Canadian Studies in Population 40, no. 3–4 (2013): 193–203.

The age pattern of retirement: Comparisons of longitudinal, pseudo-cohort, and period measures

Frank T. Denton
Ross Finnie
Byron G. Spencer[1]

Abstract

Retirement patterns change. When that happens, measures based on single-period cross-sectional data may provide a misleading picture of cohort retirement behaviour. We offer two cohort-based measures. One draws on income tax records to follow actual cohorts of individuals over time, the other on a time-series of cross-sectional surveys. We find that the retirement patterns based on the two approaches are similar but differ, often substantially, from single-period patterns. While pseudo-cohort measures can be assessed more quickly and at lower cost, knowledge of differences across income groups, income replacement rates, and so on, must rely on full longitudinal records.

Keywords: Retirement patterns, longitudinal measures of retirement, pseudo-cohort measures of retirement, age-retirement patterns.

Résumé

Les profils de départ en retraite par âge évoluent. Il en résulte que les statistiques fondées sur des observations obtenues pour une seule période donnent une vision faussée du départ en retraite de chaque génération. Nous présentons deux statistiques calculées sur des suivis de générations. La première utilise les données d'impôt sur le revenu pour observer un même groupe d'individus sur plusieurs années, la seconde juxtapose plusieurs coupes transversales. Les profils issus des deux approches sont similaires mais diffèrent, parfois notablement, des observations sur une seule période. Si les statistiques de pseudo-cohortes sont plus accessibles, il faut disposer de véritables observations longitudinales pour mesurer les différences entre catégories de revenu ou d'autres statistiques comme les taux de remplacement.

Mots-clés : modèles de retraite, mesures longitudinales de retraites, mesures de retraite de pseudo-cohortes, modèles de retraite par âge.

Introduction

Those born at the beginning of the post-World-War-II baby boom have now passed 65, the age long associated with retirement, and the retired proportion of the adult population will rise in the next two decades. Of particular policy concern is the economic well-being of future cohorts of retirees and, related to that, the anticipated increases in the tax burden of providing publicly financed

1 Byron G. Spencer, Department of Economics, McMaster University. Hamilton ON L8S 4M4. Email: spencer@mcmaster.ca. With Frank T. Denton, Department of Economics, McMaster University (Hamilton); Ross Finnie, Graduate School of Public and International Affairs, University of Ottawa.

healthcare services and pension benefits. The growth in the retired portion of the population will be especially large if gains in life expectancy continue without corresponding increases in the average age of retirement; if that happens, the fraction of adult life spent beyond the working years will continue to increase.

The possibility of changes in retirement age is thus a matter of much interest. At the societal level, later access to pension entitlements, such as the recently proposed increase from 65 to 67 in the age of eligibility for OAS and GIS benefits, would directly reduce the costs of those programs. Moreover, it would also encourage delays in retirement, and that, in turn, would mean a larger labour force and hence an increase in the productive potential of the economy.[2] So, too, would other measures that would encourage or facilitate the continued participation of older workers, such as those proposed by the Expert Panel on Older Workers (2008) and by Foot and Venne (2011), some of which are assessed by Denton and Spencer (2009b).

Policy initiatives matter, but given the policy environment, the ability to sustain a lifestyle will also have an important bearing on retirement decisions as individual circumstances change. Thus, the decline in the value of assets accumulated for retirement, and the increased uncertainty associated with future private pension benefits resulting from the widespread financial problems of recent years, have caused many to rethink and revise their retirement plans.

In short, retirement patterns, and potential changes in them, are important for both individuals and public policy. It would be helpful, therefore, to have a well defined and easily understood measure of retirement with which to assess current patterns and future trends; we should be able to state unambiguously whether people are retiring at younger ages now than they did a decade ago, or whether trends are in the other direction.

Many measures of retirement have been proposed, but no consensus has emerged in which one dominates the others. New measures are suggested from time to time, but they often draw on features that are specific to a single data source, thereby limiting their comparability with other measures. This may make comparisons with previous findings difficult, and render their on-going use problematic (see Denton and Spencer 2009a for a review). As one example, Statistics Canada produces age of retirement series based on responses to the *Labour Force Survey*: respondents who are not working and who left their last job within the year prior to being surveyed are asked why they left that job, and one response category is "retired."[3] However, this provides only a cross-sectional measure; as such it is not likely to represent the age pattern of retirement for any actual cohort, nor does it provide estimates of the proportion who retire at each age (the flow), or the proportion who have retired by each age (the stock).

Retirement, then, is a somewhat fuzzy concept. Nevertheless, it would be useful to have a single consistent, cohort-based measure that would make it possible to compare retirement patterns over time and between countries or regions. Such a measure could be the starting point for further analyses that would increase our understanding of the determinants of retirement and inform policy initiatives. However, our purpose here is somewhat less ambitious; we propose instead, and implement, two types of measures of the transition from work to retirement that apply to older popula-

2 Denton and Spencer (2011) have assessed the potential impact of delayed retirement (and additional years of work) on the productive capacity of the economy.

3 Further information is provided in Gower (1997) and Bowlby (2007). Respondents are in the survey for six consecutive months, and the age of retirement series is based only on responses in the first of the six months. Carriere and Galarneau (2011) have proposed an alternative measure: drawing on that same question, but using all rotation groups rather than just the first one.

tion cohorts. Either or both could be adapted for use in jurisdictions that have suitable longitudinal administrative records and on-going household labour force surveys.

Taking a cohort perspective is important, since retirement patterns tend to change over time and across cohorts. Recent examples of U.S. cohort studies include Cahill et al. (2006) and Coile and Gruber (2007), both of which used data from the Health and Retirement Study, albeit with differing indicators of retirement. In the present paper we draw on the large *Longitudinal Administrative Databank (LAD)* compiled from Canadian income tax records to derive earnings-based, cohort-specific measures of retirement from individual records.[4] For comparison, we draw also on time series of cross-sectional age-group data from the Canadian *Labour Force Survey* to construct consistent measures of labour force retirement for pseudo-cohorts—measures which, by definition, follow age groups rather than the same individuals over time. We conclude that while the *LAD* file provides a richer resource for understanding the changes that have taken place, the two sources do yield quite similar overall patterns of the transition from work to retirement. This finding is useful, since repeated cross-sections of data are more commonly available from survey sources, which facilitates comparisons across jurisdictions.

Two approaches to the measurement of retirement patterns

Pseudo-cohort measures based on cross-sectional household survey data

We confine our attention to the population 50 and over, and start by considering measures based on responses to Statistics Canada's *Labour Force Survey* (*LFS*). The *LFS* is a monthly sample survey of more than 50,000 households, designed to provide basic information about the labour force and its characteristics. It is generally similar to the *Current Population Survey* in the United States and to corresponding surveys in other OECD countries and elsewhere.

Drawing on the confidential LFS master files, we calculate annual average rates of labour force participation by single years of age, separately for each sex, from 1976 to 2006.[5] That provides us with the basis for estimating transition-to-retirement rates for successive pseudo-cohorts (or simply cohorts, where the context makes the meaning clear) and for making comparisons with annual cross-sectional or period rates.

We find marked differences between cohort and period patterns in some cases. As one example, the *period* age profile on the upper left side of Figure 1 shows the male participation profile as it was in 1976, from ages 52 to 72. (The rates have been indexed to 100.0 at age 52 to facilitate comparison with other measures.) The *cohort* profile shows how participation rates actually evolved for those who were 52 in 1976, 53 in 1977, and so on, up to 72 in 2006.

Since very few people enter the labour force after age 50, those who are in the labour force at that age constitute the population that we define as being "at risk of retirement". The proportion subsequently withdrawing from the labour force, and hence no longer active, can be interpreted as an approximation to the proportion that has retired; that is shown in the lower panel of the figure. Changes in the proportion show the average age pattern of withdrawal from the labour force, which we interpret as the transition to retirement. It is evident that this transition was much more gradual for the cohort than would have been predicted using the 1976 period rates. The difference is most

4 Other studies that have used that database to obtain cohort measures of retirement include Tompa (1999) and Wannell (2007), although their measures differ from one another and from what is proposed here.

5 The LFS master files were accessed in the Statistics Canada Research Data Centre at McMaster University.

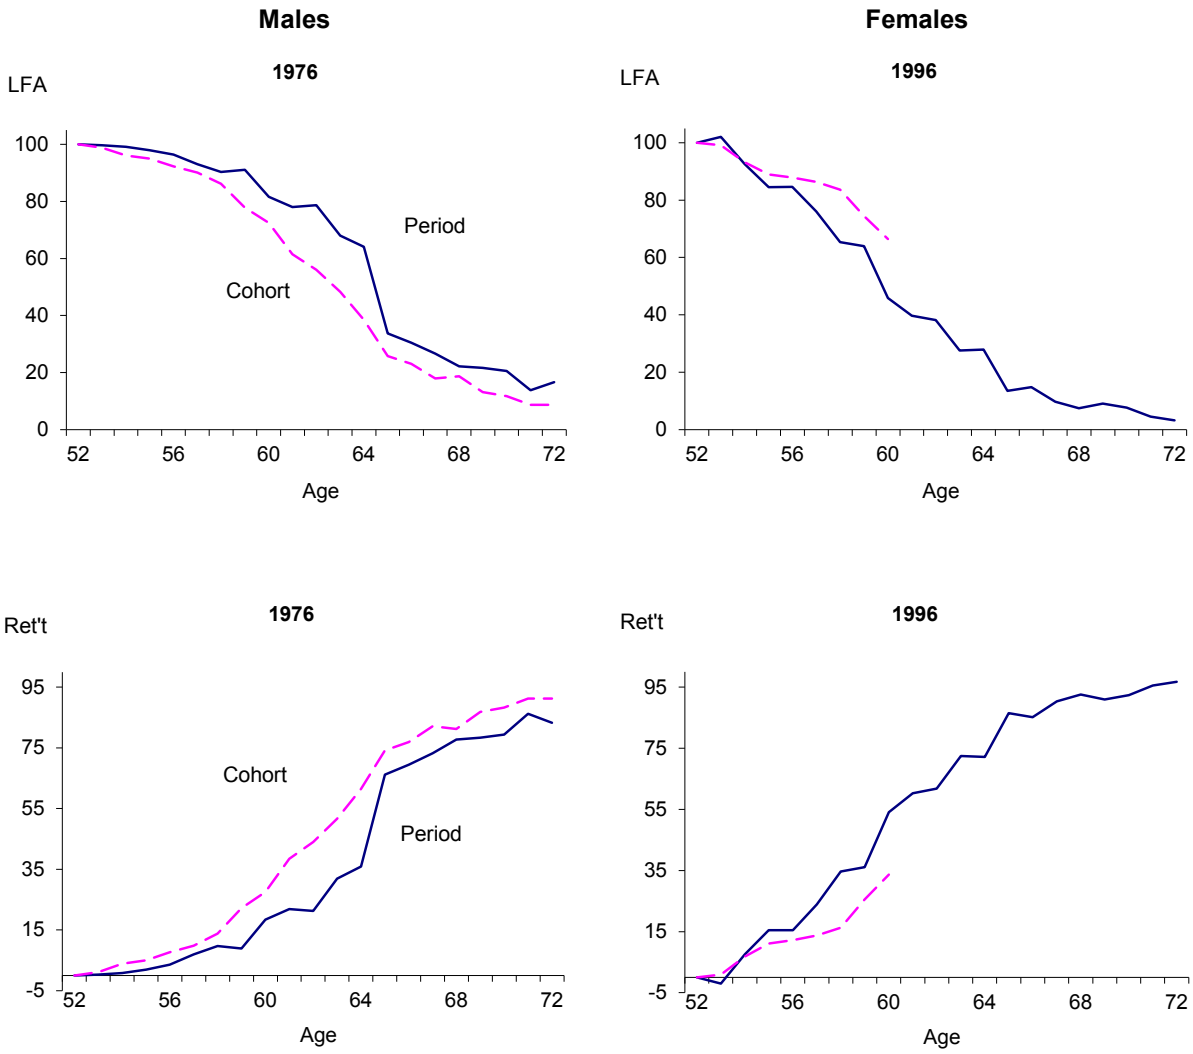

Figure 1. LFS period and cohort indexes of labour force activity (LFA; age 51 = 100) and retirement (Ret't), selected years and cohorts.

pronounced at age 62, where the cohort profile indicates that 40 per cent of those aged 50 in 1976 had actually retired while the period profile (based only on 1976 data) indicates only 20 per cent. Of particular interest, perhaps, is that there was no precipitous drop in participation, and hence no sharp increase in retirement, between the ages of 64 and 65. These are important differences, and point to the weakness in using period data to infer cohort patterns.

The right side of Figure 1 relates to females. The period rates in this case are from 1996, and we compare them to the rates for the cohort aged 52 in 1996. The observed experience for the 1996 cohort is much shorter, extending only up to age 62 in 2006, the last year of our data. The participation rates were much higher for this cohort when it was in its late 50s, and the retirement rates commensurately lower than would have been anticipated from the 1996 cross sectional profile. The difference is emphasized in the lower panel, which shows the estimated proportion retired. Such cohort/period differences again show the importance of basing retirement rates on cohort experience, and the errors that can arise when period rates are interpreted as if they applied to cohorts.

Cohort measures based on longitudinal administrative data

We work also with Statistics Canada's *Longitudinal Administrative Databank*. The *LAD* consists of a random 20-per cent sample of all taxpayers who filed Canadian income tax returns in any year, starting in 1982.[6] Information is added each year as new returns are filed, and the sample is augmented with 20 per cent of first-time tax filers. Individuals are included for all years in which they filed returns. By 2006, the most recent year for which we had data for this analysis, there were more than 4.9 million individuals in the sample.

The *LAD* has much to recommend it. Indeed, the very large sample size, its longitudinal nature, and the detailed and accurate information about income—in total and by source—that it provides year by year make it an appealing foundation for the analysis of earnings-based measures of cohort retirement, how patterns of retirement have changed over time for successive cohorts, and how they vary by level of earnings and such other individual and family characteristics as may be observable from income tax records.

Our approach to the choice of observations is as follows. We take the notion of retirement to be irrelevant before the age of 50. We first select all tax filers aged 50 with significant income from employment in 1982 and follow them until 2006, or until they died or were otherwise lost from the sample because they failed to file returns.[7] (The overall retention rate in the *LAD* is very high for the 1982 cohort: by the end of the data period, in 2006, we can account for 92.2 percent of all males who filed income tax returns in 1982, when they were aged 50, and 87.7 per cent of all females.) We then do the same for tax filers aged 50 in 1983, tax filers aged 50 in 1984, and so on, thus building up income histories for a series of successive cohorts, each identified by the year in which it reached the age of 50. (As one would expect, the retention rates are even higher for later cohorts than for the 1982 cohort.)

We exclude from the sample those few individuals who died or were lost before reaching age 52. We exclude also those who had any income from farming or fishing at ages 50, 51, or 52, since the notion of retirement is conspicuously vague for those occupations.[8] For each tax filer remaining in our observation set, average earnings from employment at ages 50 to 52 is then calculated as the arithmetic mean of the earnings at those three ages. In order to limit the analysis to individuals with significant labour market attachment, we exclude those for whom this average is less than $10,000 in constant

6 The *LAD* files are of course confidential. All calculations based on them were carried out at Statistics Canada. The following description is drawn largely from Statistics Canada's Longitudinal Administrative Data Dictionary (Cat. No. 12-585-XIE).

7 For this analysis, income information is imputed for those few (about 0.8 per cent of the sample) who failed to file income tax returns for either a single year or two years in a row, but then filed again. The imputation is based on a simple averaging of each component of the income information, including the total, as reported in the year preceding and the year following missing value(s). This is done to reduce possible sample selection bias related to occasional failure to file returns. Such imputation would be inappropriate if the typical reason for not filing was a much lower than average level of income in the affected year, but we have no way of assessing whether that was the case.

8 The exclusion of farmers and fishers from the analysis is somewhat arbitrary. However, it is intended to recognize that retirement is less well defined for self-employed persons in these occupations, since the point at which a self-employed farmer or fisher "retires" is more difficult to determine than in most other cases, especially when the relationship between gross and net income is somewhat problematic to start with. In fact, though, the exclusion has only a negligible effect on our results, so as a practical matter it can be ignored.

(2006) dollars.[9] That figure is arbitrary, but it may be thought of as representing about the amount that would be earned by someone working roughly half-time at a legislated minimum wage rate.

The next step is to identify those who have retired, as indicated by a major and sustained reduction in employment income. For each tax filer, the ratio of employment income at each subsequent age to average employment income at ages 50–52, denoted by R, is calculated for each year, for the maximum period permitted by the data. A tax filer is said to have retired at the age at which R first falls below a threshold level, R^*, provided that that condition continues to be satisfied in each of the subsequent two years.[10] We have experimented with several alternative values of R^*, ranging from 0.00 to 0.50 (Denton et al. 2011). Thus, at one extreme, a person would be deemed to have retired only if he or she had no income at all from employment ($R^*=0.00$); at the other, the same person would be classified as retired even if income from employment was just under half as great as its average level when he or she was 50–52 ($R^*=0.50$).

We note and emphasize that what we measure here is *first* retirement. It is possible that an individual may retire by our criterion, but subsequently may return to work. However, our criterion is rather demanding, inasmuch as earned income must remain below the threshold ratio for three successive years.

In total, we have 26 cohorts, each defined by the year in which their members reached the age of 50. In what follows, we focus attention on the 1982, 1987, 1992, and 1997 cohorts, whose transitions to retirement we can follow to ages 72, 67, 62, and 57, respectively. These cohorts are generally representative of those near them.

Results

Figure 2 compares the retirement index based on the *Labour Force Survey* for the 1982 cohort with corresponding indexes based on the *LAD* for four values of R^*. It is evident from the figure that the *LFS* and *LAD* age profiles are very similar, for both males and females, and especially for low values of R^*. That is as one might expect, since a respondent to the *LFS* with even a very low level of earnings would be classified as being in the labour force and hence not retired. Thus, we might anticipate that not being in the labour force and being retired would correspond fairly closely to the case of $R^*=0.00$, as we see it does. The *LFS* series relates to pseudo-cohorts, while the *LAD* series relates to true cohorts (the same individuals followed over time). That the two series are in such close agreement provides evidence to support the use of pseudo-cohorts derived from times series of cross-sectional data to study retirement patterns when longitudinal data are not available. Both approaches also demonstrate the importance of taking an explicit cohort approach rather than assuming that the period rates will continue to apply, since they often give quite misleading indications of cohort patterns.

With $R^*=0.00$, both series indicate that about 40 per cent of males who had significant labour force attachment when they were in their early 50s had retired by the age of 61, and almost 80 per cent by the age of 66. For females, the proportions retired at each age are somewhat higher, which-

9 All income measures are adjusted for inflation using the consumer price index, and expressed in dollars of 2006. Employment income includes net income from self-employment.

10 Note that this calculation tells us the age reached during the first full year of retirement, not age at the exact date of retirement within a year. A tax filer would be deemed to be retired at the youngest age x at which the specified condition is satisfied. By way of example, a person would be deemed to have retired at 63 if the retirement condition is satisfied at each of ages 63, 64, and 65. In addition, a person would be deemed to have retired at age 63 if the condition is satisfied at age 63 and the person is dead or lost from the sample at age 64, or if it is satisfied at ages 63 and 64 and the person is dead or lost at age 65.

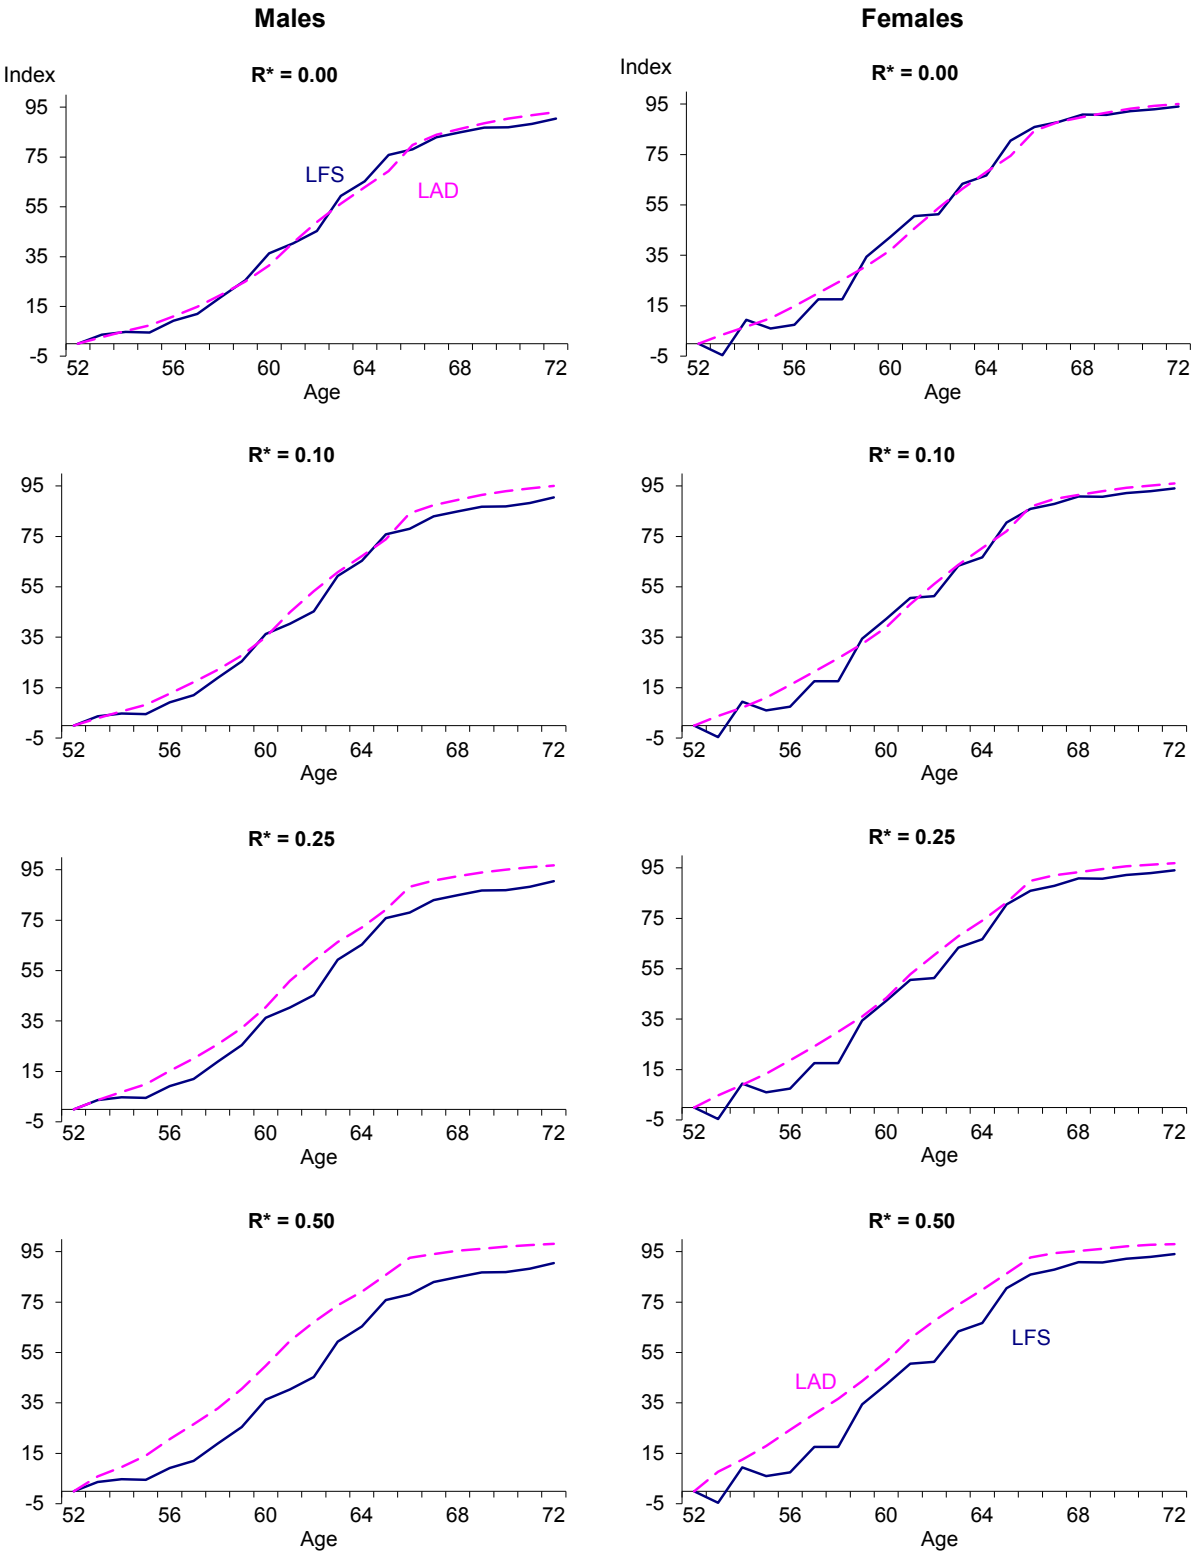

Figure 2. Comparison of LFS and LAD retirement indexes, 1982 cohorts, alternative R values.*

ever series is used, but the age pattern of retirement is generally similar. Higher levels of R^* mean that higher earnings are consistent with being classified as retired, based on the *LAD*. (For example, if $R^*=0.50$, a person would be counted as retired even with earnings that were almost half as great as they were at ages 50–52, in real terms.) It is not surprising, then, that the age retirement profile based on higher values of R^* lies above the one based on the *LFS*.

Figure 3 provides similar comparisons for the 1987, 1992, and 1997 cohorts, but restricted to the case of $R^*=0.00$. The age range for which comparisons can be made is necessarily reduced by the length of the period covered by the data, but the general point remains: the two measures provide very similar indications of the transitions from work to retirement.

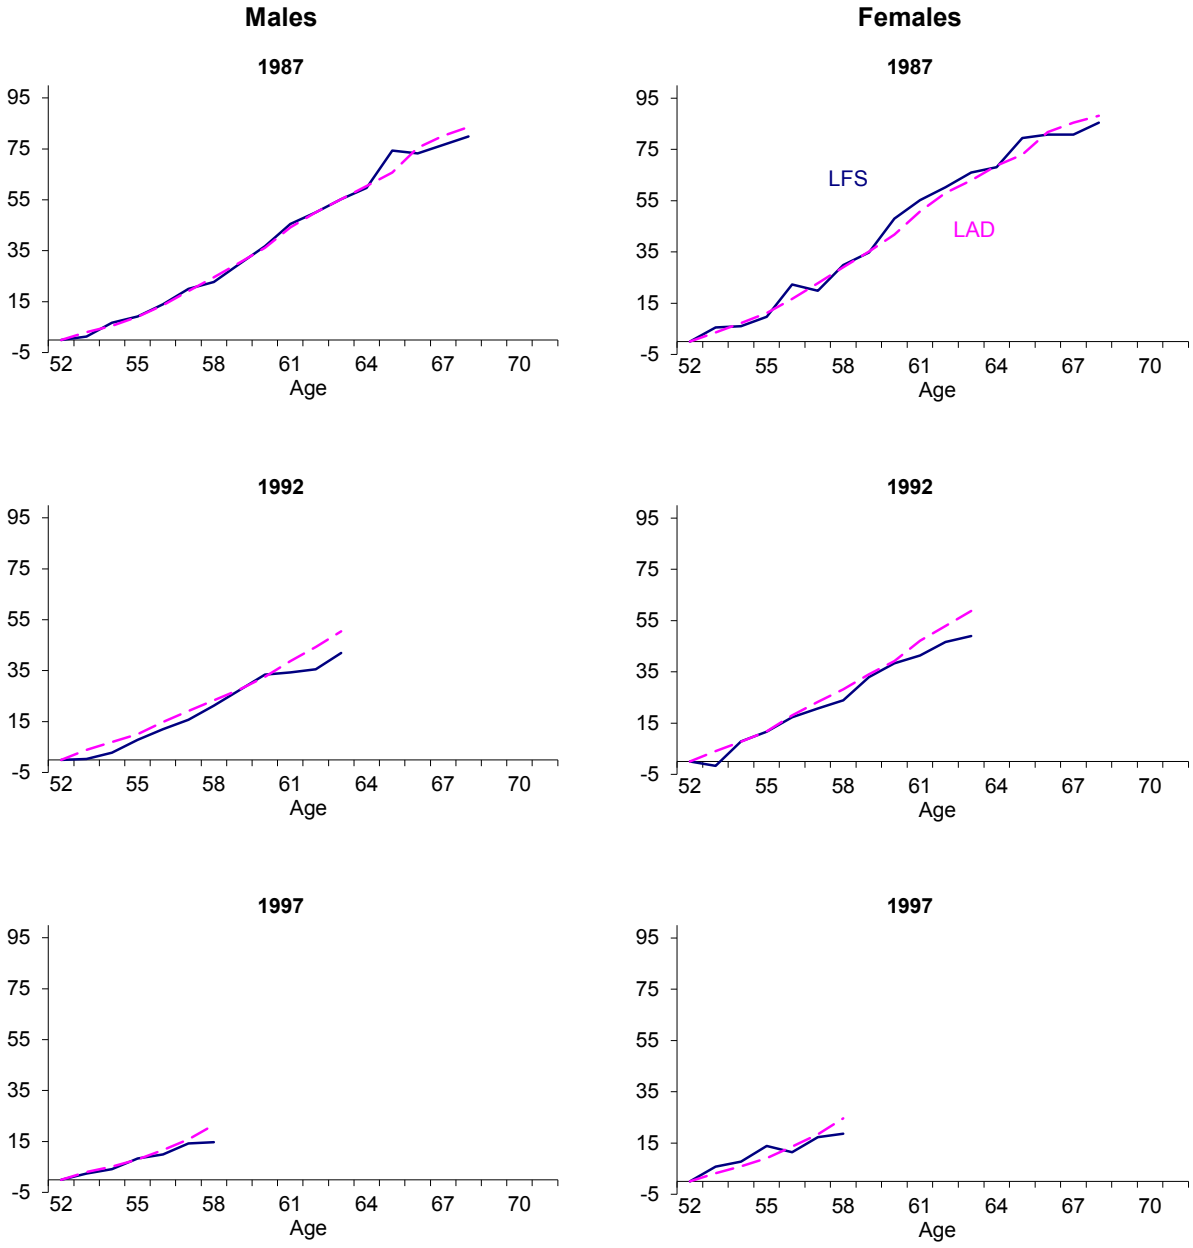

Figure 3. Comparison of LFS and LAD retirement indexes, selected cohorts, R = 0.00.*

Concluding remarks

Any assessment of changes over time in the age pattern of retirement must be based on an agreed concept of retirement, and how it is measured. The same comment applies to comparisons across countries. We have demonstrated that measures based on the realized experience of actual cohorts can be much more meaningful than measures based on data for only one period, and the assumption that individuals continue to retire at the ages implied by the cross-section information of that period. While the proportion retired increases with age, the age pattern can differ from one cohort to another, and we have shown that these differences can be substantial. When that happens, the actual age-transition path for a cohort will differ from what is suggested by the work/retirement pattern of any one period, and it is the cohort experience that is of greater interest.

We have proposed two measures of retirement. The first, a *pseudo-cohort* measure, draws on a time series of annual cross-sectional surveys; retirement is defined by withdrawal from the labour force and measured by the age pattern of reduction in average labour force participation rates. The second, a *true cohort* measure, draws on administrative income tax records, in which the same individuals are tracked over time; retirement is defined, in that case, at the individual level by a sufficiently large and sustained reduction in employment income, and aggregate rates are obtained by summation.

Using Canadian data, the two measures are found to provide quite similar results when the "sufficiently large" reduction in employment income in the second measure is close to 100 per cent. Furthermore, the two measures bear striking similarities, even when the second one counts as retired those whose real (constant dollar) income from employment is sustained at levels up to 25 or even 50 per cent of what it was when they were in their early 50s.

That is useful information: it tells us that plausible longitudinal measures of the age pattern of retirement can be obtained in a timely way, and need not be based on access to the much more costly full longitudinal tracking of individuals, with the inevitable time lags that are involved. Without doing the research, we could not have known that the two approaches would produce such similar estimates of age profiles. We find it reassuring to be able to derive meaningful and very similar retirement profiles from two such dissimilar data files. We note also that the close similarity of the two measures suggests that retirement as we measure it using the *LAD* is very largely voluntary. That is because measures based on the LFS would count any older worker who became unemployed, but who preferred not to retire, as still in the labour force, provided that he/she was looking for work. If continuing to seek employment is the norm for those who would prefer not to retire, the fact that the two measures are close would suggest that our *LAD*-based measure includes little involuntary retirement.

Even though the pseudo-cohort approach has merit, there are some important advantages associated with the longitudinal earnings-based approach. They derive from the fact that actual cohorts of individuals are being followed, a feature that makes it possible to gain insights that are not feasible when looking at pseudo-cohorts. By way of example, when working with longitudinal data such as the tax records used here, it is possible to assess relative post-retirement well-being (by comparing income from all sources before and after retirement), and to ask how that may vary with the age at which individuals retire, with their position in the income distribution prior to retirement, with respect to other individual or family characteristics (see, for example, Denton et al. 2009), or with differences in the state of the economy. Furthermore, individual characteristics and experience might be related to movements into retirement, and could be analysed, taking into

account relevant factors such as individual preparation for retirement (e.g., income levels and savings behaviour in the pre-retirement period)—something that is not possible when working with the *LFS* pseudo-cohorts.

Another advantage of using longitudinal data is that the measure of retirement itself is much less rigid, i.e., with the pseudo-cohort approach, based on the *LFS*, individuals are deemed to be still "in the labour force" if they are either working for any amount of pay or looking for work. With longitudinal administrative records, it is possible to adopt a more subtle measure of retirement, and to compare retirement rates based on alternative criteria, such as level of earnings is deemed to be consistent with retired status.

We conclude that while we cannot be sure that our findings would apply elsewhere, in Canada, at least, overall retirement rates may be estimated quite accurately using pseudo-cohorts derived from cross-sectional labour force survey data. Longitudinal tax data can allow for more extensive analysis of retirement patterns and factors that affect them, but to realize such benefits one must have access to a large database that follows individuals over time. In the absence of such a longitudinal base, our results suggest that reliable estimates of at least the main features of cohort retirement patterns can be based on the pseudo-cohort approach, using labour force data.

Acknowledgements

The authors thank Human Resources and Skills Development Canada (HRSDC) for its financial support, including the arrangements for access to the *LAD* files at Statistics Canada. They are grateful also to SSHRC for its financial support of the SEDAP (Social and Economic Dimensions of an Aging Population) Research Program through its Major Collaborative Research Initiative. We thank Yan Zhang who worked with the *LAD* files at Statistics Canada to provide the basic tabulations, Christine Feaver who produced the final figures and commented on the interpretation of results, and two anonymous referees who made a number of helpful suggestions.

References

Bowlby, G. 2007. Defining retirement. *Perspectives on Labour and Income* 8(2):15–19. (Ottawa: Statistics Canada, Cat. No. 75-001-XIE).

Cahill, K.E., M.D. Giandrea, and J.F. Quinn. 2006. Retirement patterns from career employment. *The Gerontologist* 40(4):514–523.

Carriere, Y., and D. Galarneau. 2011. Delayed retirement: A new trend? *Perspectives on Labour and Income* 23(4):3–16. (Ottawa: Statistics Canada, Cat. No. 75-001-X).

Coile, C., and J. Gruber. 2007. Future social security entitlements and the retirement decision. *Review of Economics and Statistics* 89(2):234–246.

Denton, F.T., and B.G. Spencer. 2011. Age of pension eligibility, gains in life expectancy, and social policy. *Canadian Public Policy* 37(2):183–200.

———. 2009a. What is retirement? A review and assessment of alternative concepts and measures. *Canadian Journal on Aging* 28(1):63–76.

———. 2009b. Population aging, older workers, and Canada's labour force. *Canadian Public Policy* 35(4):481–492.

Denton, F.T., R. Finnie, and B.G. Spencer. 2011. Income replacement in retirement: Longitudinal evidence from income tax records. *Canadian Journal on Aging* 30(3):1–17.

———. 2009. *Patterns of retirement as reflected in income tax records for older workers*. Research Report No. 434. Hamilton: Research Institute for Quantitative Studies in Economics and Population, McMaster University.

Expert Panel on Older Workers. 2008. *Supporting and Engaging Older Workers in the New Economy*. Ottawa: Human Resources and Skills Development Canada. Accessed via http://www.hrsdc.gc.ca/eng/publications_resources/lmp/eow/2008/page19.shtml (retrieved May 1, 2012).

Foot, D.K., and R.A. Venne. 2011. The long goodbye: Age, demographics, and flexibility in retirement. *Canadian Studies in Population* 38(3–4):59–74.

Gower, D. 1997. Measuring the age of retirement. *Perspectives on Labour and Income* 9(2):11–17. (Ottawa: Statistics Canada, Cat. No. 75-001-XPE).

Tompa, E. 1999. *Transitions to Retirement: Determinants of Age of Social Security Take Up*. SEDAP Research Report No. 6. Accessed via http://socserv.socsci.mcmaster.ca/sedap/p/sedap6.PDF (retrieved May 5, 2012).

Wannell, T. 2007. Young pensioners. *Perspectives on Labour and Income* 8(2). (Ottawa: Statistics Canada, Cat. No. 75-001-XIE).