

Population Aging, Labor Demand, and the Structure of Wages

Michael Papadopoulos, Margarita Patria, and Robert K. Triest

Abstract:

One consequence of demographic change is substantial shifts in the age distribution of the working-age population. As the baby boom generation ages, the usual historical pattern of a high ratio of younger workers relative to older workers has been replaced by a pattern of roughly equal percentages of workers of different ages. One might expect that the increasing relative supply of older workers would lower the wage premium paid for older, more experienced workers.

This paper provides strong empirical support for this hypothesis. Econometric estimates imply that the size of one's birth cohort affects wages throughout one's working life, with members of relatively large cohorts (at all stages of their careers) earning a significantly lower wage than members of smaller cohorts. Estimated elasticities of wages with respect to the relative size of one's own cohort are generally between -0.05 and -0.10, and are of similar magnitude for men and for women. Our results suggest that cohort size effects are quantitatively important and should be incorporated into public policy analyses.

JEL Classifications: J11, J21, J26

Keywords: labor market demographics, population aging, wage structure, social insurance

Michael Papadopoulos is a Ph. D. student in the economics department at the New School for Social Research. His email address is papam427@newschool.edu. Margarita Patria, who holds a Ph. D. in economics from Boston College, is a principal at Charles River Associates. Her email address is mpatria@crai.com. Robert K. Triest is a vice president and economist in the research department at the Federal Reserve Bank of Boston. His email address is robert.triest@bos.frb.org

The research reported herein received support from The Atlantic Philanthropies and from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium (RRC). The authors thank Jamie Lee and Brendan Mackoff for very helpful research assistance, Suzanne Lorant for expert editing, and Melinda Pitts, Ron Lee, Alicia Munnell, Joe Quinn, Steve Sass, Tony Webb, and two anonymous referees for helpful comments.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/wp/index.htm>.

The findings and conclusions expressed are solely those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Boston, the Federal Reserve System, Charles River Associates, The Atlantic Philanthropies, the SSA, any agency of the federal government, or the RRC.

This version: May 31, 2017

1 Introduction

Along with virtually all other developed countries, the United States has experienced a radical transformation of its labor markets due to a profound demographic shift. As many have documented, the growth rate of the American working-age population has already dropped substantially, with corresponding decreases in the growth rates of labor supply and potential output. As the large baby boom generation continues to age, the elderly dependency ratio (the ratio of those over 65 to the population aged 15 to 65) will increase dramatically. As a consequence, labor supply may grow at a slower rate than labor demand, possibly putting upward pressure on wages and creating tight labor-market conditions. Often overlooked, however, is the fact that the age composition of the labor force has also changed dramatically. The traditional pattern of a greater number of young workers than of older workers has been replaced by a situation where the age distribution of the labor force is fairly flat.

The effect these changes may have on the labor-market opportunities of older workers is not immediately obvious. Although labor-demand conditions may be favorable for workers in general, the large size of the baby boom cohort compared with younger cohorts may place the baby boomers at a relative disadvantage. The same crowding effect that depressed the boomers' wages when they were young (Welch (1979) and others) may continue to haunt them into their 60s. Although there may be upward pressure on wages

in general, the relative glut of older workers may depress their wages relative to those of their younger colleagues.

What happens to the wages of older workers, and the structure of wages more generally, as the population ages has potentially important implications for public policy. Many analysts are convinced that longer working lives must be a key component of any solution to providing for the consumption needs of the old as the traditionally defined dependency ratio increases. The efficacy of this solution depends, in part, on the wage rates that older workers command in labor markets. If the wages of older workers fall as their ranks become crowded with the baby boomers, then continued work may seem a less desirable option to those contemplating retirement, and the earnings of those who do continue working will not go as far in financing their consumption.

How the wage structure changes as the baby boomers age also has potential implications for forecasting future payroll tax revenue and Social Security benefits. To the extent that the boomers' wages have been depressed due to cohort size effects, their exit from the labor market may affect aggregate earnings growth. More generally, the earnings trajectories of those currently in the middle of their careers, as well as those younger workers just starting out, will likely be affected by changes in the age distribution of the population. And changes in earnings trajectories will, of course, result in changes in payroll tax payments and eventual Social Security benefits.

This paper investigates empirically the effects of changes in the age distribution of the working-age population on the structure of wages. In particular,

we examine the hypothesis that cohort crowding not only affects the wages of large birth cohorts as they enter the labor market but continues to exert downward pressure on the wages of large cohorts as they approach retirement age. We find strong support for this hypothesis. The size of an individual's birth cohort affects that individual's wages throughout his or her working life, with members of relatively large cohorts at all stages of their careers earning a significantly lower wage than members of smaller cohorts. Our results suggest that cohort size effects are quantitatively important and should be incorporated into public policy analyses.

2 Previous Research

There is a sizable research literature examining how changes in the age distribution of the labor force affect the structure of wages. The unifying idea underlying this literature is that workers with different amounts of labor-market experience are imperfect substitutes in production. Workers acquire human capital through on-the-job training and through learning-by-doing; More experienced workers will generally perform somewhat different tasks than those performed by younger workers, and compared with younger workers, will tend to play different roles within a firm's organization of production. As the supply of labor with a given level of experience increases, the wages of workers in that group will tend to decrease relative to the wages of workers with different experience levels. The smaller the degree of substitutability

between workers with different experience levels, the greater the change in relative wages that will result from a given change in relative supplies. Variation over time in the relative supplies of workers at given levels of experience is essential for estimating the degree of substitutability; therefore, most of the studies on this topic are based on examining how the relative wages of men with differing levels of labor-market experience changed as the baby boom generation entered the labor market. Building on this work, our study uses more recent data than were available to previous researchers and also analyzes changes in the wages of women as well as those of men. By using wage data extending through 2015, we observe the effects of the oldest baby boomer birth cohort as its members move through the bulk of their careers, up to age 69. The added variation in relative cohort sizes associated with the recent data is very useful in identifying empirically the effects of changes in the age distribution, and allows us to find direct evidence of the impact that cohort crowding has had upon the baby boomers' wages as they approach retirement.

In a remarkably prescient analysis written well before the first baby boomers started entering the labor force, Easterlin (1961) notes that the labor-market fortunes of workers are inversely related to the relative size of their birth cohort. Easterlin anticipated that as the baby boomer generation entered the labor market, they would face less favorable conditions than the cohort that preceded them.¹

¹Easterlin's main focus was on how economic conditions affect fertility, and he correctly

An early, and very influential, econometric examination of the baby boom’s effect on relative wages is Welch’s [1979] famous study of “The Baby Boom’s Financial Bust.” Using data from the March income and demographic supplements to the Current Population Survey (CPS) from 1968 to 1976, he finds that the wages of young white men were reduced relative to those of older men as the baby boomers started entering the labor market. Noting that the range of potential substitution possibilities is too large to be investigated without some structure, Welch imposes the restriction that substitution between workers with different degrees of educational attainment is independent of their experience levels. Welch estimates the effect of own cohort size on wages allowing for an interaction between cohort size and labor-market experience. The resulting econometric estimates suggest that the relative wage reductions associated with being a member of a large cohort are concentrated in the early years of workers’ careers. A concurrent study by Freeman (1979) reaches a similar conclusion — relative wages of young workers were depressed due to cohort crowding effects. Freeman finds that the effect of the baby boom generation’s entry into the labor market on the premium paid to older workers was especially large for college-educated men.

Berger (1985) generally follows Welch’s [1979] methodology, but uses additional years of data and a somewhat less restrictive econometric specification. Like Welch, Berger finds that entry-level wages are reduced by cohort

predicted that the relatively unfavorable conditions created by the entry of the baby boom-cohort into the labor market would depress fertility rates.

size, but unlike Welch's, his estimates indicate that the cohort size effect grows with labor-market experience. Using a factor analytic technique to decompose the composition of the labor force into a small number of factors, Murphy, Plant, and Welch (1988) find a pattern similar to that found by Welch (1979) — that the depression of wages associated with being in a large cohort is concentrated early in one's career.

Katz and Murphy (1992) examine the role of cohort size in explaining changes in experience differentials as part of a larger framework exploring how shifts in both labor demand and labor supply affect the structure of wages and conclude that although the increasing wage differentials associated with labor-market experience in the 1970s and 1980s is consistent with the trend of an increasing share of young workers in the labor force, the exact timing of the changes do not match up well. Murphy and Welch (1992) also find that supply shifts alone cannot fully explain changes in the experience differential.

3 Empirical Patterns

Following most previous research on this topic, we use data from the annual income and demographic supplement to the March Current Population Survey (CPS). Unlike previous researchers, who observed data for a more limited span of time, we use data for the years 1964–2016. The March supplement survey collects income information for the preceding year, so our wage data


span 1963 through 2015.

3.1 Changes in the Age Distribution

Figure 1 shows that striking changes have occurred in the age distribution of the working-age population (here defined as ages 18 through 65) over the past 40 years. Each panel of the figure shows the frequency distribution of the working age population for a given year. A growing population is associated with a downward-sloping line, while a stable population produces a horizontal line (with each annual birth cohort making up roughly 2.1 percent of the working-age population). Barely discernible in 1966, the emergence of the younger baby boomers into their working years is very apparent in the graph for 1976, where young adults greatly outnumber middle-aged and older workers. In 1986, when the youngest baby boomers turned 20, one can see the start of a hump-shaped distribution forming, as the post-boom “baby bust” generation started to enter their working years. The hump moves to the right between 1986 and 1996, producing an unusual situation in which middle-aged workers outnumber those in both older and younger cohorts.

The 2006 age distribution looks somewhat similar to the 1966 distribution. However, unlike in 1966, when the baby boomers were about to enter the labor force, the working-age population distribution increasingly approximated a uniform distribution in subsequent years. The 2016 panel shows a population that is fairly evenly distributed over all age ranges, with only a modest downward tilt associated with people in their 50s and 60s. The days

Figure 1: Changes in the Distribution of the Working-Age Population over Time



of a large ratio of older to younger workers seem to be over for the foreseeable future.


3.2 Age and Labor-Market Experience

In order to understand the implications of the changing age distribution for the structure of wages, we need to analyze how the relative supplies of workers with differing levels of educational attainment and labor-market experience have changed. Our analysis classifies individuals into five educational attainment categories: those who did not complete high school, high-school graduates, some college (one to three years completed), four-year college graduates, and those with post-college graduate education. We aggregate the CPS data into groups defined by this classification of educational attainment and years of labor-market experience.

Researchers using the CPS have typically used age minus years of schooling as a measure of potential labor-market experience. However, particularly for women, actual labor-market experience is likely to be significantly less than potential labor-market experience, with the difference between the two measures strongly dependent on birth cohort. Unfortunately, the CPS does not include information on actual work experience, so we need to impute this information. In constructing the experience variable, we follow a nonparametric cohort “splicing” approach similar to that used by Herd (2005) in a study of the effect of a minimum Social Security benefit on benefit adequacy for women. We interpolate population microdata from the decennial census

to form synthetic labor-experience histories by gender and educational attainment for birth cohorts from 1900, and then impute the resulting measure of labor-market experience to our CPS-based observations.²


Figure 2: The Evolution of Labor-Market Experience for Women



The importance of using actual rather than potential labor-market experience is illustrated by Figures 2 and 3. Over time, there has been a

²Details are provided in the appendix.

Figure 3: The Evolution of Labor-Market Experience for Men



sharp increase in the average years of labor-market experience for women at any given age for all levels of educational attainment. However, the rate of change deaccelerated sharply over time, and the lines for the 1960 and 1970 birth cohorts are sometimes difficult to distinguish in the figure.³ The greatest increases in labor-market experience have been for older women, reflecting the cumulative effect of increases in labor-force participation at all ages. For all birth cohorts, average labor-market experience at any given age increases with educational attainment. Over time, the smallest increases in labor-market experience have been at the extremes of the distribution of educational attainment. Women who fail to complete high school accumulate relatively little labor-market experience, and women who extend their education beyond college on average accumulate a great deal of experience, but in both cases there have been only relatively minor changes over time.

In contrast to the pattern for women, the relationship between men's age and average labor-market experience shows little change over time, and the lines for the four birth cohorts shown in the figure are difficult to distinguish. The main qualification to this statement is that for less-educated men there has been some reduction in the accumulation of labor-market experience.


Overall, Figures 2 and 3 show that while potential labor-market experience tends to greatly overstate actual labor-market experience for pre-baby boom women, this phenomenon is not quantitatively important for men. Our econometric estimates showing that key labor-demand parameters are very

³For this reason, we do not show birth cohorts more recent than 1970.

similar for men and women depend critically on making the labor-market experience adjustment for women.


3.3 Changes in the Distribution of Labor-Market Experience

Figure 4: Changes in the Distribution of Labor-Market Experience for Men



Following Welch (1979), researchers have generally assumed that substitution possibilities between workers with different experience levels are greater

Figure 5: Changes in the Distribution of Labor-Market Experience for Women



within educational attainment groups than between groups. Pursuing this assumption, changes in the distribution of labor-market experience within the same educational attainment groups are especially relevant for analyzing which changes in the relative labor supplies are most likely to affect relative wages. Figures 4 and 5 are similar to Figures 2 and 3, but show the frequency distribution of labor-market experience separately for men and women at two different levels of educational attainment, high-school graduate and college graduate.⁴ The patterns in Figures 4 and 5 differ from those in Figure 1 primarily because of changes in the average levels of educational attainment over time.⁵ As a result, the impact of the baby boom on the age distribution will differ across educational groups. The entry of the baby boom generation into the labor force had a larger initial impact on the distribution of college-educated workers than it did on high-school graduates — the oldest baby boom cohort was not only much larger in overall numbers than were earlier birth cohorts, but its members were also much more likely to attend and complete college. The relative cohort size of the oldest baby boom college graduates then decreased over time as the pre-baby boom cohorts were replaced by the younger, even-more-highly educated, baby boomers. In recent years, the experience distributions of the high-school graduates and college graduates have converged, and in the future the combined distribution will

⁴The lines shown in these figures are smoothed versions of the underlying data. A locally weighted scatterplot smoothing (LOWESS) procedure was used, with a bandwidth of 0.2.

⁵For women, increases in labor-force participation over time also play an important role.

increasingly resemble a uniform distribution.


3.4 Changes in the Experience Premium

It is evident that cohort size has a large impact on the wage rates of older workers relative to younger workers, as shown in top panel of Figure 6, which charts from 1963 to 2015 the median wage rates of full-time, full-year male workers with 22 to 25 years of labor-market experience relative to the median wage rates of those with two to five years of experience.⁶ The experience premium is shown only for men in this panel because average experience for women is consistently less than 22 years for all education-birth-year groups. Our wage measure is based on individual average hourly earnings, which is annual wage and salary income divided by total hours worked. Total hours worked per year is computed by taking the product of weeks worked the previous year and usual hours worked per week. The median of individual hourly wages within education-experience groups for each year is used as the group wage measure.⁷ For all levels of educational attainment up to post-college, there is a very clear pattern of the experience premium first rising and then falling with the aging of the baby boomers. The post-college group is an exception to this, and may reflect the relatively low post-college degree

⁶We define full-time, full-year workers as those who report working at least 45 weeks in the previous year and that they normally work at least 35 hours per week. We use the CPI-W series to express nominal wage rates in 2004 dollars. As described in the appendix, experience is imputed for each gender-education-age-birth-year group. A LOWESS procedure, with bandwidth equal to 0.2, was used to smooth the plotted data.

⁷The median, rather than the mean, of individual wages is primarily used to lessen the impact of outliers.

Figure 6: Changes Over Time in the Experience Premium



completion rate among baby boomers, along with technical change that has increased demand for workers with advanced degrees.

The bottom panel of Figure 6 shows changes in the the premium paid for workers with 10 to 13 years of labor-market experience relative to those with two to five years of experience. By focusing on the return to this more limited quantity of labor-market experience, we are able to display the premium for both male and female workers in this panel.⁸ However, the difference in the relative supplies of the two experience ranges shown in the lower panel is generally much smaller than the corresponding difference shown in the top panel, and as a result it is difficult to discern the effect of changing relative supplies on the experience premium in the lower panel. A comparison of the upper and lower panels suggests that the exaggeration of the cross-sectional experience premium caused by the entry of the baby boomers was much more pronounced at relatively senior experience levels.

4 Econometric Specification

The patterns shown in Figures 4, 5, and 6 provide strong evidence that the distribution of labor-market experience within educational groups is one determinant of the wage differential between more experienced workers and less experienced workers. This section presents an econometric specification

⁸We are not able to display the experience premium for the "less-than-high-school" women for early years due to the insufficient average work experience of this group in those years.

that allows us to estimate this relationship more formally.

Following previous researchers, we specify an aggregate production function treating workers with differing degrees of educational attainment as imperfect substitutes.⁹ Within each educational group, workers with differing levels of labor-market experience are imperfect substitutes. Our formal specification is:

$$Y_t = \left(\sum_j \theta_j E_{jt}^\rho \right)^{1/\rho}, \quad (1)$$

where Y_t is aggregate output in year t , E_{jt} is the number of workers with educational attainment j used in production in year t , $\rho = 1 - \frac{1}{\sigma_E}$, and σ_E is the elasticity of substitution between workers with differing educational attainments. Each E_j quantity is, in turn, a C.E.S. aggregator over workers with differing degrees of labor-market experience:

$$E_j = \left(\sum_k \alpha_k E_{jk}^\eta \right)^{1/\eta}, \quad (2)$$

where E_{jk} is the number of workers with educational attainment j and with k years of labor-market experience (the time subscript, t , is dropped here), $\eta = 1 - \frac{1}{\sigma_A}$, and σ_A is the elasticity of substitution between workers with differing years of labor-market experience. The wage of a worker in educational group g who has h years of labor-market experience is then:

⁹Our specific specification most closely follows that of Card and Lemieux (2001).

$$w_{gh} = \frac{\partial Y}{\partial E_{gh}} = \theta_g \alpha_h \left(\frac{E_{gh}}{E_g}\right)^{\eta-1} E_g^{\rho-1} \left(\sum_j \theta_j E_j^{\rho_j}\right)^{\left(\frac{\rho-1}{\rho}\right)}. \quad (3)$$

Taking logarithms and rearranging yields an equation for the log wage of workers in educational group g with h years of experience that is linear in the log of the supply of labor with h years of experience relative to the total supply of labor with educational attainment g :

$$\ln(w_{gh}) = \ln(\theta_g) + \ln(\alpha_h) + (\eta-1)\ln\left(\frac{E_{gh}}{E_g}\right) + (\rho-1)\ln E_g + \left(\frac{\rho-1}{\rho}\right)\ln\left(\sum_j \theta_j E_j^{\rho_j}\right). \quad (4)$$

We use this equation as the basis for our estimated labor-demand relationship. The data we use for estimation is organized such that each observation is specific to a given gender-educational attainment-birth cohort combination.¹⁰ We specify that α_h , the main experience productivity effect, follows a piecewise linear spline in years of labor-market experience, with nodes at three, six, nine, and 15 years of experience.¹¹ This approach is more flexible than the quadratic specification often adopted.

We interact the relative cohort size term, $\ln\left(\frac{E_{gh}}{E_g}\right)$, with a set of indicator variables corresponding to the five segments of the experience spline in order to investigate whether the relative cohort size effect changes as cohort mem-

¹⁰Each observation is formed by aggregating the underlying CPS micro data by gender, birth year, and educational attainment.

¹¹An earlier version of this paper that used data only for men included nodes at 20 and 30 years of experience, but the distribution of female labor-force experience made it necessary to set our highest node at 15 years.

bers gain experience, a point of contention in the earlier research literature on the effect of the baby boom generation’s entry into the labor market. One difficulty in estimating the C.E.S. specification is that although the inputs in the model are defined in terms of years of labor-market experience, our observations are defined by year of birth. As a consequence, the inputs in the empirical implementation of the model are delineated by potential rather than actual years of labor-market experience, although the terms in the α_h spline are measured as mean years of actual labor-market experience.¹²

Goldin (1992) emphasizes that cohort-specific effects have played an important role in women’s expectations and attitudes toward their careers. The major social changes that occurred over much of the past century have resulted in qualitative changes in socialization and in the labor-market opportunities young women expect to be open to them. And for both men and women, over time there have been improvements in the human capital acquired prior to beginning one’s career. To control for these effects, we allow for $\ln(\theta_g)$ to vary linearly with year of birth. Because these birth cohort effects are allowed to vary by educational attainment, they also capture changes over time in the relative demand for workers by educational attainment.

We also allow for a time-trend spline, with kink points at five-year intervals, to allow for technical change (changes in the θ_g terms) and for changes

¹²That is, although the observations are formed by aggregating by birth year and educational attainment, and so effectively by potential experience, the value of actual labor-market experience that is imputed for each observation is used in the regressions.

over time in aggregate labor supplies $((\rho - 1)\ln E_g + (\frac{\rho-1}{\rho})\ln(\sum_j \theta_j E_j^\rho))$. Although we think that technical change is the main phenomenon driving the estimated time-trend spline coefficients, the time-trend spline can also capture other time-varying phenomena not explicitly included in the model. A potentially important factor that may be reflected in the spline coefficients is the effect of the business cycle on wages. Changes in corporate human resources policies and practices during our sample period might also be captured by the spline, although some of the changing practices may be due to demographic change and be captured by the relative cohort size terms. For example, a shift in human resource policies away from promoting long-term careers within firms may reflect the relative glut of experienced baby boomers within firms, rather than an exogenous force.

Finally we assume that $\ln(\theta_g)$ incorporates a linear stochastic term. Our reported standard errors and test statistics are robust to the possibility that this term is correlated across observations in a given year, to allow for macroeconomic influences on wages, and are also robust to the presence of heteroskedasticity.

5 Regression Results

The dependent variable for all of the regressions is the natural log of the median real wage of full-time, full-year workers within cells defined by single years of potential labor-market experience, the five educational attainment

groups defined above, and single calendar years. The same March CPS data from 1964 through 2016 that were used in the figures are also used here.

Falaris and Peters (1992) and Connelly and Gottschalk (1995) provide evidence that educational investment decisions are affected by demographically induced changes in relative wages, so it is likely incorrect to treat relative cohort size within educational attainment groups as exogenous. In addition, sampling error may be a significant factor in measured changes in relative cohort size over time. To address these problems, we use relative cohort size, defined over all educational groups with the same birth year, as an instrument for relative cohort size defined within educational groups. Overall relative cohort size is very likely to be exogenous in this context, and has smaller sampling variation due to the larger number of observations used in its estimation.

Table 1 presents instrumental variables regression results for the combined sample of men and women. An indicator variable for “female” is included in the specification, and the birth year effect is also allowed to vary by gender. The subsample used for estimation is limited to workers between the ages of 22 and 62 in order to reduce the influence of selection out of the labor force due to schooling or retirement. The measure of relative cohort size used in these regressions is gender specific: it is the ratio of full-time workers with a given gender, level of educational attainment, and birth year to all full-time workers with the same gender and level of educational attainment.¹³

¹³To smooth over sampling variation, cohort size is calculated as a five-year centered

Table 1: Estimation Results : Instrumental Variable Regressions

Variable	Less Than High School		High-School Graduate		Some College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.141**	(0.046)	-0.044*	(0.019)	-0.065**	(0.011)
Cohort Size, 3-6 years	-0.142**	(0.045)	-0.046*	(0.020)	-0.060**	(0.011)
Cohort Size, 6-9 years	-0.144**	(0.045)	-0.045*	(0.019)	-0.063**	(0.011)
Cohort Size, 9-15 years	-0.143**	(0.046)	-0.048*	(0.019)	-0.068**	(0.011)
Cohort Size, 15+ years	-0.147**	(0.045)	-0.054**	(0.020)	-0.070**	(0.012)
Female	-1.089**	(0.110)	-0.375**	(0.033)	-0.293**	(0.034)
Year of Birth (Male)	-0.005**	(0.001)	0.009**	(0.001)	0.013**	(0.001)
Year of Birth (Female)	-0.003**	(0.001)	0.011**	(0.001)	0.014**	(0.001)
Experience Spline, 0-3 years	0.093**	(0.017)	0.098**	(0.006)	0.145**	(0.005)
Experience Spline, 3-6 years	0.029**	(0.005)	0.061**	(0.005)	0.063**	(0.004)
Experience Spline, 6-9 years	0.019**	(0.005)	0.042**	(0.003)	0.043**	(0.003)
Experience Spline, 9-15 years	0.007**	(0.002)	0.032**	(0.002)	0.037**	(0.002)
Experience Spline, 15+ years	-0.007**	(0.002)	0.009**	(0.001)	0.011**	(0.001)
Time Spline, 1964-1969	0.022**	(0.004)	0.003	(0.003)	0.003	(0.004)
Time Spline, 1970-1974	0.011*	(0.005)	-0.002	(0.004)	-0.009*	(0.005)
Time Spline, 1975-1979	-0.013**	(0.004)	-0.023**	(0.005)	-0.037**	(0.007)
Time Spline, 1980-1984	-0.004	(0.003)	-0.013**	(0.005)	-0.009	(0.005)
Time Spline, 1985-1989	-0.001	(0.003)	-0.011**	(0.003)	-0.013**	(0.003)
Time Spline, 1990-1994	-0.215**	(0.003)	-0.029**	(0.003)	-0.031**	(0.004)
Time Spline, 1995-1999	-0.010	(0.007)	-0.020**	(0.008)	-0.025**	(0.009)
Time Spline, 2000-2004	0.046*	(0.022)	0.036	(0.023)	0.037	(0.025)
Time Spline, 2005-2009	0.060**	(0.019)	0.063**	(0.021)	0.066**	(0.023)
Time Spline, 2010+	-0.009	(0.007)	-0.026**	(0.008)	-0.035**	(0.008)
Intercept	1.841**	(0.309)	1.415**	(0.151)	1.133**	(0.087)
N		4346		4346		4346
R ²		0.849		0.915		0.894
F _(23,52)		1453.37		2329.95		3927.65

Significance levels : † : 10% * : 5% ** : 1%

Table 1: Estimation Results : Instrumental Variable Regressions (2)

Variable	College Graduate		Post College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.057**	(0.006)	-0.039**	(0.005)
Cohort Size, 3-6 years	-0.053**	(0.007)	-0.033**	(0.007)
Cohort Size, 6-9 years	-0.051**	(0.007)	-0.031**	(0.008)
Cohort Size, 9-15 years	-0.054**	(0.007)	-0.036**	(0.009)
Cohort Size, 15+ years	-0.057**	(0.008)	-0.040**	(0.010)
Female	-0.325**	(0.037)	-0.089**	(0.047)
Year of Birth (Male)	0.014**	(0.001)	0.019**	(0.002)
Year of Birth (Female)	0.015**	(0.001)	0.018**	(0.002)
Experience Spline, 0-3 years	0.131**	(0.005)	0.123**	(0.008)
Experience Spline, 3-6 years	0.065**	(0.006)	0.088**	(0.010)
Experience Spline, 6-9 years	0.055**	(0.007)	0.050**	(0.006)
Experience Spline, 9-15 years	0.036**	(0.002)	0.034**	(0.003)
Experience Spline, 15+ years	0.013**	(0.001)	0.021**	(0.002)
Time Spline, 1964-1969	0.014**	(0.006)	0.002	(0.005)
Time Spline, 1970-1974	-0.013**	(0.006)	-0.015*	(0.006)
Time Spline, 1975-1979	-0.046**	(0.008)	-0.053**	(0.006)
Time Spline, 1980-1984	-0.011*	(0.006)	-0.010†	(0.006)
Time Spline, 1985-1989	-0.005**	(0.004)	-0.011*	(0.005)
Time Spline, 1990-1994	-0.023**	(0.004)	-0.011*	(0.005)
Time Spline, 1995-1999	-0.024**	(0.008)	-0.020*	(0.009)
Time Spline, 2000-2004	0.043†	(0.026)	0.042	(0.027)
Time Spline, 2005-2009	0.072**	(0.023)	0.077**	(0.025)
Time Spline, 2010+	-0.031**	(0.008)	-0.040**	(0.009)
Intercept	1.303**	(0.095)	1.170**	(0.124)
N		4344		4296
R ²		0.889		0.824
F _(19,52)		975.49		909.16

Significance levels : † : 10% * : 5% ** : 1%

The relative cohort size coefficients can be interpreted as $\eta - 1$ in the context of the C.E.S. specification or more simply as the elasticity of wages with respect to the relative size of one's own cohort. The estimated elasticities are uniformly negative, which confirms that belonging to a relatively large cohort is associated with depressed wages. These elasticities are also sizable in magnitude, generally hovering around -0.05 for high-school graduates and for college graduates. For all of the educational groups, the relative cohort size effect varies relatively little with years of labor-market experience, implying that relative cohort size is roughly as important to the wages earned late in one's career as earlier on in one's work life.

This sheds light on a point of contention in the earlier literature, where studies were mixed on whether the relative cohort size effect is greater in the early or later stages of careers. Although the magnitude of the relative cohort size coefficients remain roughly constant over one's career, the effect of cohort size on one's earnings will vary over one's career for any given birth cohort and will also vary across birth cohorts. The reason for this is that relative cohort size changes with birth cohort and with years of labor-market experience (as well as educational attainment), as shown in Figures 4 and 5. So, the magnitude of the effect of relative cohort size on wages may be greater either early in one's career or later in one's career, depending on how relative cohort size evolves for one's birth cohort and educational attainment group.

moving average, with weights equal to $1/9, 2/9, 1/3, 2/9,$ and $1/9$.

The labor-market experience spline coefficients generally imply that although real wages increase rapidly with labor-market experience, there is a sharp drop in the growth rate of earnings as this experience increases. Real wage rates tend to increase very slowly with experience beyond 15 years of experience.¹⁴

In interpreting the relative cohort size coefficients, it is important to remember that they are capturing the effects of cohort size on monetary compensation and that they omit the effect on fringe benefits. A massive switch from defined-benefit pensions, in which workers' pension accruals are typically concentrated in the years when they are approaching retirement, to defined-contribution pensions occurred starting in the 1980s (Munnell and Sundén (2004)). As a result, inclusion of the value of employer-paid pension accruals in our measure of compensation would likely accentuate the estimated effect of relative cohort size on compensation.

Tables 2 and 3 are similar to Table 1, but display results for the regression model estimated separately for men and women. For women who did not complete high school, the highest experience group (15 years or more) was eliminated, as there were no observations in this experience range.

Surprisingly, the relative cohort size effects for women are similar to those for men shown in Table 1, with elasticities of roughly the same order of magnitude for both men and women. Men's and women's wages are depressed

¹⁴We do not include knots in the spline at points beyond 15 years because the experience is measured as within-group averages and generally does not extend much beyond 15 years for women.

Table 2: Estimation Results : Instrumental Variable Regressions, Male Only

Variable	Less Than High School		High-School Graduate		Some College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.056 [†]	(0.031)	-0.046 [†]	(0.024)	-0.062**	(0.017)
Cohort Size, 3-6 years	-0.068*	(0.030)	-0.051*	(0.024)	-0.054**	(0.016)
Cohort Size, 6-9 years	-0.068*	(0.028)	-0.050*	(0.023)	-0.052**	(0.016)
Cohort Size, 9-15 years	-0.068*	(0.029)	-0.051*	(0.023)	-0.056**	(0.016)
Cohort Size, 15+ years	-0.073*	(0.029)	-0.060*	(0.024)	-0.062**	(0.017)
Experience Spline, 0-3 years	0.065**	(0.022)	0.055**	(0.014)	0.168**	(0.010)
Experience Spline, 3-6 years	0.056**	(0.010)	0.044**	(0.007)	0.078**	(0.007)
Experience Spline, 6-9 years	0.041**	(0.010)	0.014*	(0.005)	0.050**	(0.006)
Experience Spline, 9-15 years	0.031**	(0.005)	0.002	(0.005)	0.037**	(0.006)
Experience Spline, 15+ years	0.018**	(0.005)	-0.019**	(0.004)	0.019**	(0.006)
Time Spline, 1964-1969	0.003	(0.005)	0.033**	(0.004)	-0.007	(0.005)
Time Spline, 1970-1974	-0.001	(0.005)	0.021**	(0.004)	-0.009	(0.006)
Time Spline, 1975-1979	-0.032**	(0.005)	-0.001	(0.006)	-0.044**	(0.007)
Time Spline, 1980-1984	-0.023**	(0.005)	0.006	(0.005)	-0.014*	(0.006)
Time Spline, 1985-1989	-0.010*	(0.004)	-0.004	(0.004)	-0.021**	(0.006)
Time Spline, 1990-1994	-0.040**	(0.004)	-0.010*	(0.005)	-0.039**	(0.006)
Time Spline, 1995-1999	-0.021*	(0.008)	0.002	(0.007)	-0.026**	(0.009)
Time Spline, 2000-2004	0.032	(0.021)	0.056*	(0.024)	0.033	(0.027)
Time Spline, 2005-2009	0.052**	(0.019)	0.085**	(0.021)	0.063**	(0.024)
Time Spline, 2010+	-0.021*	(0.09)	-0.007	(0.008)	-0.037**	(0.009)
Year of Birth	0.011**	(0.007)	-0.012**	(0.003)	0.018**	(0.004)
Intercept	1.159**	(0.402)	2.918**	(0.303)	0.789*	(0.313)
N		2173		2173		2173
R ²		0.773		0.850		0.866
F (21,52)		221.11		383.54		308.05

Significance levels : † : 10% * : 5% ** : 1%

Table 2: Estimation Results : Instrumental Variable Regressions, Male Only (2)

Variable	College Graduate		Post College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.057**	(0.008)	-0.043**	(0.009)
Cohort Size, 3-6 years	-0.057**	(0.009)	-0.048**	(0.012)
Cohort Size, 6-9 years	-0.065**	(0.011)	-0.052**	(0.015)
Cohort Size, 9-15 years	-0.076**	(0.012)	-0.058**	(0.017)
Cohort Size, 15+ years	-0.084**	(0.012)	-0.065**	(0.017)
Experience Spline, 0-3 years	0.141**	(0.013)	0.097**	(0.022)
Experience Spline, 3-6 years	0.060**	(0.008)	0.051**	(0.017)
Experience Spline, 6-9 years	0.041**	(0.010)	0.046**	(0.017)
Experience Spline, 9-15 years	0.036**	(0.008)	0.017	(0.017)
Experience Spline, 15+ years	0.012	(0.008)	-0.001	(0.018)
Time Spline, 1964-1969	0.006	(0.009)	0.014	(0.014)
Time Spline, 1970-1974	-0.006	(0.008)	0.010	(0.020)
Time Spline, 1975-1979	-0.048**	(0.010)	-0.044**	(0.016)
Time Spline, 1980-1984	-0.009	(0.007)	0.011	(0.018)
Time Spline, 1985-1989	-0.015†	(0.008)	-0.001	(0.015)
Time Spline, 1990-1994	-0.028**	(0.007)	0.005	(0.017)
Time Spline, 1995-1999	-0.023*	(0.011)	-0.003	(0.018)
Time Spline, 2000-2004	0.044†	(0.027)	0.059*	(0.029)
Time Spline, 2005-2009	0.073**	(0.024)	0.091**	(0.031)
Time Spline, 2010+	-0.034**	(0.010)	-0.028	(0.018)
Year of Birth	0.015*	(0.007)	0.005	(0.015)
Intercept	1.169*	(0.474)	2.104*	(1.047)
N		2173		2146
R ²		0.881		0.818
F (21,52)		511.53		328.04

Significance levels : † : 10% * : 5% ** : 1%

Table 3: Estimation Results : Instrumental Variable Regressions, females only

Variable	Less Than High School		High-School Graduate		Some College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.002	(0.042)	-0.083**	(0.024)	-0.087**	(0.016)
Cohort Size, 3-6 years	-0.002	(0.042)	-0.085**	(0.024)	-0.084**	(0.016)
Cohort Size, 6-9 years	-0.001	(0.042)	-0.085**	(0.024)	-0.088**	(0.016)
Cohort Size, 9-15 years	-0.005	(0.041)	-0.086**	(0.024)	-0.091**	(0.016)
Cohort Size, 15+ years			-0.082**	(0.024)	-0.083**	(0.015)
Experience Spline, 0-3 years	0.038	(0.020)	0.102**	(0.007)	0.138**	(0.006)
Experience Spline, 3-6 years	-0.001	(0.10)	0.066**	(0.009)	0.068**	(0.007)
Experience Spline, 6-9 years	0.014	(0.010)	0.053**	(0.006)	0.052**	(0.005)
Experience Spline, 9-15 years	-0.010	(0.022)	0.039**	(0.004)	0.046**	(0.004)
Experience Spline, 15+ years			0.047**	(0.008)	0.038**	(0.006)
Time Spline, 1964-1969	0.028**	(0.006)	-0.007*	(0.004)	0.005	(0.006)
Time Spline, 1970-1974	0.011†	(0.006)	-0.007	(0.005)	-0.017*	(0.007)
Time Spline, 1975-1979	-0.004	(0.006)	-0.026**	(0.007)	-0.037*	(0.008)
Time Spline, 1980-1984	0.003	(0.004)	-0.014*	(0.006)	-0.010	(0.007)
Time Spline, 1985-1989	-0.001	(0.004)	-0.010*	(0.005)	-0.013**	(0.004)
Time Spline, 1990-1994	-0.009*	(0.004)	-0.031**	(0.005)	-0.033**	(0.004)
Time Spline, 1995-1999	-0.006	(0.007)	-0.026**	(0.009)	-0.032**	(0.009)
Time Spline, 2000-2004	0.051*	(0.022)	0.034	(0.023)	0.033	(0.025)
Time Spline, 2005-2009	0.059**	(0.020)	0.053**	(0.020)	0.059**	(0.022)
Time Spline, 2010+	-0.006	(0.008)	-0.030**	(0.008)	-0.039**	(0.008)
Year of Birth	-0.002	(0.002)	0.014**	(0.002)	0.017**	(0.001)
Intercept	1.922**	(0.309)	0.704	(0.239)	0.563**	(0.156)
N		2173		2173		2173
R ²		0.707		0.871		0.838
F _(19,52)		209.23		328.93		254.68

Significance levels : † : 10% * : 5% ** : 1%

Table 3: Estimation Results : Instrumental Variable Regressions, females only (2)

Variable	College Graduate		Post College	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Cohort Size, 0-3 years	-0.056**	(0.007)	-0.030**	(0.008)
Cohort Size, 3-6 years	-0.049**	(0.008)	-0.016	(0.011)
Cohort Size, 6-9 years	-0.043**	(0.008)	-0.009	(0.012)
Cohort Size, 9-15 years	-0.042**	(0.009)	-0.011	(0.014)
Cohort Size, 15+ years	-0.032**	(0.010)	-0.011	(0.014)
Experience Spline, 0-3 years	0.115**	(0.005)	0.103**	(0.012)
Experience Spline, 3-6 years	0.059**	(0.008)	0.074**	(0.008)
Experience Spline, 6-9 years	0.049**	(0.009)	0.020**	(0.005)
Experience Spline, 9-15 years	0.026**	(0.002)	0.011*	(0.005)
Experience Spline, 15+ years	0.029**	(0.004)	0.012*	(0.006)
Time Spline, 1964-1969	0.026**	(0.006)	0.019**	(0.006)
Time Spline, 1970-1974	-0.016**	(0.006)	-0.008	(0.007)
Time Spline, 1975-1979	-0.041**	(0.008)	-0.033**	(0.007)
Time Spline, 1980-1984	-0.010†	(0.006)	0.000	(0.007)
Time Spline, 1985-1989	0.010**	(0.004)	0.011	(0.007)
Time Spline, 1990-1994	-0.013**	(0.004)	0.003	(0.007)
Time Spline, 1995-1999	-0.019*	(0.008)	-0.007	(0.010)
Time Spline, 2000-2004	0.047†	(0.026)	0.052†	(0.028)
Time Spline, 2005-2009	0.070**	(0.023)	0.088**	(0.024)
Time Spline, 2010+	-0.027**	(0.008)	-0.025**	(0.009)
Year of Birth	0.011**	(0.001)	0.006*	(0.002)
Intercept	1.272**	(0.095)	1.942**	(0.170)
N		2171		2150
R ²		0.845		0.790
F _(19,52)		419.26		363.86

Significance levels : † : 10% * : 5% ** : 1%

by cohort-crowding effects by similar magnitudes throughout their careers. Although the elasticity estimates are qualitatively similar, Wald tests of the hypothesis that the relative cohort size coefficients are the same for men and women easily reject that hypothesis (at the 0.0001 significance level or less for each educational attainment level).¹⁵


5.1 Changes in Relative Cohort Size Over Time

It is not immediately obvious from the regression results how changes in the distribution of labor-market experience affect the life-cycle wage profile of a given birth year cohort because the relative size of a given birth cohort changes over time. This fact is illustrated in Figure 7, which shows relative cohort size over time for four birth cohorts: those born in 1940, 1950, 1960, and 1970.

Looking first at the data for college graduates shown in the bottom panels of the figure, one sees that the baby boomers born in 1950 comprised an exceptionally large fraction of the college-educated labor force when they first entered the labor market, but the cohort's relative size decreased over time as even larger cohorts from the middle years of the baby boom subsequently entered their working years. Those born relatively late in the baby boom, in 1960, were a smaller fraction of the labor force when they first entered

¹⁵The test was conducted by estimating instrumental variables regressions in which the gender indicator variable is interacted with the relative cohort size and experience spline variables. The hypothesis that the coefficients on both the relative cohort size and age spline variables are equal for both genders was also rejected at the 0.0000 significance level for each educational attainment level.

Figure 7: The Evolution of Relative Cohort Size



the labor market than the early baby boomers were at the same stage of their careers. As the entire baby boom generation matured and increasingly made up the bulk of the college-educated work force, the relative size of any given baby boom birth year cohort shrank. This is reflected in the gradual convergence of the lines for the 1950 and 1960 birth cohorts in Figure 7. Note that the lines for these two baby boom cohort years are always well above those for both the pre-baby boom 1940 birth cohort and the post-boom 1970 cohort.

The patterns are somewhat different for high-school graduates. Because of changes in the distribution of educational attainment over time, the 1960 birth cohort was a larger fraction than the 1950 birth cohort of the high-school-educated labor force at all levels of labor-market experience. Unlike the case of college-educated men, where the early baby boomers had exceptionally large relative cohort sizes, the later baby boom birth cohorts were a larger fraction than the early baby boom birth cohorts of the high-school-educated labor force at all levels of labor-market experience.

The cohort size effects have interesting implications for how one interprets the relationship between wages and labor-market experience. The regression coefficients for the labor-market experience spline reflect what the wage-experience profile would be for a birth cohort that had a constant relative size (within education groups). In a growing population, a given cohort's relative size will shrink with age. If the relative cohort size coefficients were negative and constant over experience levels, this would result in any given cohort's

wage-experience profile being steeper than the experience spline coefficients indicate. A decrease in the population growth rate would eventually flatten the wage-experience profile, producing cohort wage-experience profiles closer to those implied by the experience spline coefficients. At a given point in time, the cross-sectional wage-experience profile will reflect the pattern of relative cohort sizes experienced by the birth cohorts in the labor force at that time and will generally differ from both any given cohort's wage-experience profile and from the constant relative cohort size wage-experience profile implied by the experience spline coefficients.

6 Policy Implications

Our findings have potentially important implications for public policy. The decrease in birth rates at the end of the baby boom and ongoing increases in life expectancy have combined to create funding problems for retirement income systems. Social Security is projected to deplete its trust funds in 2034 (The Board of Trustees, Federal Old-age and Survivors Insurance and Federal Disability Insurance Trust Funds (2016)), traditional defined benefit pension plans cover a dwindling share of new retirees, and many older workers appear to have savings inadequate to supplement Social Security and pension benefits to maintain their standard of living in retirement. Postponing retirement to later ages is sometimes proposed as a solution both to the Social Security funding problem (through increasing the age at which workers qualify for full

benefits) and to the problem of savings inadequacy. In essence, postponing retirement to a later age shrinks the length of time over which savings must be stretched and potentially also allows the build-up of a greater stock of savings before retirement. The logic of this argument is sound, but the findings of this paper suggest that older workers will likely experience diminished labor-market prospects (relative to earlier generations) as they attempt to postpone retirement.

Our finding that cohort crowding within educational attainment groups depresses wages at all levels of labor-market experience implies that boomers have faced a wage penalty, relative to the counterfactual situation of their being members of a smaller birth cohort, at all stages of their careers. This will likely continue as boomers age into their 60s and 70s and means that delaying retirement will be less lucrative for boomers than it would have been if they had been members of a relatively smaller generation. This attenuates, although certainly does not eliminate, the efficacy of delaying retirement as a strategy to bolster savings and enhance retirement income security. Although boomers are disadvantaged relative to earlier cohorts in this regard, future cohorts will face largely the same situation. Very slow population and labor-force growth imply that the ratio of older workers to younger workers has permanently increased. Like the boomers, future cohorts will find that there is a relative abundance of older and experienced workers as they approach retirement.

The mechanism through which cohort crowding has decreased the eco-

economic return to labor-market experience is important. Our estimates are derived through a model of labor demand in which workers with different levels of experience, within educational attainment groups, are imperfect substitutes in production. Workers are paid the value of their marginal product, which depends both on the direct effect of accumulated experience on productivity and on the the size of the worker's group relative to other groups with which the worker is substitutable. Our estimates indicate that the effect of labor-market experience on wages (and productivity, under the assumptions underlying our model) accumulates at a decreasing rate as workers gain experience, although this pattern becomes less pronounced as educational attainment increases. Generations prior to the baby boom experienced a more sharply upward-sloping wage-experience profile than that generated by the direct effect of experience on productivity, because in an economy with a high rate of labor-force growth, experienced workers were in short supply relative to entry-level workers (and so the relative cohort size effect augmented the direct experience effect in generating the total cross-sectional return to labor-market experience). Compared with earlier generations when there were relatively few experienced workers to serve as managers and mentors to entry-level workers, the new normal is for there to be more experienced workers than would have traditionally been needed for such roles.

Although our model specification assumes that workers are paid the value of their marginal product, in the real world it is likely that social norms, legal constraints, morale factors, and other frictions result in wages deviating

from marginal product. For older workers, this might induce impediments to employment if traditional wage norms associated with upward-sloping wage-experience profiles cause wages to exceed marginal productivity as workers age. Decreased employment stability, including semi-forced retirement and difficulty in securing new employment after job displacement, might be a consequence of this phenomenon.

This suggests a need for public policy to address obstacles to earnings and employment stability of older workers. Demographic change has simultaneously increased the desirability of workers prolonging their working lives and also, by making experienced workers abundant relative to entry-level workers, made it more difficult and less lucrative for experienced workers to maintain their labor earnings. Social insurance policy might usefully pivot from focusing mainly on income replacement after full retirement, as is the case with Social Security, to also addressing the problem of diminished or unstable earnings prior to retirement.

7 Conclusion

The age distribution of the American working age population has flattened and now approaches a uniform distribution. The historical pattern of a relatively large number of inexperienced young people working alongside a relatively small number of more experienced older workers has been replaced by a labor force where older and younger workers are roughly equal in number.

One consequence of this change in the relative supplies of older and younger workers is to reduce the experience premium that older workers command in the labor market.

Large birth cohorts depress their own wages relative to those of other cohorts in the labor force at the same time, with reductions roughly equal in magnitude for men and women. The increase in the cross-sectional labor-market experience premium induced by the baby boom generation's entry into the labor market has been offset by a decrease in the cross-sectional experience premium as the baby boom progresses through middle age and approaches retirement.

These results imply that older workers may continue to face relatively unfavorable labor-market conditions. Although the slowing of labor-force growth may create tight labor markets, the pecuniary benefits of labor-market tightness will disproportionately accrue to younger, less experienced workers. Loss of defined benefit pensions and increases in Social Security's normal retirement age may result in baby boomers retiring at older ages than the birth cohorts that immediately preceded them, but the boomers will suffer from the same cohort crowding as they consider retirement as they did earlier in their careers.

References

- Berger, Mark C. 1985. "The Effect of Cohort Size on Earnings Growth: A Reexamination of the Evidence." *Journal of Political Economy* 93 (3):561–573.
- Card, David and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *Quarterly Journal of Economics* 116 (2):705–746.
- Connelly, Rachel and Peter Gottschalk. 1995. "The Effect of Cohort Composition on Human Capital Accumulation across Generations." *Journal of Labor Economics* 13 (1):155–176.
- Easterlin, Richard A. 1961. "The American Baby Boom in Historical Perspective." *American Economic Review* 51 (5):869–911.
- Falaris, Evangelos M. and H. Elizabeth Peters. 1992. "Schooling Choices and Demographic Cycles." *Journal of Human Resources* 27 (4):551–574.
- Freeman, Richard B. 1979. "The Effect of Demographic Factors on Age-Earnings Profiles." *Journal of Human Resources* 14 (3):289–318.
- Goldin, Claudia. 1992. *Understanding the Gender Gap: An Economic History of American Women*. Oxford University Press.
- Herd, Pamela. 2005. "Ensuring a Minimum: Social Security Reform and Women." *The Gerontologist* 45 (1):12–25.

- Katz, Lawrence F. and Kevin Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1):35–78.
- Munnell, Alicia H. and Annika Sundén. 2004. *Coming Up Short: The Challenge of 401 (k) Plans*. Brookings Institution Press.
- Murphy, Kevin, Mark Plant, and Finis Welch. 1988. "Cohort Size and Earnings in the USA." In *Economics of Changing Age Distributions in Developed Countries*. 39–58.
- Murphy, Kevin and Finis Welch. 1992. "The Structure of Wages." *Quarterly Journal of Economics* 107 (1):285–325.
- The Board of Trustees, Federal Old-age and Survivors Insurance and Federal Disability Insurance Trust Funds. 2016. *The 2016 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds*. United States Government Printing Office.
- Welch, Finis. 1979. "Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust." *Journal of Political Economy* 87 (5):S65–S97.

A Imputing Labor-Market Experience

Past labor-market experience is not reported in the CPS, so we imputed average years of full-time labor-market experience based on synthetic labor-participation histories that we constructed by gender and educational attainment for each single-year birth cohort.

The synthetic labor-participation histories are based on decennial population census micro-data (IPUMS) samples for the years 1940, 1950, and 1960, along with March CPS data for 1964–2016. Census data prior to 1940 do not include information on educational attainment comparable to that which we use in this study, and so could not be used. For the years available (1940, 1950, 1960, and 1964–2016) we calculate the mean full-time employment ratio for each year/gender/age/education group. That is, for each cell in the year/gender/age/education matrix we find the percentage of people working full-time (which we define as working 45 or more weeks per year and 35 or more hours of work in the previous week; in 1960, we treat 40 or more weeks per year as full time due to data limitations). Because we lack data for years prior to 1940, we assume that full-time participation rates were constant from 1900 to 1940. We use linear interpolation to impute mean full-time participation rates for years between the decennial censuses and between 1960 and 1964. The final step in the imputation is to create a running sum of the full-time participation rates for each birth-year cohort (by gender and educational attainment). This yields a measure of mean years of full-time

labor-market experience for each birth-year cohort/age/gender/educational attainment combination.