

How Do Public Pensions Affect Retirement Incomes and Expenditures?: Evidence over Five Decades from Canada

Kevin Milligan
University of British Columbia
and
David A. Wise
NBER

Prepared for the 16th Annual Joint Meeting of the Retirement Research Consortium
August 7-8, 2014
Washington, DC

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA), funded as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the federal government, the University of British Columbia, the National Bureau of Economic Research, or the Center for Retirement Research at Boston College.

Abstract

We study the income and expenditures of Canadian elderly families over the five decades from 1960 to 2010. We first document the tremendous changes in income and expenditures using available microdata sources. Through simulations, we then display the impact of the public pension system over a fifty year period. Our analysis reveals three important findings. First, we document that the expansion of Canada's public pension system over the last 50 years has coincided with a large improvement in elderly living standards, measured by income or consumption. Second, we causally relate these changes using an instrumental variables strategy exploiting variation across ages and years in the Canadian system. We find strong evidence that public pensions have lifted income. For expenditures, the evidence is more mixed but there is strong evidence of improvements in reducing expenditure poverty. Third, taking our estimates on the effect of the pension system on income poverty, we perform counterfactual simulations by applying the public pension system of different decades to data from the 2000s. We find that the 2010 system reduces age 70 to 79 relative income poverty by 88 percent relative to the system in 1960. For relative expenditure poverty, the corresponding reduction is 56 percent.

Introduction

Around the world, the performance of public pensions is critical to the wellbeing of the elderly. Empirically assessing the importance of the contribution of public pensions, however, is a difficult task. The systems of most countries were put in place before the advent of detailed microdata, which inhibits analysis of their introduction. Even when reforms have been made in more recent years which might allow an analysis, it can be difficult to exploit these reforms because most people within a country are treated the same. This makes it difficult to form a comparison group to separate general time trends from the impact of the reform.

The case of Canada has several advantages for evaluating the overall impact of public pensions. First, the main earnings-related benefit was introduced in 1966, meaning that the roll-out of that system can be observed using fairly comparable microdata sources spanning 1969 to 2010. Second, a series of several smaller reforms taking place in particular years is available to provide even more policy variation to help examine the impact of pensions on income and expenditure.

In this paper we carry out an analysis of the impact of public pensions on retirement income and expenditure using the case of Canada. We first document the tremendous changes in income and expenditures over a four-decade period using available microdata sources. Through simulations, we then display the impact of the public pension system over the fifty-year period 1960 to 2010. We calibrate our simulations to capture the exogenous institutional variation across age and time. These simulated benefits are then used as instruments in a regression analysis of the impact of public pension benefits on incomes and expenditure. We close with a counterfactual policy

simulation of how poverty in the 2000s would have looked had Canada imposed the public pension system prevailing in each decade from 1960 to 2010.

We build here on previous work in Canada by Baker, Gruber, and Milligan (2009). We extend their approach to incorporate another decade of data, and also add new simulations that provide insight. More detailed work looking at elderly poverty rates in Canada can be found in Milligan (2008) and Schirle (2013). Two related papers look at incomes in the transition to retirement for Canada (Milligan 2010) and for the United States in Milligan (2012). A similar approach to the exploitation of cohort-based variation for the United States can be found in Engelhardt and Gruber (2006). A specific focus on consumption and income poverty can be found in Meyer and Sullivan (2010).

Our analysis reveals three important findings. First, we document that the expansion of Canada's public pension system over the last 50 years has coincided with a large improvement in elderly living standards, measured by income or consumption. Second, we causally relate these changes using an instrumental variables strategy exploiting variation across ages and years in the Canadian system. We find strong evidence that public pensions have lifted income. For expenditures, the evidence is more mixed but there is strong evidence of improvements in reducing expenditure poverty. Third, taking our estimates on the effect of the pension system on income poverty we perform counterfactual simulations by applying the public pension system of different decades to the data from the 2000s. We find that the 2010 system reduces relative income poverty by 88 percent relative to the system in 1960, for those aged 70 to 79, and relative expenditure poverty by 56 percent.

The paper proceeds as follows. We begin by describing the institutional features of the Canadian system. This is followed by an overview of the trends in retirement income and expenditures, and an exploration of our simulated benefits to highlight the extent of policy variation. We then present our regression estimates, followed by the counterfactual simulations. Finally, we offer some concluding thoughts.

Institutions

The Canadian retirement income has undergone tremendous change over the last 50 years, going from one small flat benefit for 70 year olds to a complex multi-tiered structure that incorporates flat benefits, income-tested benefits, and earnings-related benefits. In this section, we discuss the key elements of the development of the Canadian retirement income system over the 50 year period from 1960 to 2010, taking each component of the system in turn. For a more detailed overview of the Canadian retirement income system, see Milligan and Schirle (2013).

Old Age Security

In 1960 the only country-wide retirement income program was Old Age Security, which paid a flat \$55 per month benefit to seniors age 70 or more. This benefit was originally contributory, but has been funded from general revenues since 1973. The eligibility age decreased from 70 to 65 in the five years from 1966 to 1970, and has stayed at 65 since.¹ There is a residence requirement, with those between 10 and 40 years of residence occurring after age 18 receiving partial benefits. An income test for higher earners was put in place in 1989, and the benefits have

¹ The Canadian Government announced plans to move the eligibility age to 67 starting in 2023.

been indexed for inflation since 1969. The current payment rate (as of December 2013) is \$550.99 per month, which is reduced by 15 cents for every dollar of income over \$70,954. The benefits are taxable income.

Guaranteed Income Supplement

The Guaranteed Income Supplement was introduced in 1967 to complement the other components of the system by targeting benefits at lower income seniors. Those over age 65 receive an income-tested supplement which is reduced by 50 cents with every dollar of other income (excluding Old Age Security income). The benefits are indexed and they are not considered taxable income. The current maximum amount is \$747.11 per month for singles; \$990.78 for couples. The Guaranteed Income Supplement has been expanded several times in its 46 year history, which has had a substantial impact on the incomes of those at the bottom of the senior income distribution. The benefit expansions happened in 1971, 1979-80, 1984, and 2006.

Allowance

Some seniors age 60 to 64 are eligible for a supplemental non-taxable benefit in two particular circumstances. The Allowance benefit is paid if someone is married to an Old Age Security recipient or if someone is a widow or a widower. The amount is approximately equal to the sum of the Old Age Security and Guaranteed Income Supplement payments, and is reduced with other income at rates of 50 and 75 percent. In 2013, the amount is \$1,046.38 per month. The Allowance was introduced in 1975 and was increased each time the Guaranteed Income Supplement amount was increased.

Canada / Quebec Pension Plan.

The Canada Pension Plan is a contributory earnings-related plan. The province of Quebec operates a parallel and highly similar plan called the Quebec Pension Plan, but the discussion here will focus on the larger Canada Pension Plan. Benefits depend on earnings up to a cap (the Year's Maximum Pensionable Earnings—set at \$51,100 in 2013), measured on a monthly basis over the contributory period. The contributory period starts at age 18 (or 1966—whichever is later) and goes until benefits are taken up, or age 70. There are provisions to throw out the bottom 15 percent of months over the contributory period. The final pension is calculated using three parts. First, the base replacement rate is 25 percent of covered earnings. Second, for each month in the contributory period the average ratio of earnings to the Year's Maximum Pensionable Earnings is calculated to give a number between 0 and 1. Finally, the pension is updated to current dollars using the average of the last five Year's Maximum Pensionable Earnings numbers. So, the pension is the product of $(0.25) \times \text{Average Earnings} \times \text{YMPE5yearAverage}$. The pension is taxable income and price-indexed once initiated. The amount paid in retirement benefits was phased in over 1967 to 1976, with benefits being multiplied by 10% in 1967, 20% in 1968, and so on up to 1976 when full benefits were paid.

The main eligibility age for the Canada Pension Plan is 65. In 1987, an early retirement provision was put in place allowing for benefits to be taken up as early as age 60 and as late as age 70. An actuarial adjustment of 0.5 percent per month (up to 30 percent) was applied to benefits taken

early or later. These actuarial adjustments were updated in 2011. The Quebec Pension Plan introduced a similar early-retirement option three years earlier, in 1984.

Importantly, there is a spousal survivor benefit that is paid to the surviving spouse in the event of death. The survivor benefit is 60 percent of the benefit that was received by the deceased spouse for those age 65 and over, and 37.5 percent of the benefit plus a flat amount for those under age 65. When the surviving spouse has a Canada Pension Plan pension of his or her own, the combined benefit cannot exceed the maximum Canada Pension Plan payment.

Other Retirement Income

Canadians also receive much government income from other sources. The largest other source is employment-related pension plans. In addition, tax-assisted savings accounts allow for income tax to be deferred for funds placed in the accounts. Both of these sources of funds are taxable when received, and are important for figuring out eligibility for the income-tested guaranteed income supplement.

Empirical Strategy

The empirical strategy we employ uses simulations to exploit the variation in the institutions of the Canadian retirement income system through time. We use a simulator to project retirement income of an elderly family, using four pieces of information: a given earnings history, an age path for other pension or capital income, a family structure, and an age-retirement path. For the

simulations, each of these four pieces is fixed, so that the output from the simulations has variation that only embodies the differences in the structure of the retirement income system in Canada across ages and years. Below, we first give some more detail on the simulations, then explain how we use the output of the simulations in our regression analysis.

Simulation methodology

The aim for the simulations is to produce a projection of public pension retirement income that reflects the rules and laws that were in place for each age-year combination in our data. To do this, our approach builds the projection by taking inputs that are common across cohorts and putting those common inputs through the simulator that incorporates time-varying rule and program changes. In this way, the resulting estimates will not embody potentially endogenous individual or cohort-specific characteristics such as changing earnings patterns, marriage rates, or retirement rates. Instead, only the institutional variation across ages and years will drive the variation in simulated pension income.

The four inputs that form the building blocks for our simulation are an earnings history, an age path for pension or capital income, a family structure, and an age-retirement path. We describe each of these in turn below. Common across these building blocks, we use 1990 as a focal year, as our data (as described below) are centered around 1990. A 65 year old in 1990 was born in 1925, so birth cohorts 1920 to 1930 are used as the base cohorts for all of the simulations.

For the earnings histories, we construct a history for each sex, age, and marital status (married or single) for a given year of birth. Included are earnings from all sources, including self-employment. In each of these cells, we record nine decile cutoffs (10th percentile, 20th percentile, etc. up to 90th percentile). Our income data only go back as far as 1971, so we project earnings back from 1971 by applying an average wage growth series to the 1971 data before forming our cells.² Because the Canada Pension Plan only counts earnings from 1966, we only require the back-casted component of the earnings histories for a maximum of five years, from 1966-1970. The common cohort that we select to form the simulated public pension benefits takes an average of the earnings histories for the birth cohorts from 1920 to 1930, centered around 1925. In 1990, these cohorts are between the ages of 60 and 70.

An imputation of non-labor income is necessary to calculate income taxes and to impute Guaranteed Income Supplement properly. For this purpose, we follow a similar procedure as described above for earnings, with separate series by year of birth, marital status, and sex. We impute non-labor income (comprised of investment income and employer-sponsored pension income) to each of these cells for each age from 55 to 80. The decile cutoffs for this path for non-labor income are then recorded.

Given an earnings history and a path for non-labor income, we simulate potential public pension income at each age from 55 to 80. Retirement income at each age, though, depends on when retirement occurred because the Canada Pension Plan benefits are actuarially adjusted for age of retirement. So, a given 75 year old with the same earnings history and non-labor income path might have different Canada Pension Plan benefits (and also different Guaranteed Income

² We use the Historical Statistics of Canada, CANSIM series 11-516-XIE for this purpose.

Supplement benefits) depending on the age of retirement. We therefore simulate potential retirement ages from 55 to 70 for each age of potential income receipt. That is, for an observed 80 year old, we simulate benefits for the case when that person retired at age 55, 56, 57, and so on up to age 70.

We then average across the 16 potential years of retirement (ages 55 to 70) using a set of age-specific probabilities of retirement. We obtain these probabilities by comparing the proportion not working across ages within a birth cohort. For example, if the proportion not working rises from 30 percent to 35 percent at a given age, then 5 percent of the population is deemed retired at that age. We do this calculation for birth cohorts defined by year of birth, sex, earnings and decile using our annual income data. We then calculate these probabilities by comparing within a cohort across years of the survey—for example 55 year olds in 1995 compared to 56 year olds in 1996. We use a fixed set of birth cohorts from 1920 to 1930 to form the retirement probabilities that we apply to all birth cohorts in the simulation. Importantly, this means that any retirement behavior that may have changed because of, for example, the introduction of the early retirement option in 1987, will not be captured in our simulations.

We repeat this exercise for each of fourteen different family types in order to capture the important differences across different kinds of families. This is especially important when comparing across ages after 65. For example, while only 8 percent of families at age 55 to 59 are widowed women, 41 percent of families are widowed women in the age range 75 to 79. The family types we use are married, single, and single-widowed. For the two single types, we have both male and females. For the married types, we use ten different age gaps between the man and

women.³ The four single types and the ten married types add up to fourteen different family types to consider. We use data from the 1990 General Social Survey to create the age-specific probabilities for these 14 family types for all age between 55 and 80 in 5-year age groups.

The goal of the simulations is to produce age-year cells of potential benefits. To go from what has been described above to these age-year cells requires that we average over states of the world, in a few different dimensions. To reduce down to age-year cells, we require two sets of probabilities. The first is to average over different family types. We have simulated a different benefit for each earnings history, depending on family type. By averaging over the benefits in each of the family types using the probability of being in that marital state, we get an overall average. Similarly, we have 16 potential years of retirement. We use the probabilities of retirement at each age and multiply them by potential benefits assigned to each possible age of retirement to find the overall average. Finally, we have benefits over 9 different decile cutoffs. We take a simple equally weighted average over the benefits corresponding to each of these decile cutoffs.

These calculations and simulations can all be summarized in one large summation in the following equation.

$$B_{ay}^{sim} = \sum_{decile=1}^9 \frac{1}{9} \sum_{age=55}^{70} \sum_{family=1}^{14} \Omega_{ay}(Y_{rdfb}, C_{abfd}) \times P_a^f \times P_{adf}^r \quad (1)$$

³ The age gaps are: man older by 10, 8, 5, 4, 3, 1 years; same age; woman older by 1, 2 years.

In equation (1) we make use of the following notation:

B_{ay}^{sim}	Simulated benefits for those age a in year y .
Ω_{ay}	‘Benefit function’ which determines benefits at age a in year y , given the laws in place and an earnings history and a non-labor income level.
Y_{rdfb}	Earnings history up to retirement age r for decile d and family type f and birth cohort b .
C_{abfd}	Non-labor income at age a for birth cohort b and family type f and decile d .
P_a^f	Probability of being in family type f at age a .
P_{adf}^r	Probability of having retired at age r for someone currently at age a in decile d and family type f .

The fraction $1/9$ is in equation (1) because we want to take a simple unweighted average of the benefits across the 9 different lifetime earnings deciles. By averaging across these deciles, we can pick up more of the non-linearities in the benefit function (for example, the income tested Guaranteed Income Supplement) that we would miss if we just focused on the median or the mean income level.

Regression Specifications

The regression analysis aims to explain the pattern of income and consumption across age and year cells as a function of public retirement pension benefits. Our main specification controls

both for age fixed effects (α_a) and for year fixed effects (τ_y), along with a small set of characteristics (X_b) that varies by birth cohort. Finally, there is an i.i.d. disturbance term ε_{ay} .

The base regression thus takes the following form.

$$Y_{ay} = \beta_0 + \beta_1 BEN_{ay} + \beta_2 \alpha_a + \beta_3 \tau_y + \beta_4 X_b + \varepsilon_{ay} \quad (2)$$

The dependent variables Y_{ay} are wellbeing measures drawn from various measures of the income and expenditure distribution. The details of these measures are provided in the next section. The BEN_{ay} variable is the *observed* total of public pension retirement income in our microdata; distinct from the simulated counterpart B_{ay}^{sim} , as the observed BEN_{ay} variable comes straight from the data and reflects each individual's actual earnings history and non-labor income. In contrast, the simulated measure B_{ay}^{sim} uses a fixed earnings history and non-labor income schedule, but allows benefits to vary across ages and years because of program rules (according to the benefit function Ω_{ay}).

Some birth cohorts may have experienced different economic trajectories, such as shocks to earnings, investment returns, or lifespan. Because of this, it is possible that these kinds of cohort level shocks could affect both the observed wellbeing measure Y_{ay} and the public pension retirement benefits BEN_{ay} . In equation (2), this would manifest as a correlation between BEN_{ay} and ε_{ay} , biasing the estimated coefficient β_1 . In light of this, we implement an instrumental variables strategy that makes use of our simulated benefits to correct for this endogeneity. Because our simulated benefits B_{ay}^{sim} don't embody any birth cohort-specific information by

construction, it only varies by age and year. With age and year dummies included in the specification, the benefit variable can be excluded from the second stage of the instrumental variables analysis. The two stages comprising this instrumental variables analysis are in the equations below.

$$BEN_{ay} = \theta_0 + \theta_1 B_{ay}^{sim} + \theta_2 \alpha_a + \theta_3 \tau_y + \theta_4 X_b + \gamma_{ay} \quad (3)$$

$$Y_{ay} = \beta_0 + \beta_1 \widehat{BEN}_{ay} + \beta_2 \alpha_a + \beta_3 \tau_y + \beta_4 X_b + \delta_{ay} \quad (4)$$

Equation (3) is the first stage, where we regress the observed benefits on the simulated benefits, along with age and time fixed effects and the vector of cohort characteristics. The second stage equation is (4), where the predicted value from (3) \widehat{BEN}_{ay} is used to explain the outcome variable. We also make use of a ‘reduced form’ specification as depicted in equation (5) that uses the simulated pension benefit B_{ay}^{sim} directly in the estimation.

$$Y_{ay} = \beta_0 + \beta_1 B_{ay}^{sim} + \beta_2 \alpha_a + \beta_3 \tau_y + \beta_4 X_b + \delta_{ay} \quad (5)$$

Data

The data we use allow us to form age-year cells of various measures of income and expenditures. Taken together, our data span a more-than-40 year period from 1969 to 2010. Below, we first describe the different sources of income and consumption data that we bring together. After that, we provide detail on the precise measures of wellbeing that we form.

Data Sources

We draw on annual surveys of income and expenditures to form our data. For both income and expenditures, there is a survey that covers the first part of our time span that was eventually replaced with a new survey in the 1990s. We describe these below.

The income surveys we employ are the Survey of Consumer Finances (SCF) and the Survey of Labour and Income Dynamics (SLID). The SCF was an intermittent survey, for which we have data starting in 1971; and continuously from 1984 to 1997. From 1971 to 1984 the survey occurred at least once every two years. The SLID begin in 1995 and continued on to 2011. For this paper, we use the years 1998 to 2010.

Both income surveys took a large sample of about 50,000 Canadians in each year and recorded detailed information about their family situation, their job and employment, and all sources of income. The SLID is distinguished from the SCF because of the use of administrative tax data to obtain the income information. Respondents were given the option to have their tax data matched to their survey responses instead of answering the section on income sources. Most took this option. The data are available at various levels of aggregation, ranging from the individual to a broad measure of the family. To maintain consistency over the entire period, we use a narrow definition of the family (the Census Family) that incorporates spouses and never-married children still living at home. For those over age 65, this means we include only the income of the elderly couple and not those of anyone with whom they live. Not only is this the only family aggregation available in early years of the SCF, it is also the appropriate choice for calculating tax liabilities.

There were also two different series of expenditure surveys over this period. First, the Family Expenditure Survey began in 1969 with a survey of a sample of Canadian families detailing their expenditure patterns. New cross-sectional surveys were taken in 1974, 1978, 1982, 1984, 1986, 1990, 1992, and 1996. This survey was replaced in 1997 with the annual Survey of Household Spending, which is available up to 2009. There were small changes in content and also a change in the unit of analysis to the household. Sample sizes varied sharply by year for the Family Expenditure Survey, ranging from about 3,700 in 1984 to 10,050 in 1996. The Survey of Household Spending started with about 15,000 observations, but by 2009 had fallen to a sample size of about 9,000.

For all our analysis, we exclude Quebec from the sample because of the differences between the Quebec and Canada Pension Plans. We use the provided sample weights to provide representative results in all cases.

We use a small set of cohort characteristics in the estimation. All of these variables are measured as an average of the values between ages 50 and 54 for a given cohort. The variables include the proportion of the cohort married, the average earnings of the cohort, the proportion with at least some post-secondary education, and the proportion with a tax-preferred savings account or employment-based pension.⁴ The point of including these variables is to try to control for cohort-specific trends that may be correlated with income or expenditure patterns. For example, later cohorts had higher lifetime earnings and education attainment than earlier cohorts. Controlling for these factors helps to distinguish cohort effects from the impact of policy.

⁴ Specifically, Registered Retirement Savings Plans and Registered Pension Plans.

Income and Expenditure Measures

Taking the income and expenditure data from the surveys, we form several ways to shed light on the distribution of wellbeing. We calculate some measures of low income and low expenditure, and then move onto standard measures of percentile cutoffs. We start below, though, with an explanation of how we formed our income and expenditure measures.

We form the income and expenditure measures in similar ways. For income, we take after-tax annual total family income as reported in the survey. We then adjust this after-tax income measure by a factor to account for family size. For this, we use an equivalence scale of 1.0 for the first adult and 0.7 for the second adult. All dollar values are then adjusted using the consumer price index to 2010 values. For expenditure we use current household expenditure as recorded in the surveys. This measure takes total household expenditures and subtracts off personal taxes, insurance payments, pension contributions, and gifts/contributions. We also adjust for family size in the same way as for income as discussed above. All expenditure values are then transformed to 2010 values using the consumer price index.

With these income and expenditure measures constructed, we take the mean by age-year cell and also record several percentile cutoffs (P10, P50, P90). In order to assess the adequacy of public pension income, we form a measure of relative poverty by taking the median of the prime age distribution (with a reference person of age 25 to 54) and drawing a poverty line at 50 percent of this median.⁵ Elderly families under this line are recorded as being in relative poverty. We use

⁵ See Milligan (2008), Baker, Gruber and Milligan (2009), or Schirle (2013) for an exploration of different measures of low-income for the elderly.

the same procedure for both income and expenditures. For income, we also use a low-income line in common use in Canada called the Low Income Cut Off (LICO).⁶ This line was last set in 1992, but can be applied to other years using the consumer price index.

Results

We present the results of our analysis in several steps. First, we explore the patterns in our dependent variables—the wellbeing measures. Second, we analyze the correspondence between our simulated measures and the actual benefits received by Canadians. Finally, we present our regression results.

Trends through time

We begin by describing the trends in the income and expenditure data for elderly families, where we define elderly families using the age of the reference person in the family or household. In several of the graphs below, we also include the median of ‘prime age’ families, where the reference person is age 25 to 54. All income and expenditure numbers have been adjusted for family size, so the reported numbers reflect results per effective adult. We go through both income and expenditure measures, looking first at percentile cutoffs from the distribution and then examine various poverty measures.

The first graph in Figure 1 shows the 10th, 50th, and 90th percentile cutoffs for elderly income by year. We also show the 50th percentile cutoff for prime age families. There has been growth in the 1971 to 2010 period across the elderly income distribution. At the bottom, the 10th percentile cutoff in 1971 was \$5,608. By 2010 this reached \$13,950. Most of the growth at the bottom

⁶ See Statistics Canada (2013) for details on the LICO.

occurred in the 1970s and 1980s, as the Guaranteed Income Supplement was increased and as newly elderly cohorts with Canada Pension Plan benefits displaced elderly cohorts who did not participate in the Canada Pension Plan. Growth at the top of the distribution followed a similar path, but with an extra spurt upward following 2000. This reflected the rise of dual earning couples, as well as cohort differences in earnings that benefited these elderly through higher Canada Pension Plan benefits. In Figure 2 we graph the same data, with each line starting at an index level of 100 in 1975. The percentage gain for income is greatest at the bottom of the income distribution, and least for the 90th percentile cutoff. This suggests a compression of the elderly income distribution over this period.

We repeat this analysis for the expenditure data in Figure 3 and Figure 4. The increases in income in the 1970s and 1980s from the earlier figures appears to have a somewhat muted effect on expenditure, with a rise of only about 20 percent over by 1990 over the levels seen in 1974. The growth in expenditure through the 2000s was widely felt across the distribution, but those at the top appeared to enjoy slightly greater growth than at the middle or bottom.

We now turn to the patterns in poverty using the low-income and low-expenditure rates described earlier. Figure 5 focuses on the relative income and expenditure measures, as well as the constant-line LICO for income. Through time, the proportion of the elderly under LICO has fallen from just under 40 percent to 16 percent by 2010. The relative measures both show similar, though more muted, declines from the 1970s to the mid-90s. After the mid 90s, the relative income poverty measure increases from 8.4 percent in 1995 to 14.4 percent in 2010. As can be seen in Figure 1, incomes at the 10th percentile cutoff have been growing. However, they

have not been growing as fast as those for prime-age families. For this reason, more of the elderly families are falling under the upward-trending cutoff line.

The next two graphs compare across ages and decades to get a more clear picture of where poverty has decreased. In Figure 6 we display the relative income measure by age, averaging across all available years in each of the four decades. The 1970s are the clear outlier here, with relative income poverty trending steadily upward with age. At that time, the Canada Pension Plan was just beginning to affect elderly incomes and the Guaranteed Income Supplement was at a lower level. In the 1980s, 1990s, and 2000s, there is a clear drop-off in elderly income poverty at age 65. This is consistent with the age at which Guaranteed Income Supplement eligibility and 'full' Canada Pension Plan eligibility are reached. After age 65, the 1990s show the lowest rates, with the 2000s slightly higher. In contrast, for the age 55 to 64 range, it is the 2000s that show the highest low-income rates, three to five percentage points higher than the 1990s.

The same age analysis by decade for expenditures is presented in Figure 7. The smaller sample sizes make the data noisier, but there is a clear pattern of lower expenditure poverty rates across the decades. However, in some contrast to the income data in Figure 6, there is no sharp drop-off at age 65. This suggests that individuals at those pre-65 ages may be using savings, charity, help from family, or other mechanisms to smooth their expenditures at the ages before full pension eligibility at age 65.

Overall, these patterns display a clear increase in the incomes and expenditure of the elderly over this four decade period. The patterns observed here are consistent with the timing of the impact

of the changes to the public retirement income system, suggesting that the introduction and expansion of the public measures has had some impact on elderly wellbeing.

Simulated retirement income measures

We now turn to the trends in public pension income, comparing actual data to the imputed measures that come out of our simulations. In Figure 8, we plot total simulated public pension benefits for different ages and years, summing together income from Canada Pension Plan, Old Age Security, and Guaranteed Income Supplement / Allowance. We don't make any family size adjustments for these data, and they are on a pre-tax basis. There is strong variation across ages through time. In the early 1960s, the only benefit available was Old Age Security at ages 70 and higher. To this was added Guaranteed Income Supplement and the Canada Pension Plan in 1967, both available (but small initially) at age 65. The jump in benefits at age 65 in 1970 reflects the change in eligibility age for Old Age Security down to age 65. The growth in benefits through the 1970s was driven by the phase-in of the Canada Pension Plan. In 1987, the introduction of early retirement had a large impact on benefits at age 60, and also for benefits at age 65.⁷

Benefits at age 75 flatten out around 1997 and those at age 70 flatten out around 1992. This is driven by the changing nature of the Canada Pension Plan benefit formula, and specifically the Years Maximum Pensionable Earnings. The YMPE grew by 36.6 percent in real terms in the decade leading up to 1987. In contrast, the real value of the YMPE was approximately constant

⁷ At age 65 there is a large jump at 1987. This happens because in the years before 1987 people retiring before age 65 took a number of years of zero earnings into their Canada Pension Plan calculation, which lowered benefits. For example, someone retiring at age 61 would have 4 years of zero earnings before the Canada Pension Plan retirement benefits could be taken up at age 65. With the advent of early retirement in 1987, those who retire before age 65 can now take up benefits earlier. This leads to a transitory bump in benefits measured at age 65.

in real terms from then until 2010, with only slight variation. This means that those hitting retirement ages since 1987 have seen little growth in benefits. In contrast, before 1987 those waiting a year to retire saw a large real increase in benefits because of the fast-growing YMPE.

The next two figures explore how well our simulations fit the actual data drawn from the surveys. The first is Figure 9, which shows simulated and actual benefits at age 70. The fit is quite good, matching both the level and the trend of benefits. The graph at other ages looks similar to this one for age 70. In Figure 10, we show the actual and simulated benefits in 2010 across all ages. Again, the fit looks very good.

We now decompose the simulated benefits into the three main benefit components at age 70. The phase-in of the Canada Pension plan between 1967 and 1976 is visible. As well, continued growth of Canada Pension Plan benefits in the 1980s can be observed up until 1992, when the YMPE growth of the 1980s stopped being reflected in the benefits. Old Age Security benefits are almost constant after full inflation indexing was introduced in the 1970s. Also visible in this figure are the increases in Guaranteed Income Supplement that is important for lower income elderly.

The final graph shows the difference in simulated public pension benefits across different lifetime earnings percentiles. Again we focus on benefits at age 70. In Figure 12, the Guaranteed Income Supplement pulls total benefits for those at the 10th percentile higher than the 50th or 90th. In the mid-1990s, however, continued growth in the value of the Canada Pension Plan exceeds

the price-indexed Guaranteed Income Supplement leading the higher percentiles to overtake the 10th percentile.

This exploration has demonstrated the strong growth in public pension benefits in Canada over the last forty years. Because of the cohort-based nature of the Canada Pension Plan, the variation across ages generated by its introduction lasted for decades as those who were 80 as late as 1990 had hit 65 during the 1967-1975 phase-in period. The role of the Guaranteed Income Supplement as a top-up is also notable, given that it leads to small differences in total benefits across those with very different earnings histories. Finally, the simulations we performed appear to match the data quite well, in spite of the limitations of common inputs (earnings histories etc.) we imposed on them.

Regression results

We present regression results first for income then for expenditures. We use age-year cells for the income regressions. For expenditures, we use 5-age groupings and years to form the cells. All standard errors are robust adjusted and clustered on age. We run reduced form regressions using the simulated policy measure and also instrumental variables regressions, using the simulated benefit as the instrument.

For the simulated benefits, we use two different sets of measures. First, we use the age-year cell mean as our primary measure. However, the mean might not do the best job of picking up what is going on closer to the tails. So, we also find the benefits corresponding to just the 10th and just the 90 percentile of the lifetime earnings distribution. For the poverty measures, we also try the

benefits corresponding to the 10th percentile. Given the lack of variation in benefits across the earnings distribution visible in Figure 12, however, it's not clear this approach will be effective.

The first set of results, for income, is in Table 1. We begin by examining how well the simulated benefits predict the actual benefits in the top-left part of the table. This is the first stage relationship expressed in equation (3). For both the simulated and actual benefits, we use cell means. The estimated coefficient of 0.942 indicates that a one dollar increase in per-capita family benefits predicts actual benefits to increase by 0.942 cents. This corroborates the results in Figure 9 and Figure 10 which seemed to indicate the simulations matched the data quite closely.

The next result in the same column is for total income. This is a measure of adjusted after-tax income, using the equivalence scales and the reported after-tax income value. Here we are estimating equation (5), the reduced form specification. The coefficient indicates a one dollar increase in the simulated benefits is predicted to increase total income by 0.497 dollars. This result is strongly significant. On the right side of Table 1, we estimate the instrumental variables equation (4), with first stage equation (3). The instrumental variables result is very similar to the reduced form result, owing to the fact that the first-stage relationship between actual and simulated benefits is so close to dollar-for-dollar. There is no 'specific' formulation of mean income, so those parts of the table are left blank.

At the 10th percentile, there is a strong 0.716 response to an extra dollar of pension income in the reduced form specification using the simulated benefit directly. The first column uses the cell mean for the simulated benefit. In contrast, we use the 10th percentile of the earnings distribution

to form the simulated benefit for the result shown in the ‘specific’ column. Here, that result is slightly lower at 0.525. The instrumental variables estimates come in at 0.760 for the mean instrument and 1.106 for the specific 10th percentile instrument. This coefficient in the vicinity of 1.0 indicates that every dollar of public pension income is a net addition to the family’s after-tax income.

For the 50th percentile, there continues to be a strong effect of pension benefits on income, with a coefficient of 0.427 with the reduced form, and 0.454 with the instrumental variables specification. The specific instrument (the benefits associated with the median income level) does not seem effective in this case. At the 90th percentile, the point estimate is not much different from what was observed for P50, but the standard error is much bigger. The specific instrument strategy doesn’t do a good job, yielding negative estimates that are difficult to interpret.

The final two measures in Table 1 are the poverty indicators. Both the constant-value poverty indicator (LICO) and the relative poverty measure show strongly negative response to more public pension income. The benefit variables for these poverty specifications are scaled by \$1,000 in order to enhance the clarity of the table. The relative poverty measure will be used in the simulations below.

The second set of regression results focusing on expenditures is in Table 2. Here we have age group cells (five ages together) rather than single years of age, reducing sample size to 95. Total expenditures show a response of 0.410 to an extra dollar of public pension benefit income,

significant at the 10 percent level. The percentile cutoff results are mostly insignificant with high standard errors. The results for relative expenditure poverty, though, are strong and quite consistent with the income results in the previous table. Using the cell-average reduced form, the estimate of -0.0204 means that a \$1,000 increase in benefits leads to a -2.04 percentage point reduction in the rate of expenditure poverty.

The income and expenditure regressions indicate a fairly strong impact of public pension income on the bottom of the income distribution. Both the level of income at the 10th percentile and the poverty measures display an impact from higher public pension income. For expenditure, the most solid results are the relative poverty coefficients, which are strong and significant.

Simulations

The final part of our analysis focuses on the changes that have occurred in Canada's retirement income system over the past 50 years. To do so, we restrict our attention to the relative income and consumption poverty measures and ask the question: what would poverty rates have been if we imposed the retirement income system of different years on the individuals we observe in our data in the 2000s? We start by taking our reduced form estimates (based on equation (5)) for the relative income poverty specification reported in Table 1 and predicting the income poverty rate for each observation in our data for 2000 to 2010. We then swap in the simulated public pension income from the first year of each decade (1960, 1970, 1980, 1990, 2000, and 2010) for the benefit variable B_{ay}^{sim} and predict the relative income poverty rate given the new value for public pension income.

The changes in the system across these five decades have been widespread. The development of the different pensions was reviewed earlier, but we recap the salient aspects here with a particular focus of the differences across ages. Old Age Security was available in 1960, but only at age 70. Before that age, there were no public pension benefits at all until the Canada Pension Plan and Guaranteed Income Supplement started paying in 1967. Both of those benefits paid at age 65, but survivor benefits under the Canada Pension Plan could be paid at ages before 65. The Guaranteed Income Supplement was expanded at several points in the 1970s and 1980s, and the Allowance was added in the 1970s. The Canada Pension Plan was phased in over the ten years from 1967 to 1976, but those who turned 65 before 1967 received nothing. So, the impact of the phasing in of the Canada Pension Plan could still be seen at age 80 for fifteen years until 1982. In 1987, early retirement was initiated and allowed for actuarially adjusted payments as early as age 60. These details are summarized in Table 3.

In Figure 13, we plot the predicted relative income poverty rates across ages from 55 to 79. There are six lines; one for each of the six decades. We will start at the top, which is chronologically earliest. The poverty rate at ages after 64 is much higher, reflecting the unavailability of Old Age Security until age 70. After age 70, there is Old Age Security, but no Canada Pension Plan or Guaranteed Income Supplement. The predicted poverty rates are high, but not unreasonable given the observed poverty rates in the available data (see Figure 5 and Figure 6).

The next two policy lines show the impact of the phasing in of the Canada Pension Plan. The line for 1970 breaks slightly lower than the line for 1960 at ages before 65 because of the availability of survivor benefits. With the 1970 policies, age 65 becomes very important as the Canada Pension Plan and Guaranteed Income Supplement are available. However, only those age 68 and higher in 1970 receive any Canada Pension Plan because of its phase-in period. For the 1980 policy line, more generous Guaranteed Income Supplement payments and also a fully-phased in Canada Pension Plan lead to a sharp drop in relative income poverty at age 65.

The final three policy lines break off from the others at age 60 because of the availability of Canada Pension Plan benefits at age 60. Predicted income poverty is lower at these ages of 60 to 64 for all three of the 1990, 2000, and 2010 lines. The final difference to point out occurs at age 68 for the 1990 policy line. For those age 68 or older in 1990, there was no early retirement option. So, those exiting the workforce before age 65 had many zero earnings years in their Canada Pension Plan calculation. In contrast, the 2000 and 2010 policy lines reflect the fact that those retiring at 60 to 64 do not have to wait until 65 to receive benefits. Before 1987, those retiring before age 65 would record zero earnings years that hurt their eventual benefit while they waited to hit age 65. Because early retirees no longer take these years of zero earnings, this makes benefits higher in 2000 and 2010 than was the case in 1990, and thus poverty drops relative to the earlier ways.

We repeat the analysis for expenditures in Figure 14. Again, we use the data from 2000 to 2009 as the base for the simulations. Because of the need to use age groupings, the results are coarser

than for income. However, the same pattern emerges. Expenditure poverty is much higher under the counterfactual 1960 system than under the 2010 system.

Overall, these simulations have shown the large impact of the expansion of Canada's public pension system over the five decades from 1960 to 2010. To put some more precise numbers on the extent of the reduction, the average relative income poverty rate for ages 70 to 79 under the 1960 system was 33.9 percent. Under the 2010 system, relative income poverty dropped to 4.1 percent, for an 88 percent drop. For expenditures, the drop between 1960 and 2010 in the simulations is from 36.3 percent to 15.9 percent, which is a 56 percent drop.

Conclusions

In this paper, we study the impact of public pension income on measures of income and expenditures. The Canadian institutional environment is rich with variation, allowing us to implement an instrumental variables strategy based on a simulated instrument that picks out the institutional variation in benefits. We have several findings of interest. First, the expansions of the Canadian public pension system have led to large improvements in income and expenditure deprivation, as we find a significant impact of public pension benefits on poverty measures. Related, we find stronger responses of income at the bottom of the distribution (at P10) than at the top. Finally, our simulation evidence indicates that the five decades of development of the Canadian public pension system has yielded large and persistent decreases in relative income poverty. For those aged 70 to 79, the relative income poverty rate fell by 88 percent using the 2010 system relative to the 1960 system, and the expenditure poverty rate similarly fell by 56 percent.

Our estimates and simulations here suggest that public pension programs can have a large impact on elderly income and expenditures. As governments around the world adjust public retirement programs in response to longer lifespans, the impact of reforms on elderly wellbeing should be taken into careful consideration.

References

Baker, Michael, Jonathan Gruber, and Kevin Milligan (1999), "Retirement Income Security and Wellbeing in Canada" NBER Working Paper No. 14667.

Engelhardt, Gary and Jonathan Gruber (2006), "Social Security and the Evolution of Elderly Poverty," in Alan Auerbach, David Card, and John Quigley, eds. *Public Policy and the Distribution of Income*, Russell Sage Press, 2006, pp. 259-287.

Meyer, Bruce D. and James X. Sullivan (2010), "Consumption and income of the poor elderly since 1960", working paper.

Milligan, Kevin (2008) "The evolution of elderly poverty in Canada," *Canadian Public Policy*, Vol. 34, No. 4.

Milligan, Kevin (2010), "Incomes in the transition to retirement: Evidence from Canada". NBER Retirement Research Center Paper No. NB 10-09.

Milligan, Kevin (2012), "How is Economic Hardship Avoided by those Retiring Before the Social Security Entitlement Age?" NBER Working Paper No. 18051.

Milligan, Kevin and Tammy Schirle (2013), "The Retirement Income System and the Risks Faced by Canadian Seniors," CLSRN Working Paper No. 120.

Schirle, Tammy (2013), "Senior Poverty in Canada: A Decomposition Analysis," *Canadian Public Policy*, Vol. 39, No. 4, pp. 517-540.

Statistics Canada (2013), "Low Income Lines, 2011-12" Catalogue No. 75F002M No. 002.

Figure 1: Incomes of the Elderly

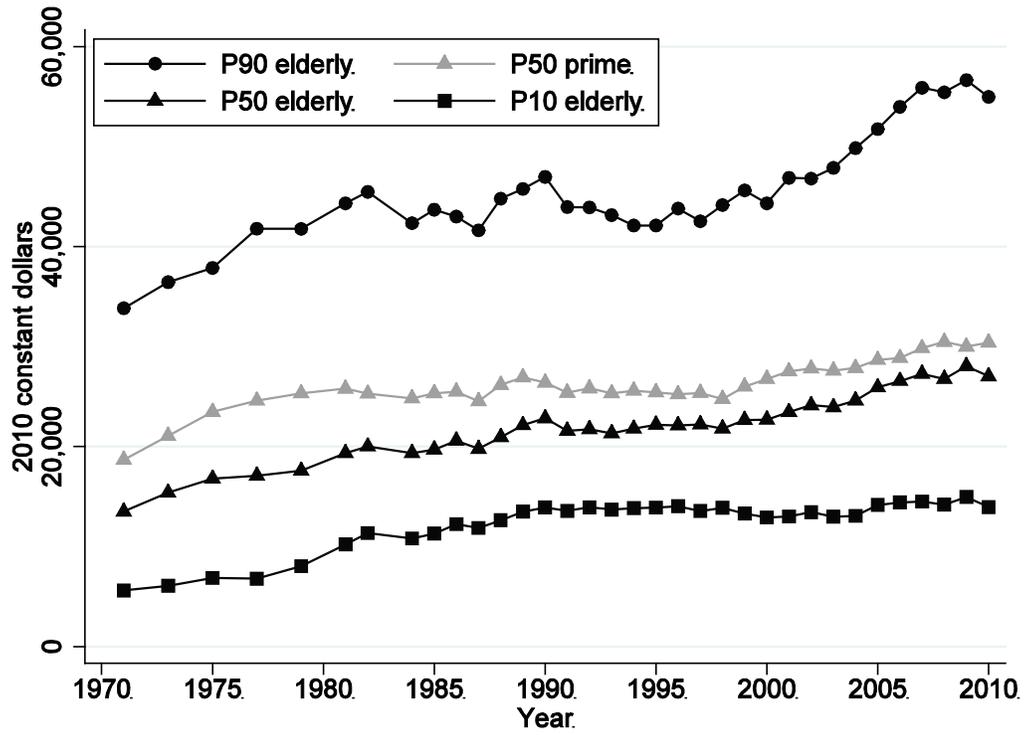


Figure 2: Indexed Income of the Elderly

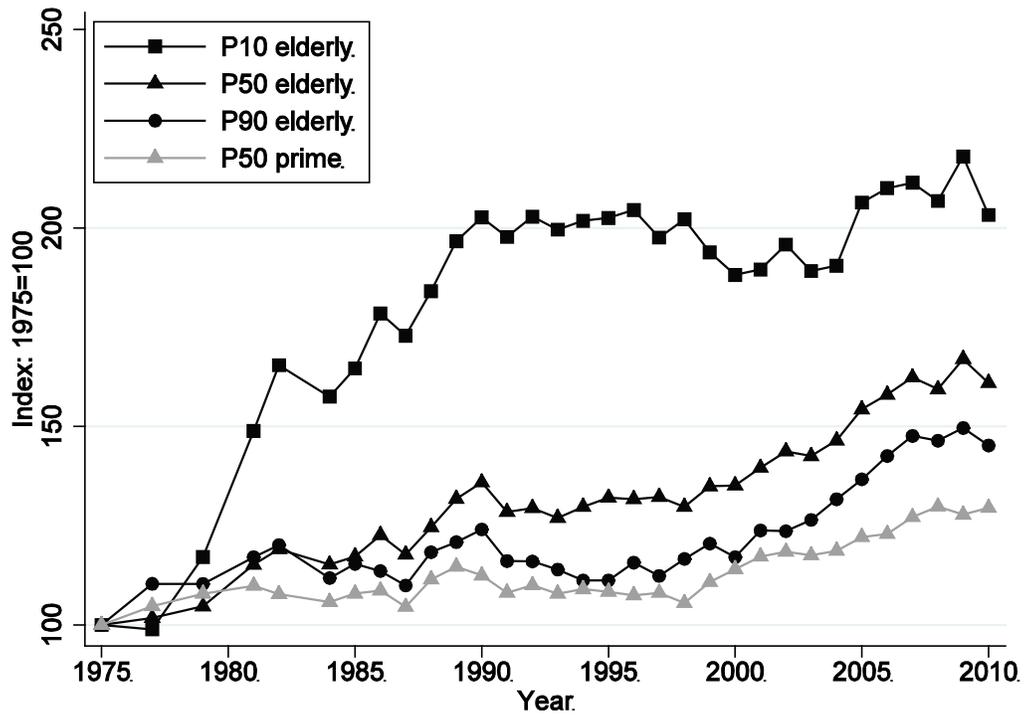


Figure 3: Expenditures of the Elderly

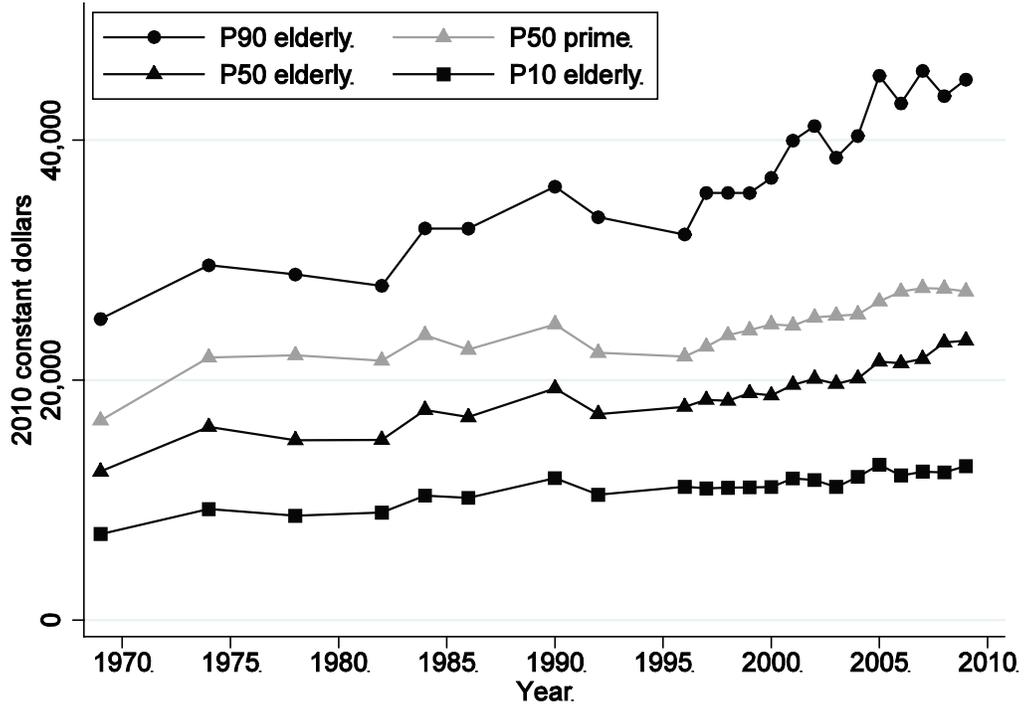


Figure 4: Indexed Expenditures of the Elderly

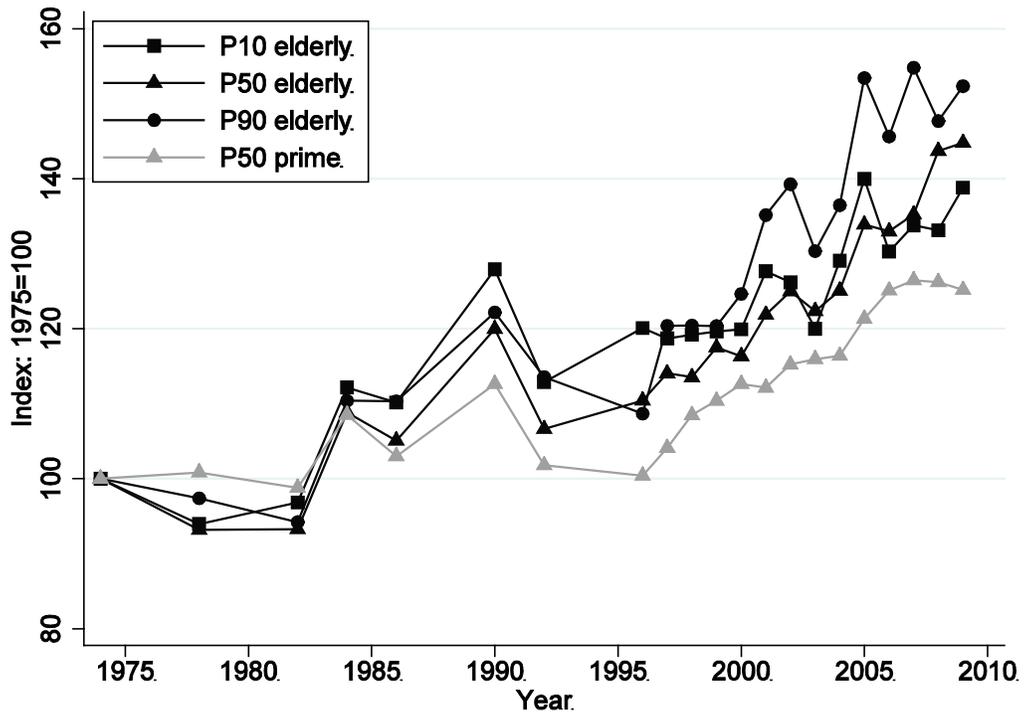


Figure 5: Poverty Measures by Year

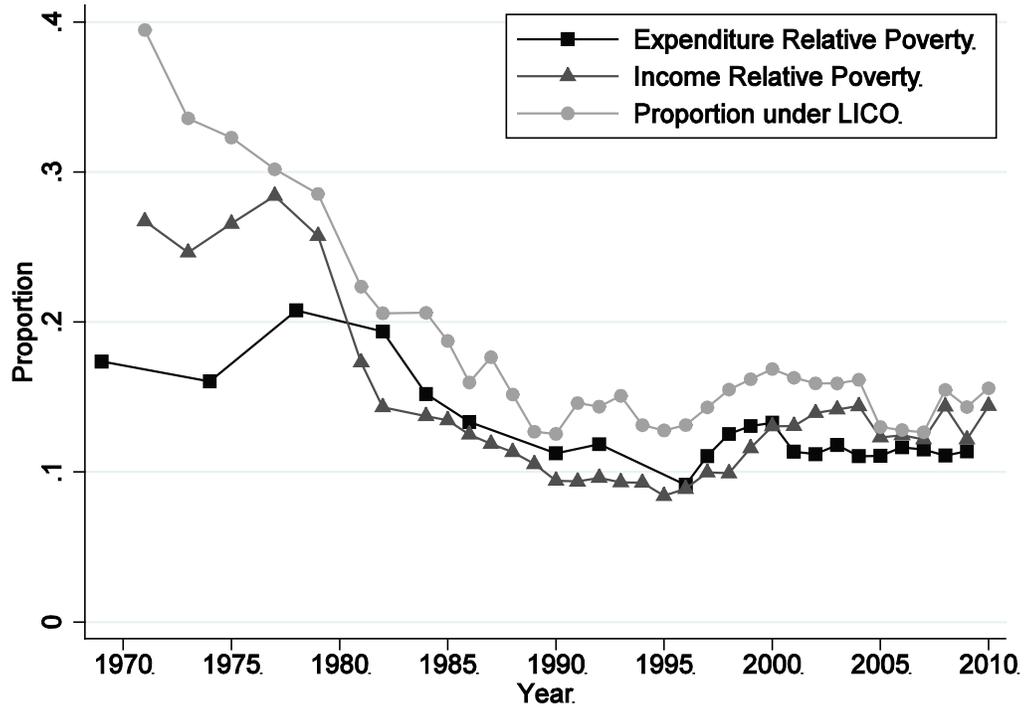


Figure 6: Relative Income Poverty by Decade

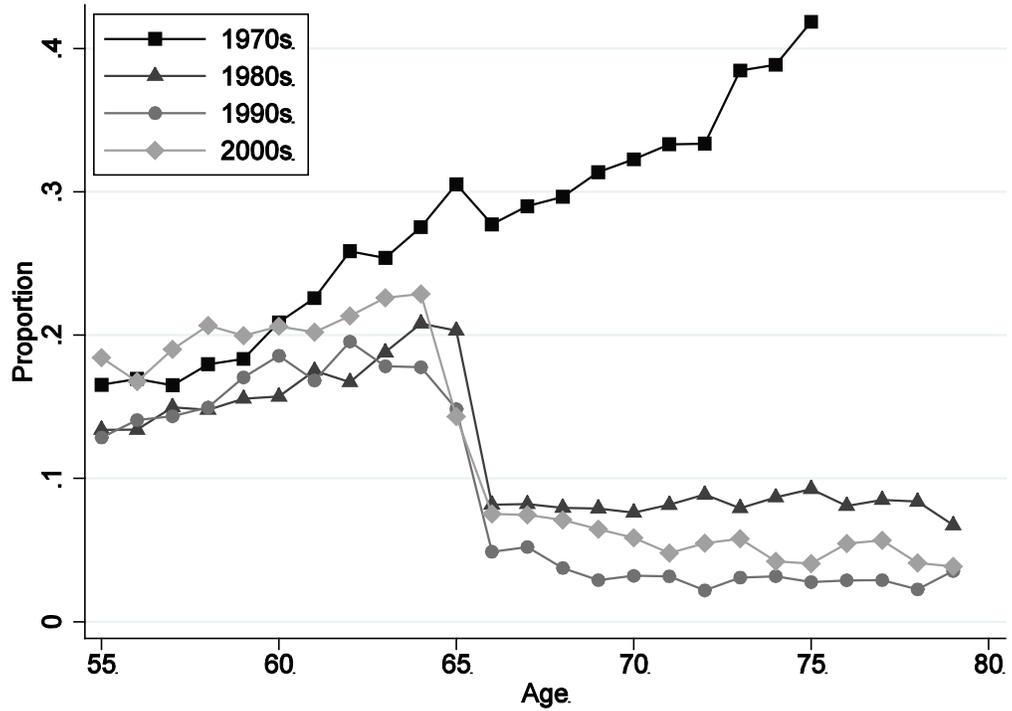


Figure 7: Relative Expenditure Poverty by Decade

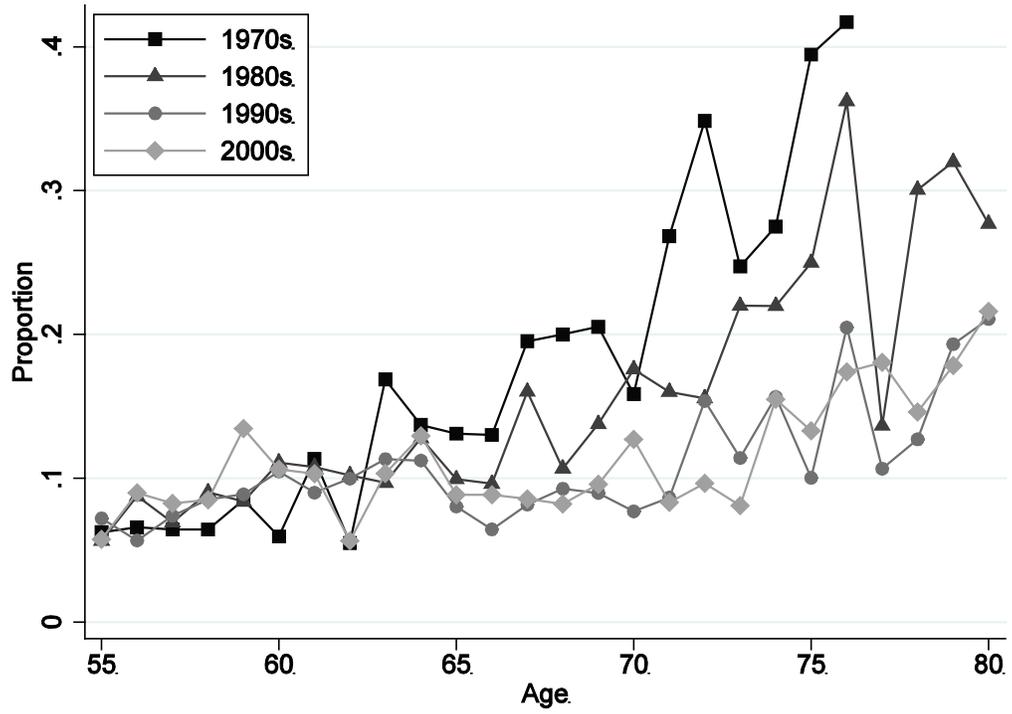


Figure 8: Simulated Public Pension Benefits by Age and Year

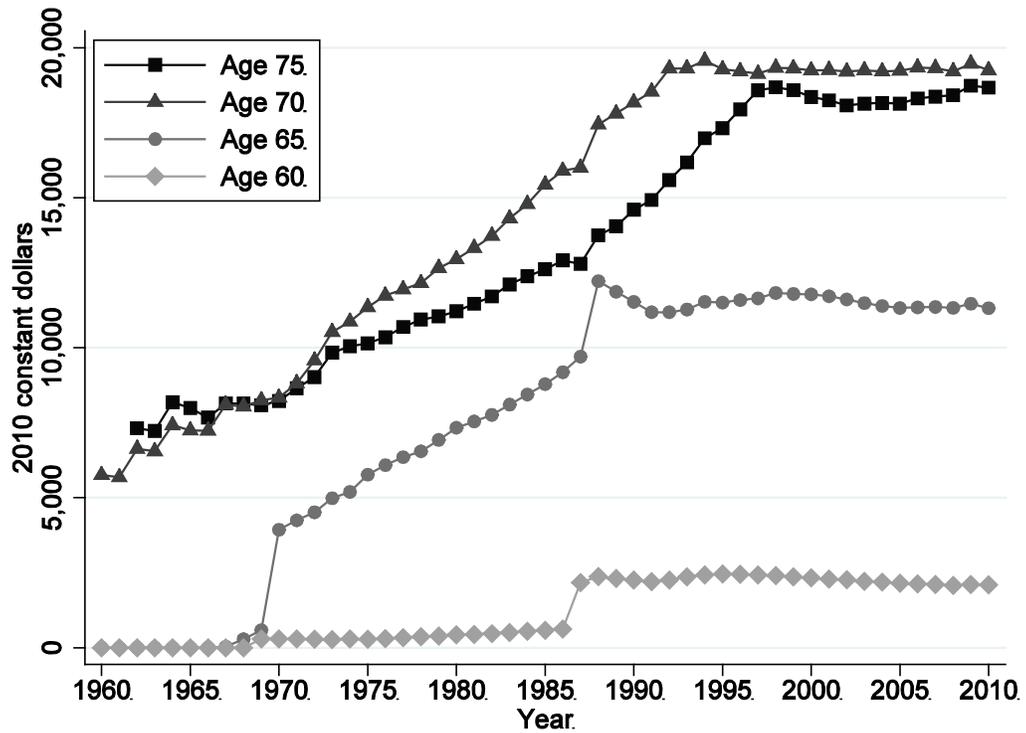


Figure 9: Simulated vs. Actual Public Pension Benefits at Age 70

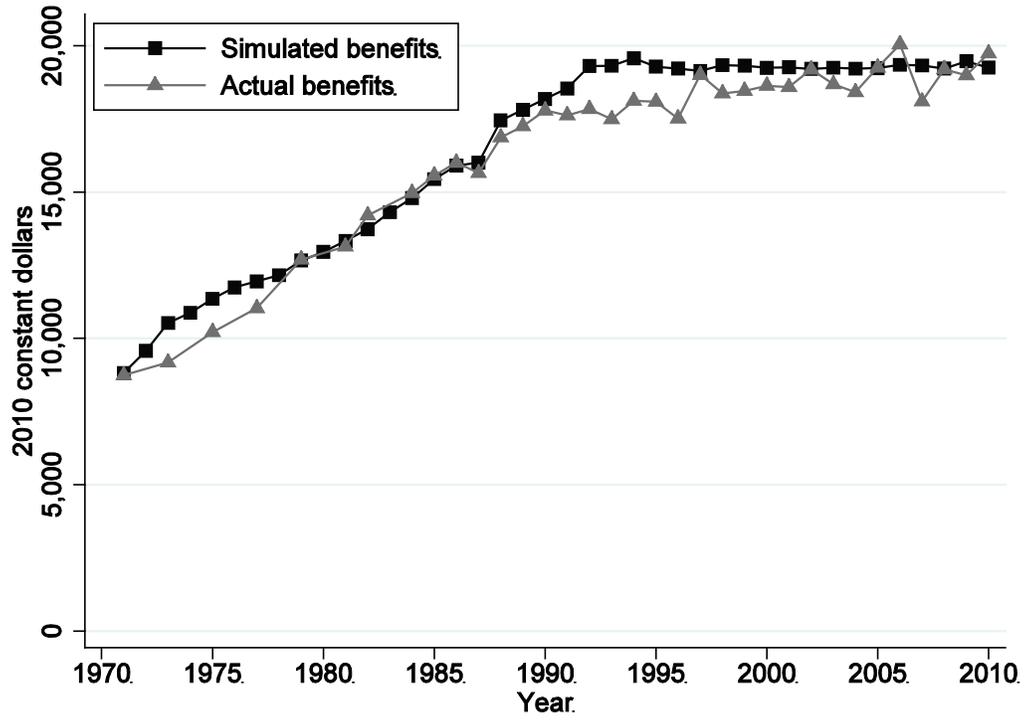


Figure 10: Simulated vs. Actual Public Pension Benefits in 2010

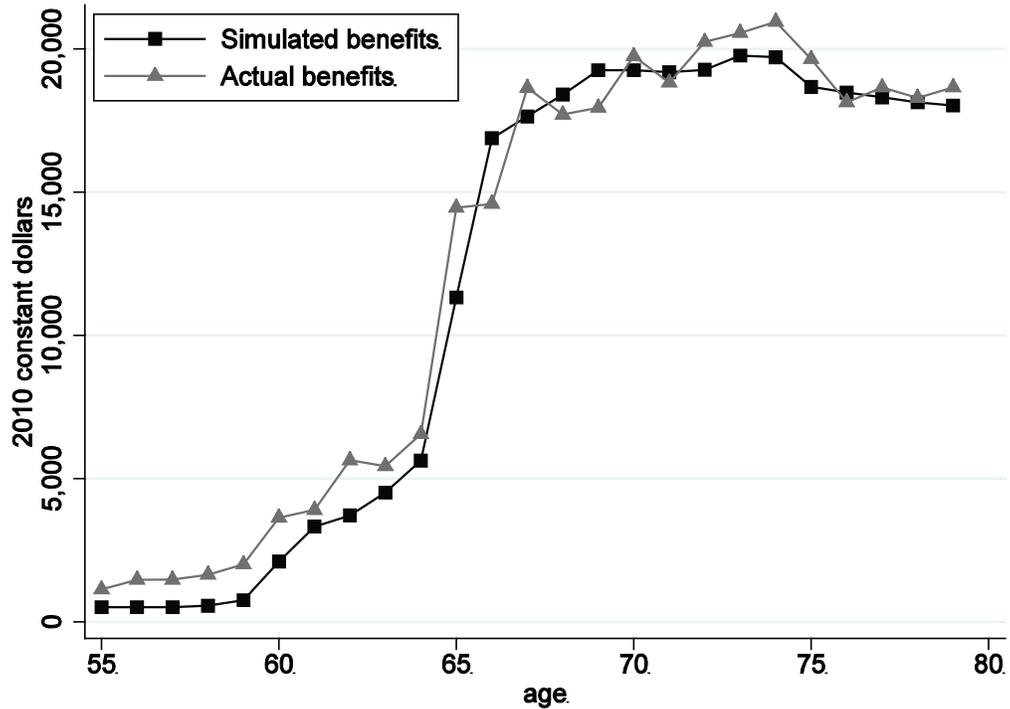


Figure 11: Simulated Composition of Public Pension Benefits at Age 70

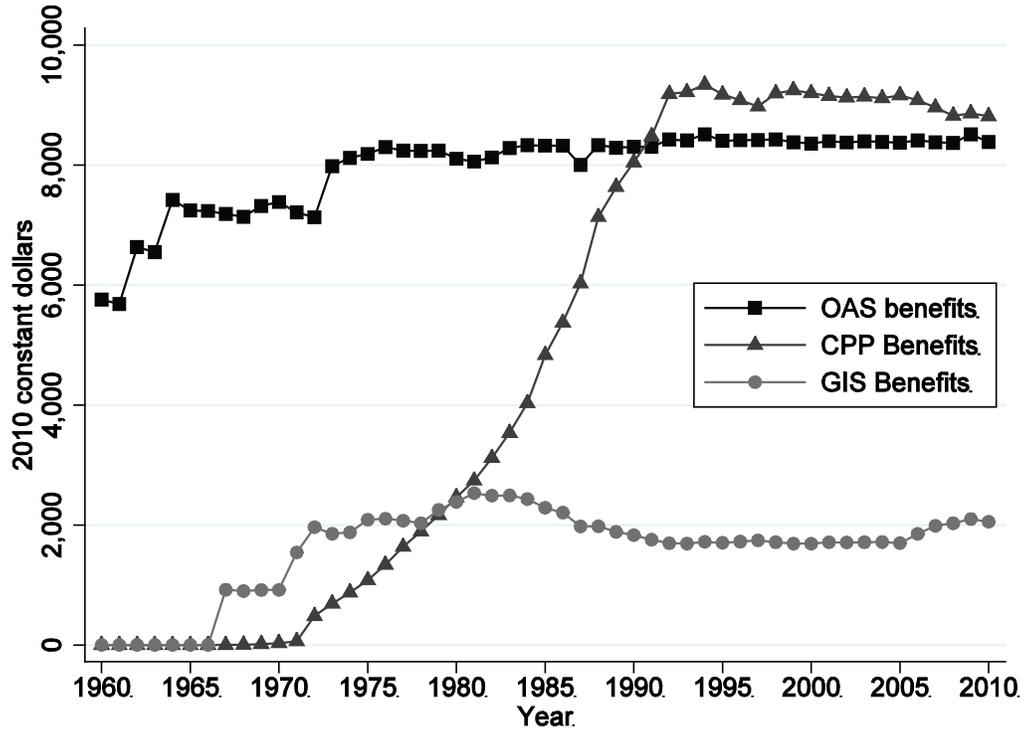


Figure 12: Simulated Total Public Pension Benefits by Lifetime Earnings Percentiles

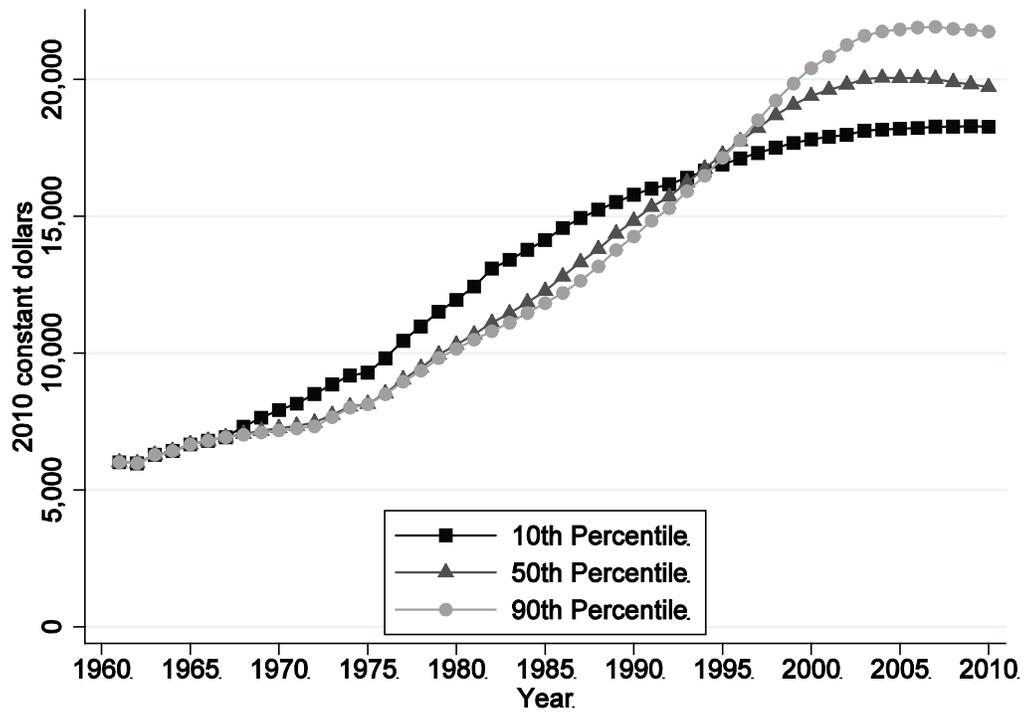


Figure 13: Counterfactual Relative Income Poverty Rates under Systems of Different Years

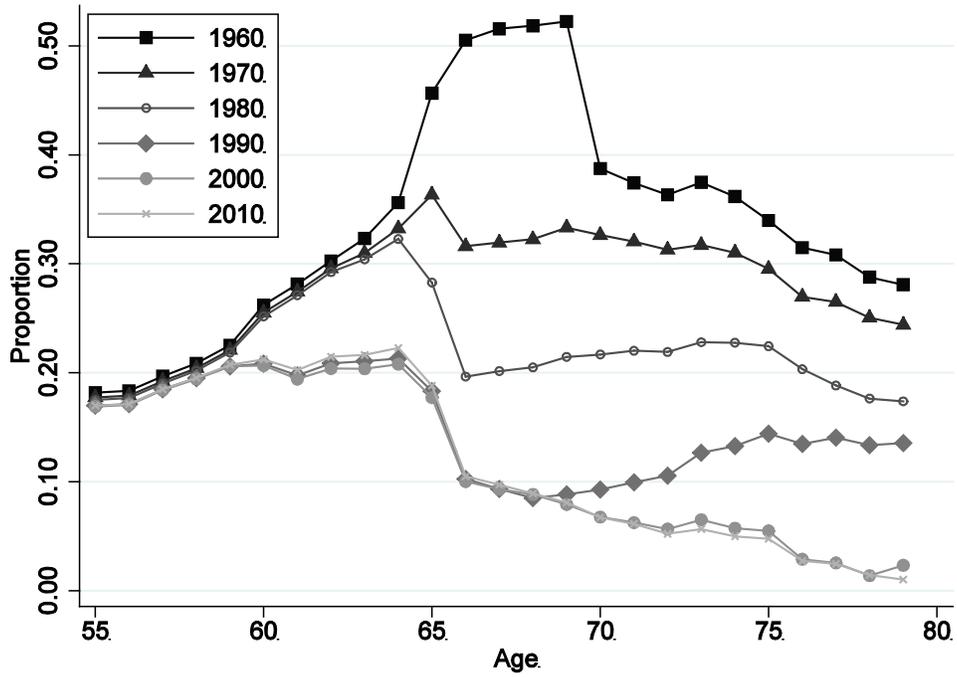


Figure 14: Counterfactual Relative Expenditure Poverty Rates under Systems of Different Years

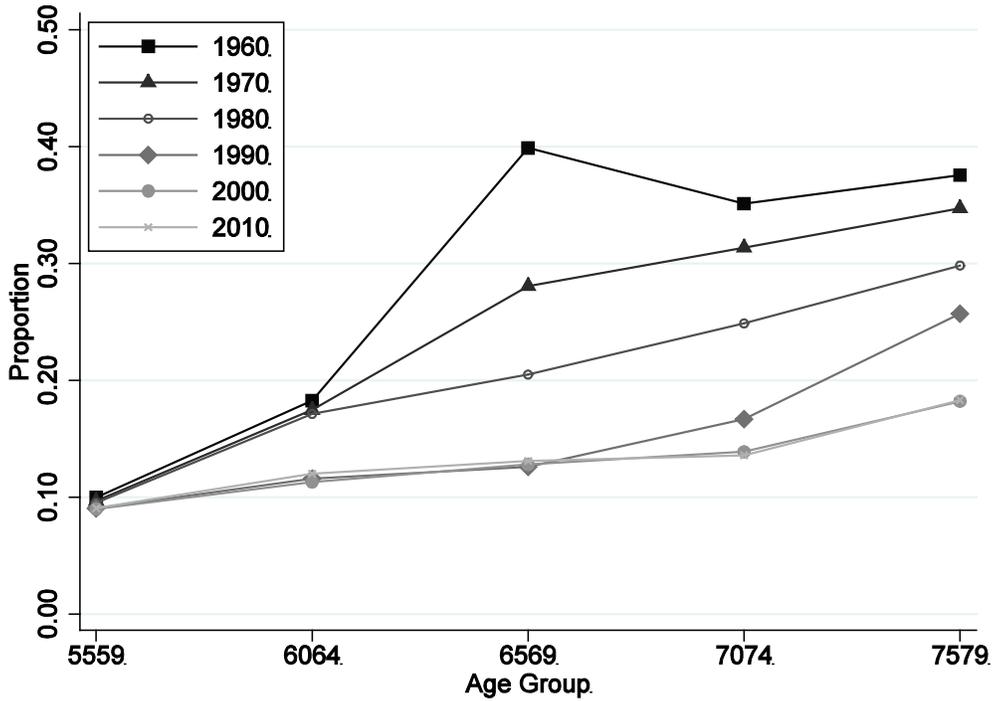


Table 1: Income Regressions

	Reduced Form		Instrumental Variables	
	Average	Specific	Average	specific
Sample size	829	829	829	829
Actual Benefits	0.942*** [0.0506]			
Total Income	0.497*** [0.138]		0.528*** [0.146]	
P10	0.716*** [0.104]	0.525*** [0.0551]	0.760*** [0.0939]	1.106*** [0.0803]
P50	0.427*** [0.119]	0.0942 [0.0579]	0.454*** [0.126]	0.259 [0.157]
P90	0.412 [0.337]	-0.377*** [0.130]	0.438 [0.362]	-0.978*** [0.341]
Under LICO	-0.0441*** [0.00470]	-0.0288*** [0.00386]	-0.0468*** [0.00408]	-0.0607*** [0.00553]
Under Relative Poverty	-0.0336*** [0.00385]	-0.0234*** [0.00323]	-0.0356*** [0.00339]	-0.0493*** [0.00435]

Notes: Reported are the coefficients on benefits, with each estimate and standard error coming from a separate regression. The regressions are on age-year cells. The 'average' specification shows the results when the age-year cell average is used. The 'specific' specification shows the results when the corresponding percentile of benefits from each age-year cell is used (P10, P50, P90 of benefits; P10 for the poverty measures). Three asterisks indicates statistical significance at the 1 percent level; two are for 5 percent; and 1 is for 10 percent. The standard errors in parentheses are clustered on age.

Table 2: Expenditure Regressions

	Reduced Form		Instrumental Variables	
	Average	Specific	Average	specific
Sample size	95	95	95	95
Total Expenditure	0.410* [0.180]		0.730 [0.516]	
P10	0.0711 [0.203]	0.117 [0.170]	0.127 [0.362]	0.169 [0.253]
P50	0.567* [0.204]	-0.023 [0.142]	1.009* [0.434]	-0.0797 [0.510]
P90	0.437 [0.777]	0.0196 [0.217]	0.779 [1.677]	0.0862 [0.958]
Under Relative Poverty	-0.0204** [0.00576]	-0.0368** [0.0132]	-0.0364* [0.0164]	-0.0535 [0.0266]

Notes: Reported are the coefficients on benefits, with each estimate and standard error coming from a separate regression. The regressions are on 5-age group-year cells. The 'average' specification shows the results when the age-year cell average is used. The 'specific' specification shows the results when the corresponding percentile of benefits from each age-year cell is used (P10, P50, P90 of benefits; P10 for the poverty measures). Three asterisks indicates statistical significance at the 1 percent level; two are for 5 percent; and 1 is for 10 percent. The standard errors in parentheses are

Table 3: Public Pension Eligibility at Different Ages Across the Decades

Year	Old Age Security	Guaranteed Income Supplement	Canada Pension Plan
1960	Age 70	Not available	Not available
1970	Age 65	Age 65+, small size	Age 65 40% phased in Survivor benefit before 65
1980	Age 65	Age 65+, medium size Allowance 60-64	Age 65 Fully phased in Survivor benefit before 65
1990	Age 65	Age 65+, full size Allowance 60-64	Age 60 with adjustments Fully phased in Survivor benefit before 65
2000	Age 65	Age 65+, full size Allowance 60-64	Age 60 with adjustments Fully phased in Survivor benefit before 65
2010	Age 65	Age 65+, full size Allowance 60-64	Age 60 with adjustments Fully phased in Survivor benefit before 65