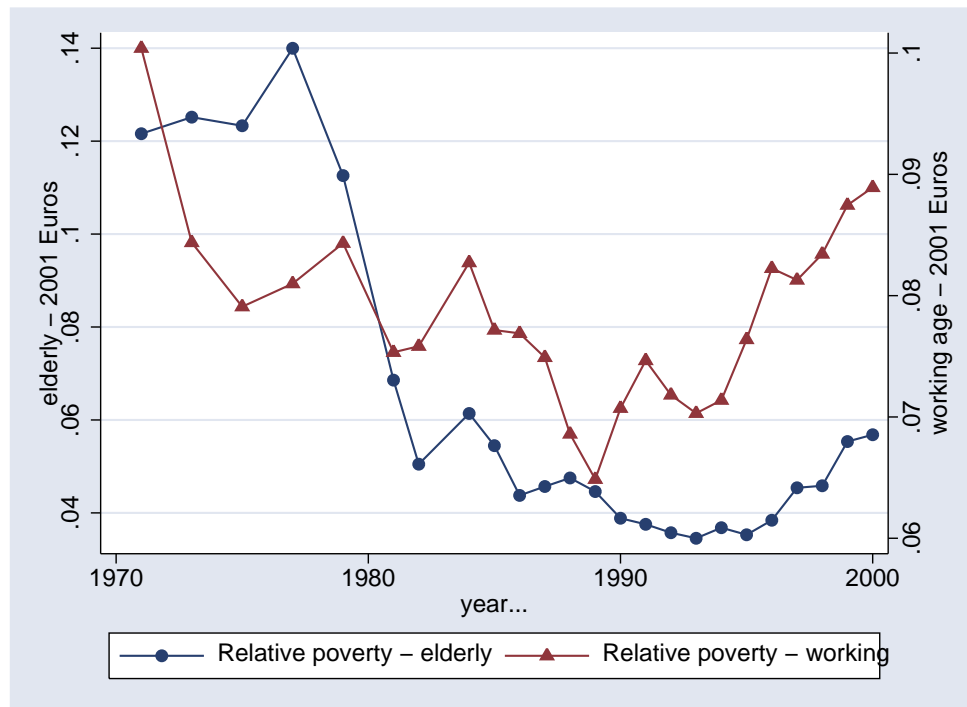


Figure 9: Absolute income poverty

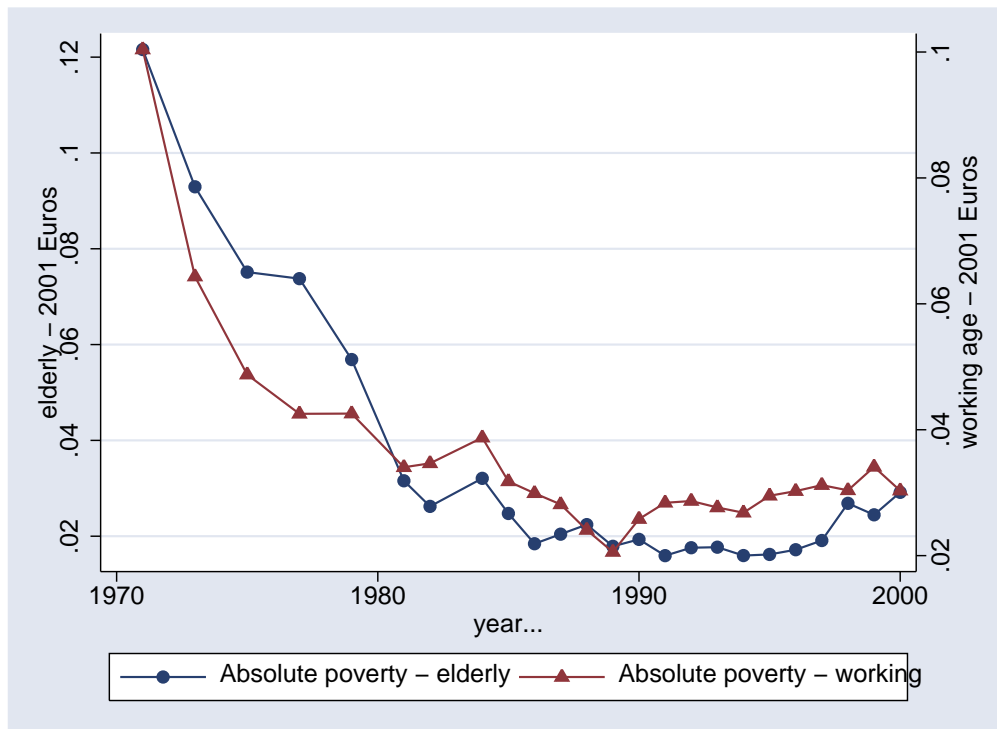


Figure 10: Income inequality 90-10

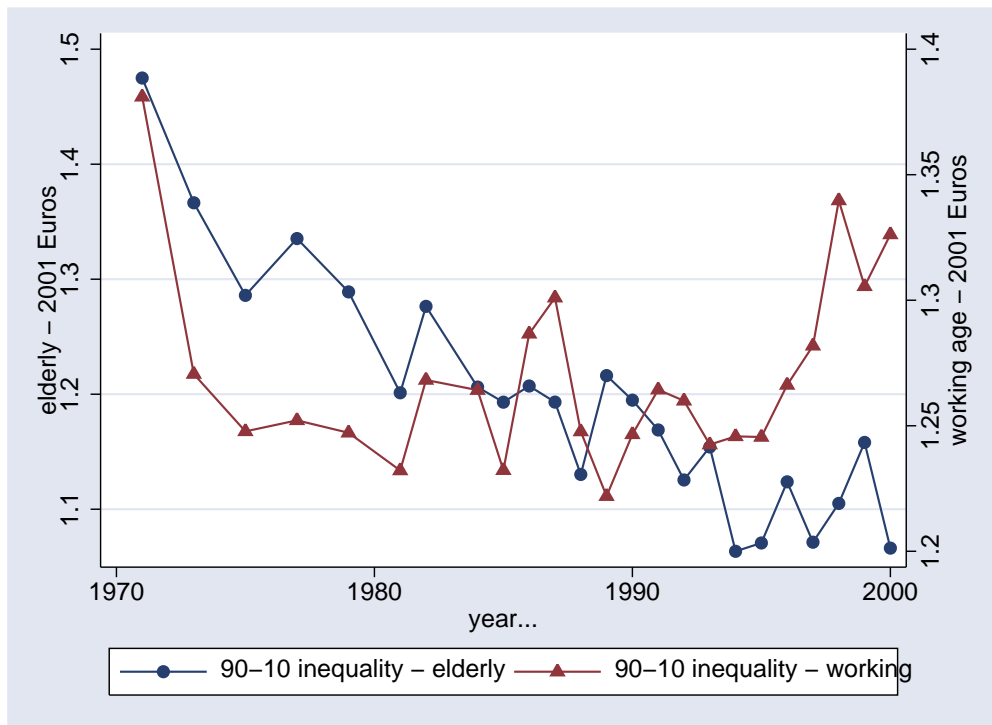


Figure 11: Total consumption

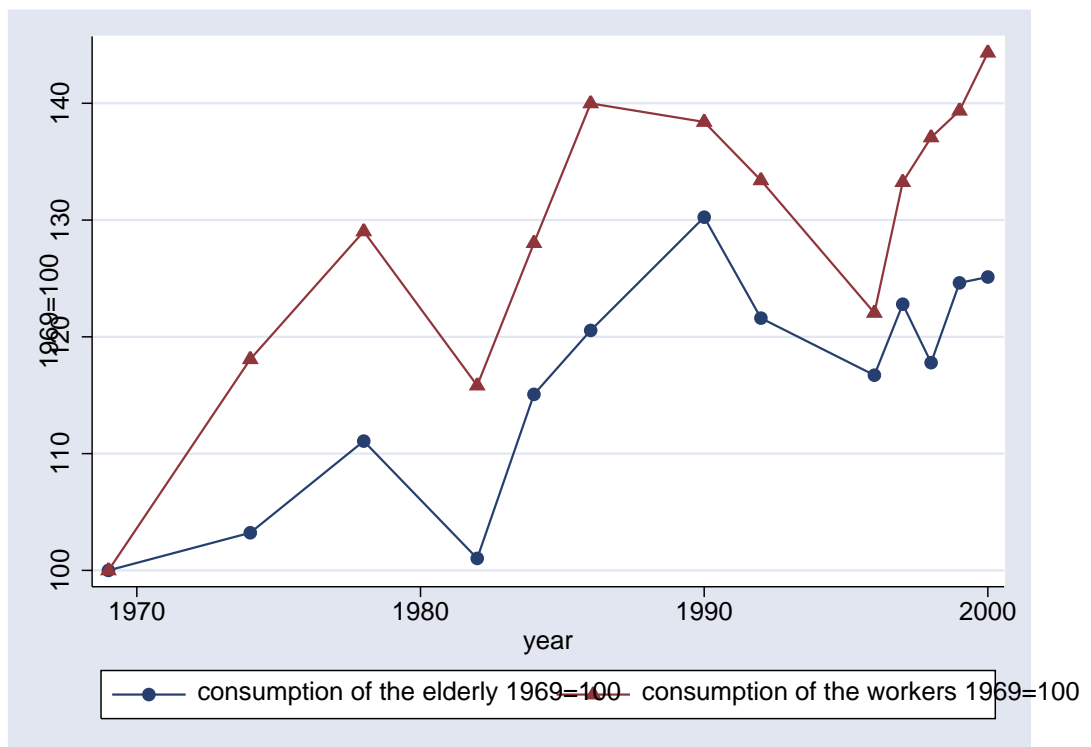
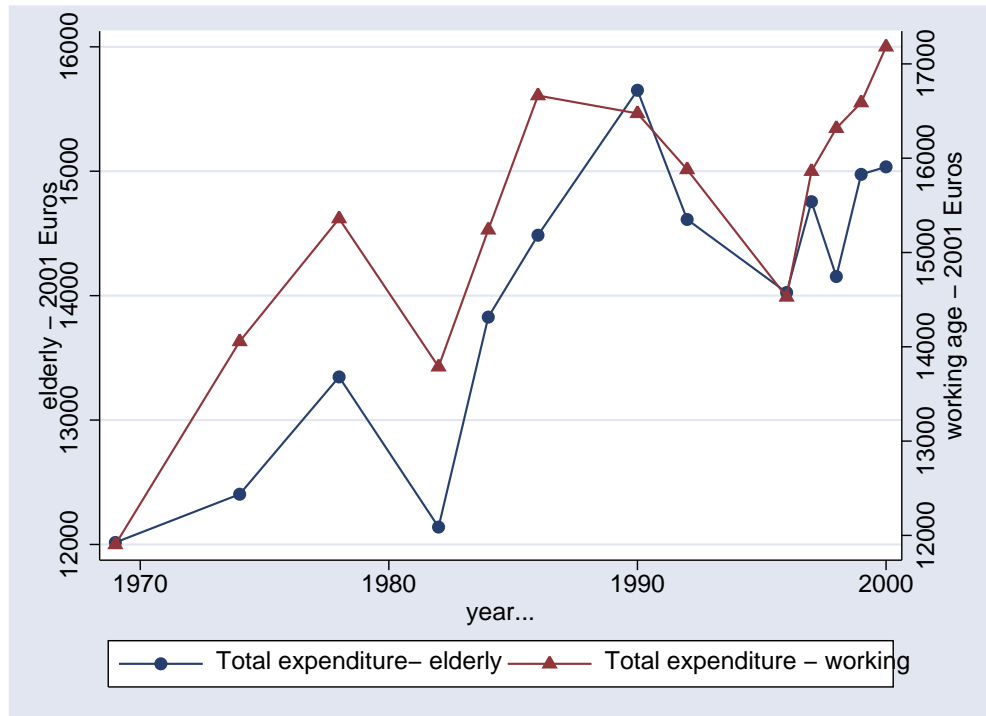


Figure 12: Relative consumption poverty

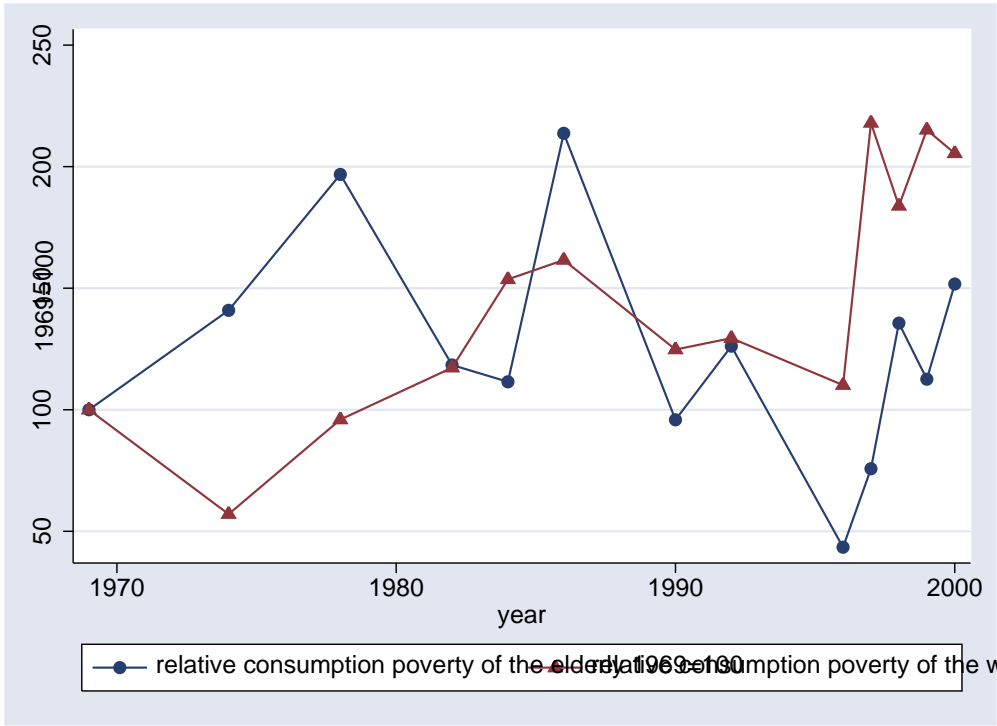
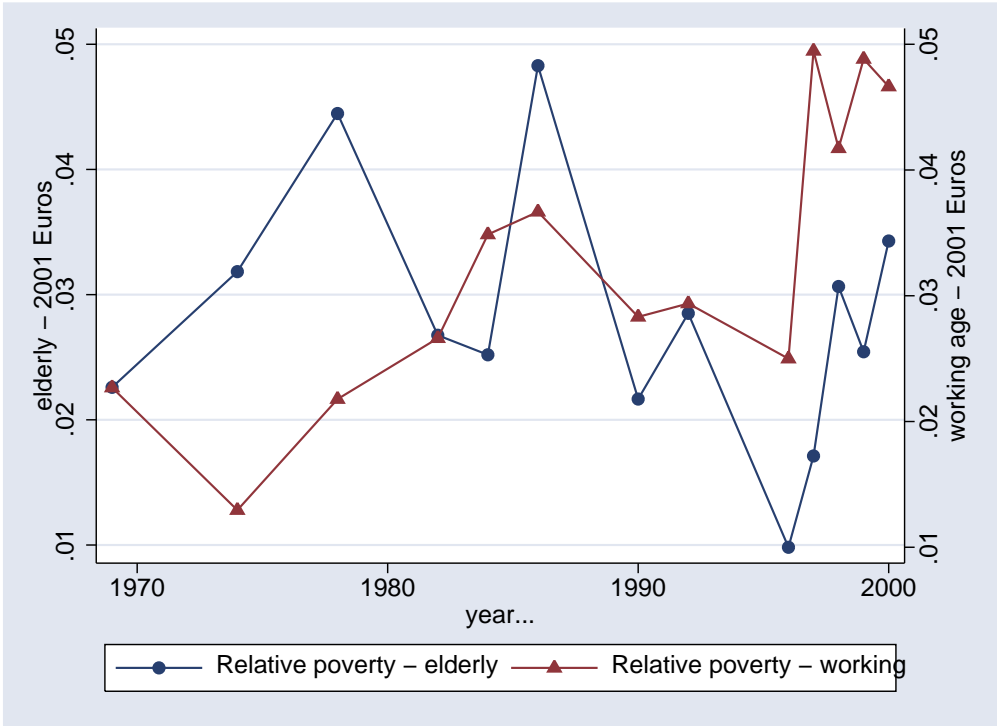


Figure 13: Absolute consumption poverty

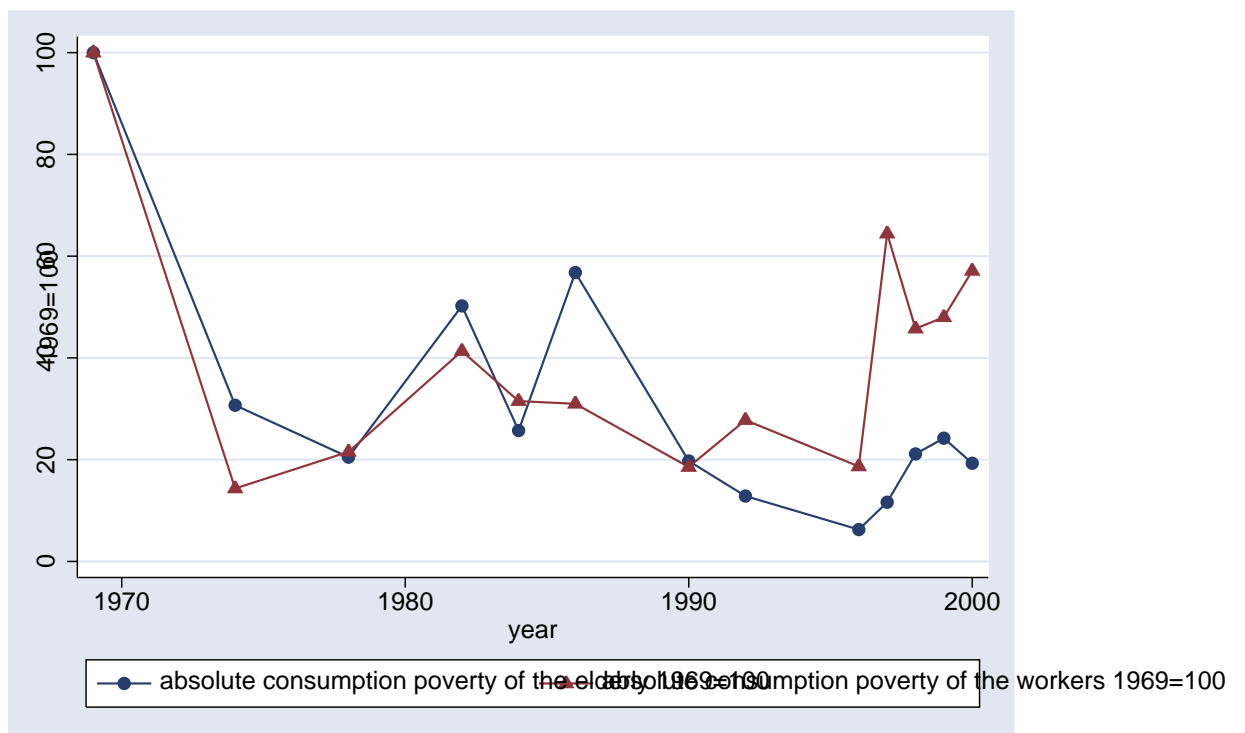
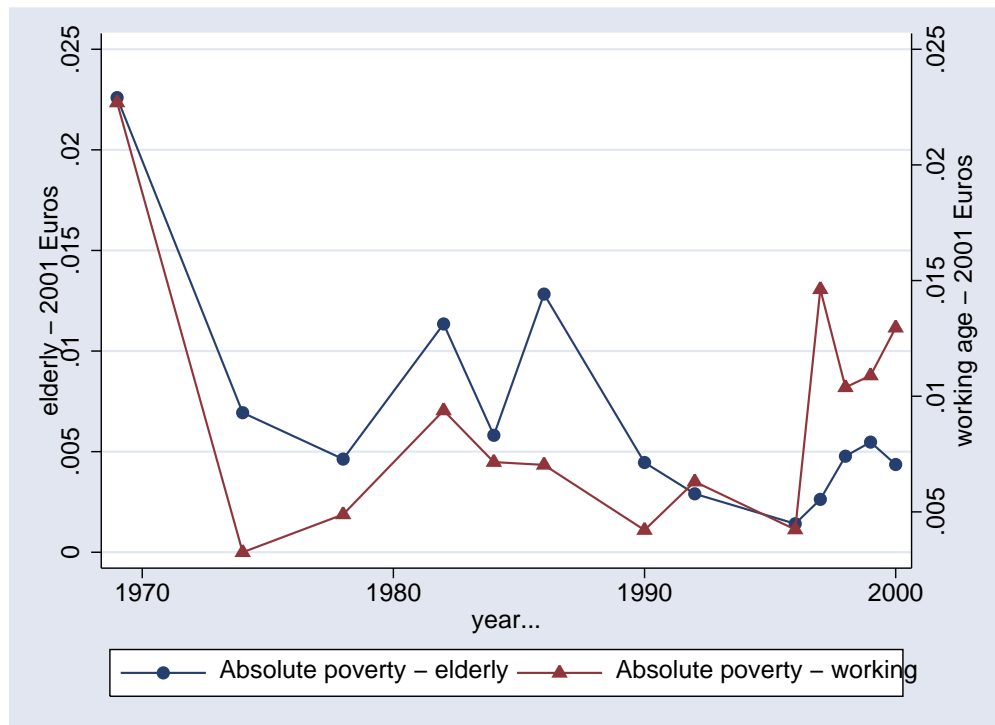


Figure 14: Consumption inequality

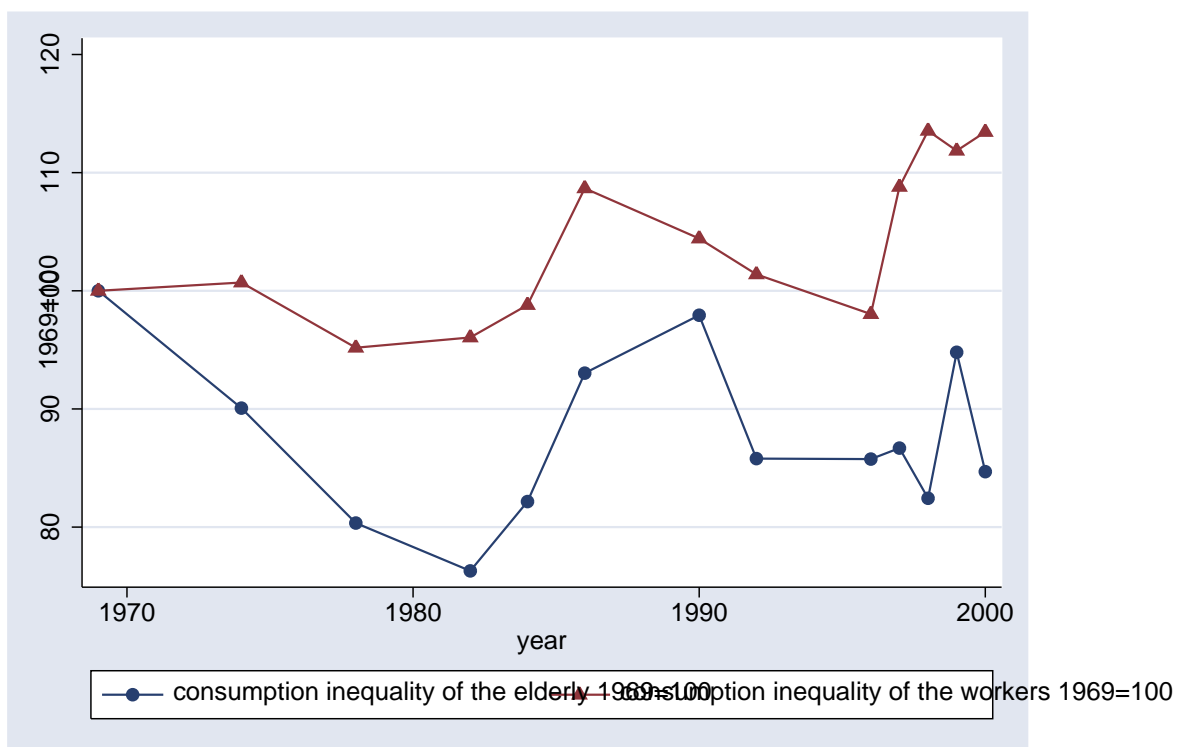
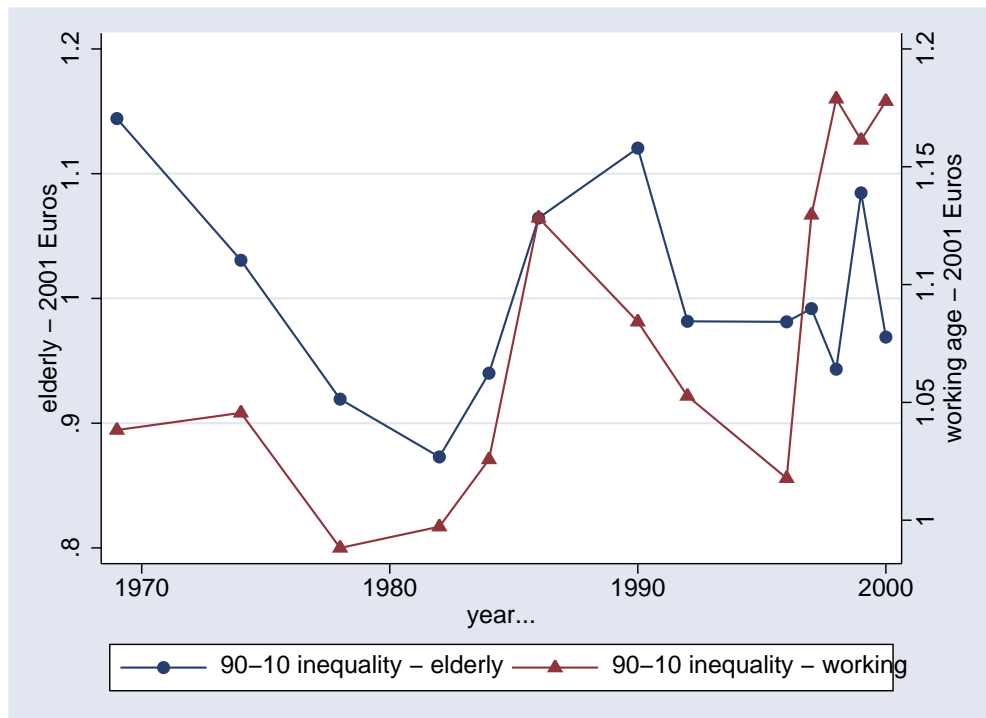


Figure 15: Very Happy

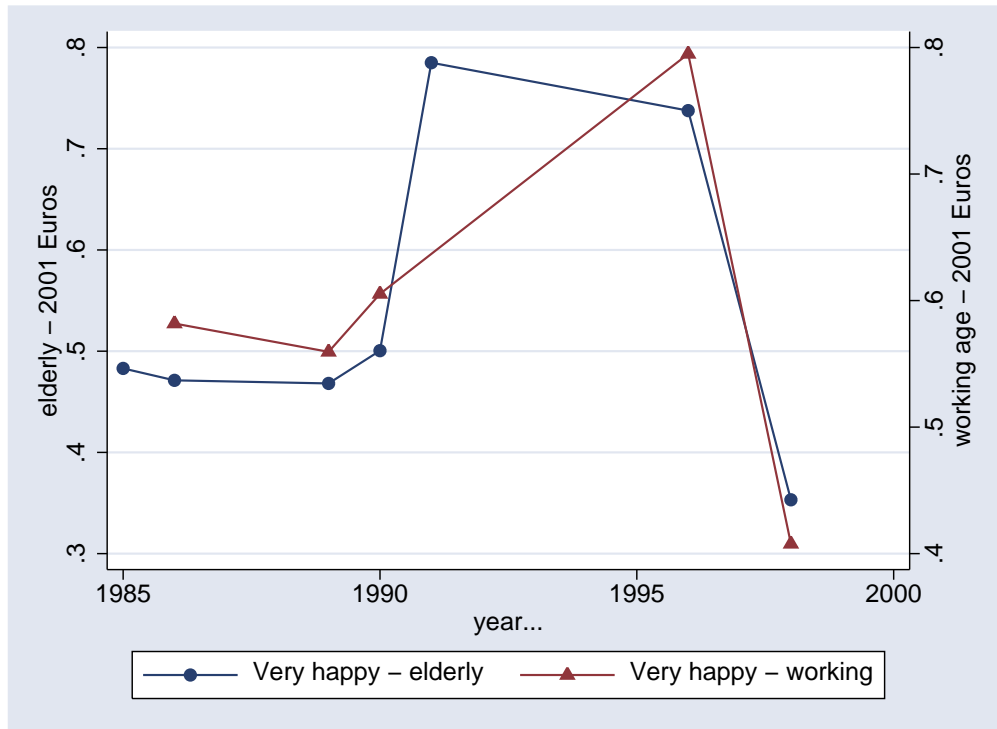


Figure 16: Unhappy

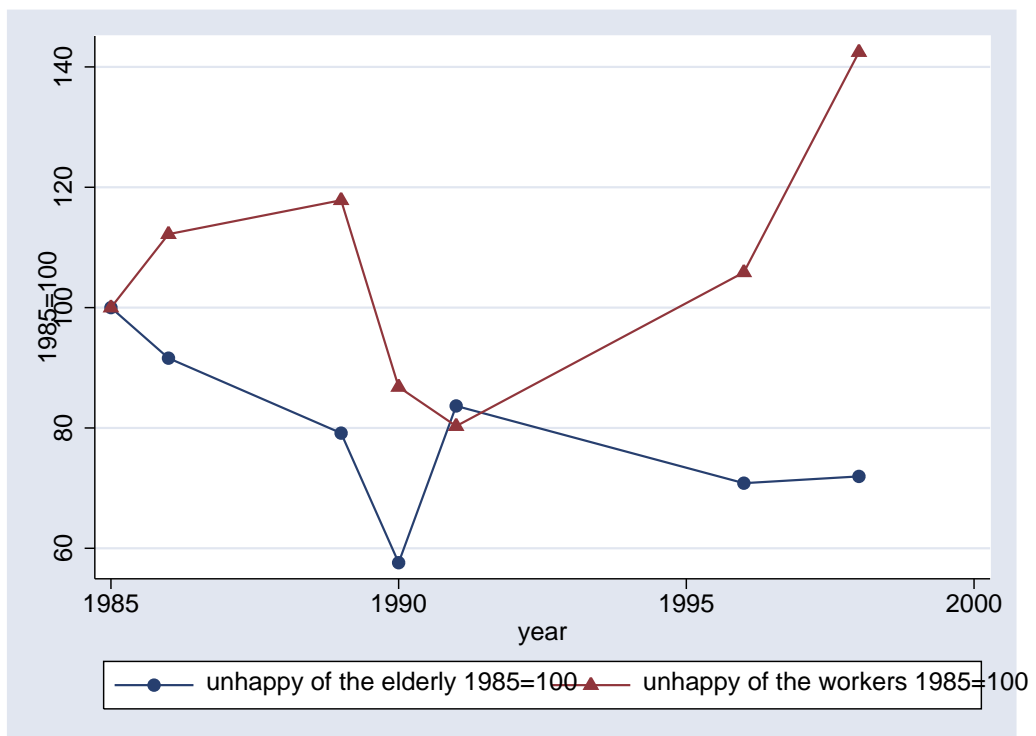
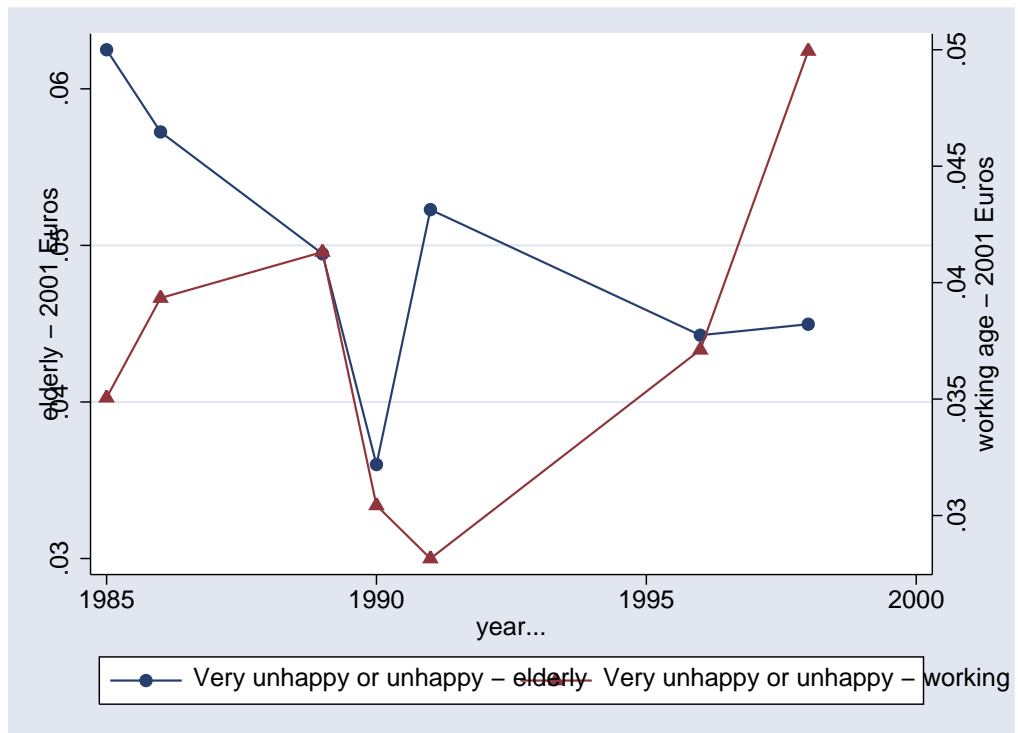


Figure 17: Benefits vs. Income Scatterplot

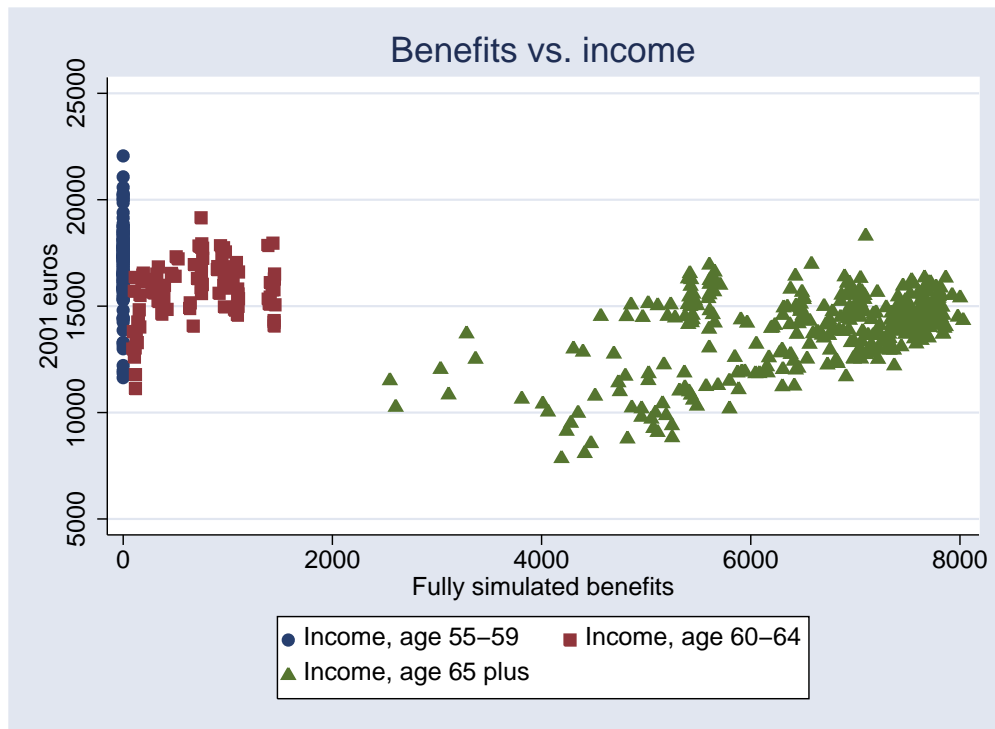


Figure 18: Residuals vs. Benefits

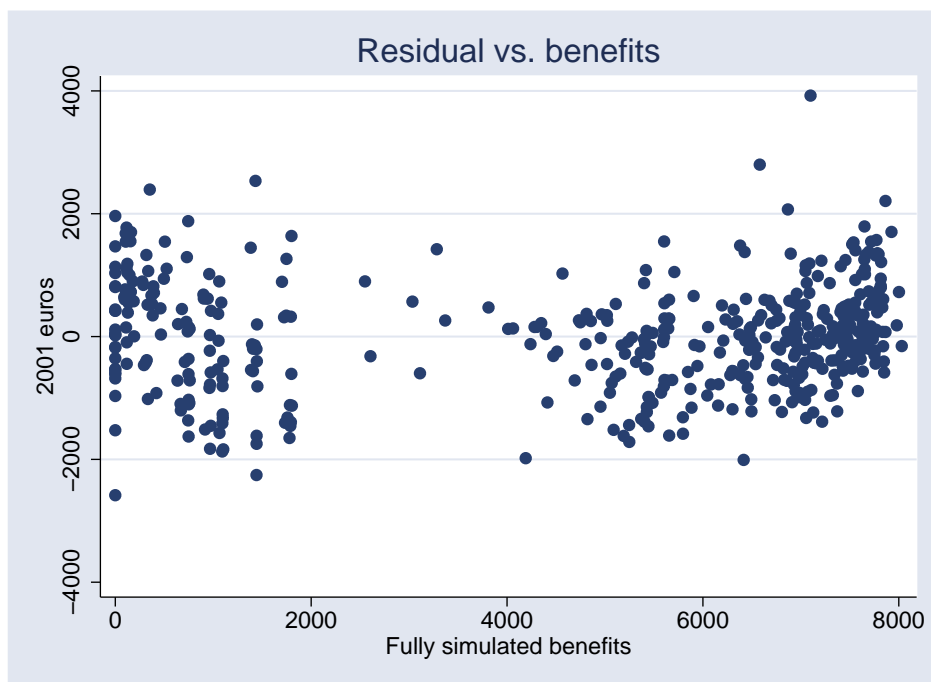


Table 1: Linear Age Specification, ages 65+

	Mean of Dep.	number	Actual		Partially simulated		Fully Simulated	
	Variable	of obs.	Benefits		Benefits		Benefits	
cohort controls			no	yes	no	yes	no	yes
working age controls			no	yes	no	yes	no	yes
Income								
SS income	6750	340	--	--	0.769 (0.053)	0.785 (0.058)	0.836 (0.056)	0.864 (0.062)
Income	14031	340	0.290 (0.088)	0.203 (0.083)	0.443 (0.084)	0.247 (0.092)	0.524 (0.090)	0.310 (0.099)
Relative income poverty	0.037	340	-0.028 (0.002)	-0.026 (0.002)	-0.024 (0.002)	-0.021 (0.002)	-0.025 (0.002)	-0.022 (0.003)
Absolute income poverty	0.019	340	-0.017 (0.002)	-0.015 (0.002)	-0.018 (0.002)	-0.013 (0.002)	-0.019 (0.002)	-0.013 (0.002)
90-10 income inequality	1.037	340	-0.089 (0.011)	-0.091 (0.012)	-0.070 (0.015)	-0.078 (0.016)	-0.064 (0.016)	-0.073 (0.017)
Consumption								
Consumption	13108	157	-0.269 (0.182)	-0.357 (0.155)	-0.217 (0.151)	-0.470 (0.157)	-0.219 (0.154)	-0.449 (0.162)
Relative consumption poverty	0.031	157	-0.001 (0.003)	0.000 (0.003)	-0.002 (0.003)	-0.004 (0.004)	-0.003 (0.003)	-0.005 (0.004)
Absolute consumption poverty	0.007	157	-0.001 (0.001)	0.000 (0.001)	-0.002 (0.002)	-0.002 (0.002)	-0.003 (0.001)	-0.003 (0.002)
90-10 consumption inequality	0.918	157	-0.047 (0.025)	-0.043 (0.025)	-0.075 (0.025)	-0.073 (0.027)	-0.071 (0.025)	-0.066 (0.028)
Happiness								
Very happy	0.618	21	-0.100 (0.077)	0.055 (0.053)	-0.007 (0.051)	0.036 (0.017)	-0.004 (0.054)	0.036 (0.017)
Unhappy or very unhappy	0.057	21	-0.009 (0.007)	0.006 (0.029)	-0.004 (0.005)	0.006 (0.014)	-0.004 (0.005)	0.005 (0.014)

cohort controls: For each cohort at age 50, take mean earnings, marital status, and education levels.

Working age controls: The mean of the dependent variable among the working age population in that year is added as a control variable.

Table 2: Age Dummy Specification, ages 65+

	Mean of Dep.	number	Actual		Partially simulated		Fully Simulated	
	Variable	of obs.	Benefits		Benefits		Benefits	
cohort controls			no	yes	no	yes	no	yes
working age controls			no	no	no	no	no	no
Income								
SS income	6750	340	--	--	0.477 (0.061)	0.468 (0.065)	0.440 (0.067)	0.449 (0.077)
Income	14031	340	0.593 (0.219)	0.565 (0.219)	0.746 (0.247)	0.733 (0.251)	1.545 (0.219)	1.522 (0.230)
Relative income poverty	0.037	340	-0.021 (0.004)	-0.022 (0.005)	-0.010 (0.005)	-0.012 (0.005)	-0.006 (0.005)	-0.007 (0.006)
Absolute income poverty	0.019	340	-0.011 (0.004)	-0.011 (0.005)	-0.004 (0.005)	-0.006 (0.006)	0.003 (0.004)	0.003 (0.005)
90-10 income inequality	1.037	340	0.003 (0.037)	-0.010 (0.036)	0.054 (0.050)	0.033 (0.051)	0.191 (0.045)	0.185 (0.048)
Consumption								
Consumption	13108	157	0.067 (0.390)	0.054 (0.377)	1.108 (0.375)	1.181 (0.411)	1.227 (0.386)	1.440 (0.474)
Relative consumption poverty	0.031	157	-0.015 (0.013)	-0.016 (0.012)	-0.030 (0.012)	-0.033 (0.013)	-0.038 (0.014)	-0.041 (0.015)
Absolute consumption poverty	0.007	157	0.004 (0.003)	0.003 (0.003)	-0.010 (0.004)	-0.013 (0.004)	-0.012 (0.004)	-0.015 (0.005)
90-10 consumption inequality	0.918	157	-0.033 (0.080)	-0.041 (0.083)	0.046 (0.073)	0.030 (0.075)	0.120 (0.069)	0.116 (0.081)
Happiness								
Very happy	0.618	21	-0.064 (0.037)	-0.097 (0.098)	-0.059 (0.045)	-0.086 (0.119)	-0.070 (0.052)	-0.086 (0.131)
Unhappy or very unhappy	0.057	21	-0.005 (0.013)	0.015 (0.032)	-0.030 (0.017)	-0.017 (0.058)	-0.034 (0.020)	-0.009 (0.071)

cohort controls: For each cohort at age 50, take mean earnings, marital status, and education levels.

Working age controls: The mean of the dependent variable among the working age population in that year is added as a control variable.

Table 3: Age Dummy Specification, ages 60+

	Mean of Dep. Variable	number of obs.	Actual Benefits		Partially simulated Benefits		Fully Simulated Benefits	
cohort controls			no	yes	no	yes	no	yes
working age controls			no	no	no	no	no	no
Income								
SS income	5075	460	--	--	0.877 (0.047)	0.868 (0.048)	0.918 (0.035)	0.930 (0.038)
Income	14494	460	0.777 (0.116)	0.710 (0.122)	0.646 (0.149)	0.551 (0.153)	1.099 (0.157)	1.011 (0.167)
Relative income poverty	0.056	460	-0.025 (0.003)	-0.024 (0.003)	-0.015 (0.004)	-0.013 (0.004)	-0.023 (0.004)	-0.021 (0.004)
Absolute income poverty	0.030	460	-0.004 (0.002)	-0.004 (0.002)	0.000 (0.003)	0.000 (0.003)	0.000 (0.003)	0.000 (0.003)
90-10 income inequality	1.165	460	-0.003 (0.018)	-0.006 (0.018)	0.038 (0.023)	0.031 (0.023)	0.062 (0.022)	0.058 (0.022)
Consumption								
Consumption	13912	222	0.467 (0.198)	0.505 (0.205)	0.550 (0.176)	0.590 (0.177)	0.596 (0.171)	0.650 (0.182)
Relative consumption poverty	0.028	222	-0.009 (0.005)	-0.010 (0.005)	-0.004 (0.004)	-0.005 (0.004)	-0.007 (0.004)	-0.008 (0.004)
Absolute consumption poverty	0.007	222	0.000 (0.002)	0.000 (0.002)	-0.001 (0.003)	-0.001 (0.003)	-0.001 (0.003)	-0.002 (0.003)
90-10 consumption inequality	0.986	222	0.050 (0.042)	0.051 (0.046)	0.052 (0.038)	0.053 (0.039)	0.058 (0.036)	0.061 (0.038)
Happiness								
Very happy	0.610	28	-0.033 (0.044)	0.006 (0.073)	0.030 (0.038)	0.032 (0.045)	0.018 (0.056)	0.022 (0.068)
Unhappy or very unhappy	0.057	28	-0.003 (0.013)	0.005 (0.015)	-0.020 (0.009)	-0.015 (0.010)	-0.028 (0.012)	-0.017 (0.014)

cohort controls: For each cohort at age 50, take mean earnings, marital status, and education levels.

Working age controls: The mean of the dependent variable among the working age population in that year is added as a control variable.