Changing Levels or Changing Slopes?

The Narrowing of the U.S. Gender Earnings Gap, 1959-1999

Catherine Weinberger and Peter Kuhn
Department of Economics
University of California, Santa Barbara
Santa Barbara, CA 93106
weinberg@isber.ucsb.edu
pjkuhn@econ.ucsb.edu

February 2008

The gender wage gap among adult full-time workers, after controlling for educational attainment and other observable characteristics, is about half the size it was in 1980. Using Census and Current Population Survey data from 1959 through 1999, we assess the relative contributions of two factors to the decline in the gender wage gap: changes across cohorts in the relative slopes of men's and women's age-earnings profiles, versus changes in relative earnings levels at labor market entry. We find that changes in relative slopes account for about one-third of the narrowing of the gender wage gap over the past 40 years. Under fairly general conditions, we argue that this provides an upper-bound estimate of the contribution of all post-school investments, including experience, to the decline of the gender wage gap.

This material is based upon work supported by the National Science Foundation under Grant No. 0120111. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We thank Solomon Polachek for extensive comments on an earlier draft of this paper, and Francine Blau for a helpful discussion. This paper also benefited from the comments of Deborah Cobb-Clark, Heidi Hartmann, Larry Kahn, Kevin Lang, Paul Latreille, and from participants in presentations at the Society of Labor Economists' Annual Meetings, the National Science Foundation, the National Bureau of Economic Research, and at the Universities of British Columbia, Washington, and Western Ontario.

1. Introduction

The gender wage gap among adult full-time workers, after controlling for educational attainment and other observable characteristics, is about half the size it was in 1980 (e.g. Blau and Kahn 2006). A well known explanation of this fact is that women today experience fewer career interruptions and take more actions--such as career-enhancing geographic mobility--to preserve and enhance their earnings capacity (e.g. Goldin 1989; O'Neill and Polachek 1993). In this paper, we argue that changes across cohorts in the degree to which women make such post-schooling investments (PSIs) in their human capital after labor market entry should be reflected in the rate at which women's earnings "fall behind" men's over the life cycle. Thus, a change across cohorts in the relative slopes of women's (versus men's) age-earnings profiles is a central, testable feature of this "PSI hypothesis." More broadly, understanding the relative contribution of changing entry wages ('levels') versus changes in age-related wage growth ('slopes') to the reduction of the U.S. gender wage gap should shed some light on the processes underlying that reduction.

This paper draws on the above insight to study the decline in the U.S. gender wage gap since 1959. Our approach complements existing work on this topic in two main ways. First, existing panel-based studies of the evolving gap tend to proceed by constructing the best PSI measures they can, for example by accumulating detailed work histories (see for example Blau and Kahn 2006). This approach can yield an incomplete accounting because surveys cannot contain complete information about the quantity and quality of prior work experience, or about whether particular actions (such as geographic mobility) were intended to be career-enhancing or not.[1] In contrast, our approach does not require an exhaustive measure of all possible PSIs,

[1] Constructing work history measures poses numerous difficult questions. For example, is tenure on the current job plus previous experience a sufficient statistic? Or is tenure on previous jobs, plus the number of and timing of

since the value of all previous PSIs is incorporated into current earnings potential at each point in time. At the aggregate level, observed changes in the relative slopes of age-earnings profiles reveal information about the contribution of PSI's to the narrowing of the gender wage gap.

The other contribution of our analysis is that, unlike studies based on existing panels, we use data that span both many years and many cohorts. Combining information from the 1960 through 2000 Censuses and the 1964 through 2004 Current Population Surveys--the data span 40 years and more than 80 birth cohorts--genuine between-cohort comparisons are possible.[2] This is important because existing analyses of the evolution of the gender wage gap have concentrated their attention on the earnings premium associated with potential (or actual) experience across a series of cross-section regressions.[3] As is well known, however, cross-sectional patterns can sometimes give a misleading picture of what happens to a cohort of workers over time.[4] In particular, what might appear in repeated cross-sections to be an increasingly strong relationship between age and women's earnings might be caused purely by a difference across generations in women's entry-level relative wages.

In sum, this paper draws a distinction between components of the gender wage gap that are present when women first enter the labor market (i.e. factors affecting earnings 'levels') and those that are reflected in women's relative rate of wage growth after labor market entry ('slopes'). Under reasonable conditions, when following a cohort of women over time, we

---

all interruptions also important? Do hours on all previous jobs matter, and how important is past industry and occupational mobility? In addition, it is hard to imagine convincing measures of changes across cohorts in the amount of "energy" women put into their jobs compared to the home (Becker 1985).

[2] The PSID, beginning in 1968, is the only panel that comes close to providing the same time coverage. But, to improve on our estimates we would need a panel that enables us to see complete histories for all earnings observations, including those in the 1960's. This means the panel would have to begin decades earlier.

[3] See for example O'Neill (2003), who presents dramatic evidence that the women's cross-sectional return to potential experience has increased.

[4] In one well-known example, Borjas (1985) showed that what appears in cross-section data to be a difference in earnings slopes between two groups (in his case, immigrants versus natives) was instead driven to a great extent by cross-cohort differences in immigrants' earnings levels at labor market entry (i.e. "cohort quality"). In another, Welch (1973) showed that low apparent rates of return to schooling among blacks in cross sectional data were artifacts of cohort effects.

argue that the latter provides a summary estimate of the contribution of all post-schooling investments to the gender wage gap. We believe that this estimate provides a complementary and –to the best of our knowledge—new perspective on why the gender gap has declined over the past four decades.

Our main results are as follows. First, consider the age-wage profile of a typical cohort in our data. According to simple models of experience and wages, we might expect the women's wages in the cohort to fall further and further behind the men's as the cohort ages, due to men's greater PSIs. Instead (and in contrast to what is observed in most cross sections), we find that women's relative wages follow a U-shaped pattern with age over the life of a typical cohort, falling behind at first, but recovering significantly after that.[5] Second, (and related) even after controlling for the level of educational attainment, a surprising share of a typical cohort's lifetime maximum gender-wage gap is present at very young ages. For example, women aged 23-32 in 1959 faced a gender wage gap (holding education constant) equal to 74% of the maximum gender wage gap ever experienced by that cohort over its remaining lifetime. Comparable figures for later cohorts are even higher. Third, once we allow for a U-shaped "baseline" age effect, our data do exhibit growth across cohorts in the relative slope of women's age-earnings profiles, consistent with the notion that rising post-schooling investments play some role in the declining gender gap. Fourth, we estimate that this growth in relative slopes accounts for about one-third of the narrowing of the gender wage gap over the past 40 years, with the remainder due to a declining wage gap at the time of labor market entry.

[5] Of course, this U-shaped pattern suggests a model in which women's investments are delayed (perhaps until after childrearing), rather than reduced (see for example Polachek 1975; Weiss and Gronau, 1981). That said, it is also consistent with models in which the level of discrimination falls toward the end of the life cycle, relative to levels prevailing at the time of cohort entry (Blau and Kahn 2000).

This estimate is robust to a number of sensitivity checks, including the data source (CPS versus Census) and the incorporation of pure "year" effects into our model.

To fix ideas, Section 2 of the paper describes a simple model in which changing slopes of women's age-earnings profiles account for the entire decline of the gender wage gap. It also describes a simple competing explanation of the declining gap based purely on cohort effects on wage levels at labor market entry, and demonstrates that these two scenarios cannot be distinguished by comparing changes over time in the slope of cross-sectional age-earnings profiles: cohorts must be explicitly followed over time. Sections 3 and 4 present our empirical results on within-cohort wage growth from Census and CPS data respectively, using a hybrid model that allows both slopes and levels to change between cohorts. Section 5 discusses the interpretation of those results and Section 6 concludes.

2. The Implications of Post-Schooling Investments: A Simple Model

This section describes two extreme-case models of the decline in the U.S. gender wage gap. The first, illustrated in Figure 1, attributes all the decline to a changing relative rate of age-related wage growth caused by gender differences in work experience and other post-schooling investments.[6] In the example depicted in Figure 1, we assume that every cohort of women earns 80 percent of the male wage on labor market entry, and (for simplicity) that women's relative wages decline linearly throughout each cohort's lifetime, the latter phenomenon due to women's lower post-schooling investments.[7] Among more recent cohorts, however, higher rates of PSIs imply that women's relative wages fall more slowly with age

[6] For convenience the time frame for these examples was chosen to mirror exactly our Census data, which follow eight ten-year birth cohorts ranging from 1897-1906 to 1967-1976 over the five census years 1959 through 1999. Cohorts in Figure 1 and subsequent figures are labeled by the year in which the cohort's median age was 27; thus for example cohort 1 --our oldest--, was born in 1897-1906, and was between the ages of 23 and 32 in 1929).
[7] Linearity is not imposed in our empirical implementation.

than in earlier cohorts. Despite the fact that women's relative wage at labor market entry is constant across generations of women in this hypothetical example, Figure 1 shows that the average level of women's relative wages in any particular year (given by a weighted average of the vertical array of points in any particular year) is rising over time; this entire increase is accounted for by the changing slopes. The gender gap in the estimated "returns" to potential experience estimated from a cross-section regression also declines across years. This is illustrated by the fact that the vertical array of points is closer together in 2000 than in 1960.

Next, for purposes of comparison, Figure 2 depicts a contrasting model of the rise in women's relative wages. In this "pure cohort effects" (or pure pre-market investments) model, the relative slope of women's age-earnings profile remains constant across all cohorts; however women's relative wages at labor market entry are changing as new cohorts enter.[8] As in Figure 1, this model also has rising women's wages over time; however in this case none of the increase is explained by changes in the returns to potential experience. More importantly, despite the fact that every cohort's age-earnings profile has the same slope, estimates from cross-section regressions will show that the gender gap in returns to potential experience declines across years (the points are closer together vertically in 2000 than in 1960).[9] We conclude that caution is required in drawing conclusions from time trends in the slope of cross-sectional age-wage profiles alone.

In the remainder of the paper we ask how the actual patterns observed over the past 40 years in the United States compare with each of the two alternative scenarios described

[8] Of course, a different alternative to the "changing slopes" model is a pure "year effects" model. We discuss the performance of such a model, and the identification issues that arise in distinguishing year effects, cohort effects and (changing) age effects in Section 5 of the paper.
[9] Mathematically, this results from Figure 2's assumption that the rate of decline in the gender wage gap is decelerating across cohorts. This would be the case, for example, if women's cohort-specific relative wage at labor market entry asymptotically approaches one from below, which strikes us as a plausible scenario.

above—changing age-wage slopes or changing relative wages at labor market entry. We then estimate a decomposition describing the comparative contributions of each to the narrowing of the U.S. gender wage gap. Compared to existing studies --many of which attempt to account for the declining gap by measuring observable investments made by women in a series of cross-sections-- this decomposition offers a complementary perspective. This is because it is unlikely that any data set will ever measure all the possible post-schooling investment differentials between women; studies attempting to do so will therefore always be vulnerable to the critique that they have failed to measure some critical, unobservable aspect of women's investments (for example, the "energy" put into the job compared to the home-- Becker 1985). By searching for the 'footprint' of all PSIs that are made after leaving school in the overall rate of age-related wage growth, we generate an alternative estimate of the role of PSIs that is not vulnerable to this critique. Aside from providing an alternative measure, this estimate may help bracket the total contribution of PSIs because --at least under reasonable conditions that are discussed later in the paper-- it places an upper bound on the role of PSIs.

3. Results: Census Data

a. Estimating Gender Gaps

Our Census samples comprise U.S. born, full-time, full-year white workers aged 23-62 in the years 1959, 1969, 1979, 1989 and 1999. Simple descriptive statistics for these samples are provided in Appendix 1; their main features are well known.[10] In our analysis, gender earnings differentials are estimated for four birth cohorts in any given year, corresponding to workers who attain the age ranges 23-32, 33-42, 43-52 and 53-62 in that year. Altogether, the analysis

---

[10] The difference in the 1999 mean gender gap between the Census and CPS data is due to measurement error, and disappears when data from surrounding CPS years (1995-2003) are aggregated.

includes at least one year of data for each of eight ten-year cohorts with birth dates ranging from 1897-1906 for the oldest cohort to 1967-1976 for the youngest. In what follows, we refer to the oldest cohort as number 1, followed by cohorts 2 through 8 in turn.

Coefficients from cross-sectional earnings regressions using the above data are reported in Table 1. Since the sample is restricted to full-time, full-year workers and detailed hours controls are included, the dependent variable should be interpreted as an hourly rate of pay. In addition to these hours controls, the Table 1 regressions include standard (and comparable) controls for education and region, plus a quadratic in age (to capture the life-cycle pattern of men's wages).[11] Thus, all of the gender wage gaps that are reported in Table 1 (and whose evolution is studied in the remainder of this paper) are conditional on education: what we are attempting to understand in this paper is why the gender wage gap narrowed even after accounting for the dramatic increase in women's education levels, especially at the college and higher levels, over the past 40 years. Finally, to allow women's wages to evolve differently over the life cycle than men's in as flexible a manner as possible, we include four gender-age interaction terms in each Census year. By construction, the gender coefficients along the diagonals of Table 1 thus describe the gender gap faced by a given cohort of women as those women are followed over time.

Several patterns in Table 1 are immediately obvious: Gender coefficients tend to be larger among older cohorts than among younger cohorts observed in the same year (vertical); gender coefficients fall if a given age group is followed over time (horizontal); but are surprisingly constant when a given cohort is followed over time (diagonal). These gender coefficients are also depicted graphically in Figure 3, with observations from the same cohort

[11] Perhaps because the age groups we consider are so broad, the results are very similar if we allow for a more flexible functional form for men's age-wage profile, such as a set of decade fixed effects.

connected by lines.[12] On closer inspection of the within-cohort trends, we also see that gender wage gaps widen for every cohort as its median age rises from 27 to 37, and fall for every cohort between the (median) ages of 47 and 57, reflecting the nonmonotonic pattern described earlier. Another, perhaps surprising feature of the within-cohort patterns is the share of a cohort's lifetime gender wage gap that is already present at the youngest ages in that table. For example, women aged 23-32 in 1959 faced a gender wage gap (holding education constant) equal to .432 log points. This corresponds to 74% of the maximum gender wage gap ever experienced by that cohort over its remaining lifetime (.587, at ages 43-52). Comparable figures for later cohorts are 83% for the cohort "entering" in 1969 and 94% for those entering in 1979.[13]

Finally, between the ages of 37 and 47, we see a widening of the gender wage gaps for the two oldest cohorts observed in those age ranges, but a narrowing for the two youngest. This more subtle pattern suggests an increase across cohorts in the overall slope of women's relative age-wage profile. This first glance at the data reveals patterns consistent with at least some role for the "flattening slopes" hypothesis, as depicted in Figure 1.

b. Modeling the Evolution of the Gender Wage Gap

To quantify the role of age-cohort interactions in explaining the recent decline in the gender wage gap, we now use the 20 age- and year-specific gender wage gaps estimated in Table 1 (and depicted in

[12] For ease of interpretation, Figures 3 and 4 transform the gender gap coefficients of Tables 1 and 2 into female relative wages (by adding one to the absolute value of the coefficient). This makes them directly comparable to Figures 1 and 2.

[13] Even though we do not observe these later cohorts for their entire lifetimes, the fact that the maximum gap is attained so early in life allows us to make these statements, provided that the maximum for these cohorts is attained in the same decade of life as in previous cohorts.

Figure 3) as data points in some simple aggregate regressions.[14] In particular, we estimate the following empirical model of women's relative wages:

$$RW_j = \alpha + \sum_{a=1}^{3} \beta_a A_j(a) + \sum_{c=2}^{8} \gamma_c C_j(c) + \sum_{c=5}^{7} \theta_c age_j C_j(c) \qquad (1)$$

In this model, RW_j describes women's relative wages in each of the 20 age x cohort cells. The effects of age on relative earnings within cohorts 1-4 are captured by $\sum_{a=1}^{3} \beta_a A_j(a)$, where $A_j(a)$ is an indicator for j belonging to age group a. Pure cohort effects on women's relative earnings are captured by the term $\sum_{c=2}^{8} \gamma_c C_j(c)$, where cohort 1 (the oldest—born in 1897-1906) is the reference category, and $C_j(c)$ is an indicator for j belonging to cohort c. Finally, for each of cohorts 5 through 7, the interaction term (modeled as the product of age and the cohort indicator) allows the effect of age on women's relative earnings to differ from its effect in cohorts 1 through 4.[15] In the regressions, age_j is scaled to measure potential experience in decades elapsed since the first observation of the cohort, so $age_j \in \{0,1,2,3\}$. If each successive cohort of women after cohort 4 had a steeper age-relative wage profile, as predicted by the PSI model, we would observe that $0 < \theta^5 < \theta^6 < \theta^7$.[16]

OLS estimates of equation (1) using the 20 estimated values of RW_j in Table 1 as data are presented in column 1 of Table 2. Overall, these results provide support for both the changing levels and changing slopes models. As predicted by the changing levels model, cohorts of

[14] We could, of course pool the microdata from all years, and estimate the influence of cohort-age interactions on the gender wage gap in a one-stage procedure. We prefer the current approach because it allows for an intuitive and visual analysis of the 20 gender-wage gaps that interest us here. Since the samples from which these 20 data points are estimated are so large, the consequences of using estimated coefficients as regressors are minimal.

[15] For practical reasons, we chose to group cohorts 1-4 together to estimate the baseline profile. Figure 4 shows little change in the age-relative wage profile within this group of cohorts.

[16] Since cohorts 1 and 8 are observed only once in these data, separate age effects on earnings cannot be estimated for them.

women born later earn significantly more, relative to men, than cohorts born earlier, even at labor market entry. And, as predicted by the PSI model, younger cohorts of women exhibit a higher rate of age-related relative wage growth than older cohorts. Together, the "composite" model in column 1 fits our data on gender wage gaps almost perfectly, with an adjusted R^2 of .96. But how can we assess the quantitative contribution of changes in the slope of age-earnings profiles across cohorts to the recent narrowing of the gender wage gap?

One way to pose this question is to ask how well we can explain recent trends without recourse to steepening age-wage profiles at all. To that end, column 2 of Table 2 restricts all the cohort-experience interaction terms of column 1 to equal zero ($0 = \theta^5 = \theta^6 = \theta^7$). According to this model, there has been a large and statistically significant increase in women's relative wages upon labor market entry across cohorts. As might be expected, the growth in estimated cohort effects is somewhat larger in this model than in column 1. Somewhat less expected, though, is the fact that relatively little explanatory power is lost by dropping the interaction terms from the model: adjusted R^2 only falls to .91 from .96. Thus, a parsimonious model that fits the past four decades of data on the U.S. gender wage gap surprisingly well has each successive cohort of women entering the labor market at a higher wage relative to men, with each cohort having the same rate of wage growth, relative to men, as every other cohort.

For comparison, column 3 of Table 2 estimates a "pure changing slopes model" (as depicted in Figure 1) where women's relative entry wages are constrained to be the same across all cohorts (and years) but their relative rate of age-related wage growth can differ across cohorts. Compared to the composite model of column 1, the cohort-experience interactions are now much stronger, as the model attempts to fit the declining gender gap across cohorts using slope terms only. Clearly, however, with an adjusted R^2 of .72 this model does a considerably worse job of fitting the data than either the composite or the pure cohort model. We conclude that, if an analyst had to choose only one of these two polar case models to describe the evolution of the gender wage gap over the last 40 years, he or she should choose the changing levels model over the changing slopes model.

c. Decomposing the decline in the gender wage gap

An alternative way to quantify the effects of steepening relative age-wage profiles on the narrowing of the gender wage gap uses the coefficients estimated in the "composite" model of column 1, Table 2 to predict what the 1999 gender wage gap might have been in the absence of these effects. By comparing the actual change in the gender wage gap with the change that would have occurred under the counterfactual assumption of no changes in slopes, we can estimate the relative importance of changing slopes to the narrowing of the gender wage gap.

The first two columns of Table 3 are similar to the regressions of Table 1, columns 1 and 5, except that only a single gender coefficient is estimated for women of all ages in 1959 (column 1) and 1999 (column 2). The estimated log wage differentials are -0.540 in 1959, and -0.278 in 1999. The third column is similar, but the dependent variable is the actual 1999 log wage minus the potential experience*cohort interactions estimated in Table 2, column 1 (zero for men and for the youngest women, .068 for women age 33-42, .067*2 for women age 43-52, and .051*3 for women age 53-62). Under the counterfactual, the 1999 log wage differential is -0.361 . In other words, about one-third of the closing of the gender wage gap between 1959 and 1999 ((.361-.278)/(.540-.278)=.32) can be attributed to changes in slopes.[17]

[17] The contribution of changing age distributions was estimated to be negligible. When the 1959 women were reweighted to match the 1999 age distribution, the estimated gender coefficient fell by only 0.002 . When *pe*cohort* interactions were estimated from a single stage regression (gender interacted with age, cohort, and the

4. Results: CPS data

To assess the robustness of our results to the data source, in this section we replicate the main aspects of our analysis using March CPS data.[18] The main advantage of these data over the Census is that they allow us to estimate an annual series of gender wage gaps, disaggregated by exact year of age rather than aggregated age group. The key disadvantage is fewer observations per year, and a later initial observation. To overcome the absence of early observations, we simply used the 1960 Census data as a proxy for 1960-1963 CPS observations, so that we can compare CPS with Census results over the same 1960-2000 observation window.[19]

CPS estimates are based on 41 years (1960-2000) and 40 ages (23-62), broken down into 10-year, 5-year, or 1-year cohorts and age groups. Rather than run dozens of cross-section regressions with between 4 and 40 gender*age coefficients apiece, we use the other common two-stage procedure to compute the gender gaps. In the first stage, male-only regressions are run for each year, with controls identical to those used in the Table 1 regressions. Gender gaps for each individual woman are computed as the difference between her actual earnings and the earnings predicted if she were paid as much as a man with the same observable characteristics. The data points used in the second stage regressions are mean estimated log wage differentials for each age group*cohort cell, estimated from male-only regressions.[20]

---

*pe*cohort* interactions; year interacted with education, region, and hours per week), the estimated contribution of changing slopes fell from 0.32 to 0.29 .

[18] As with the Census, we observed annual earnings for the year preceding the survey date. These data were collected in 1964-2004, so we have earnings observations for 1963-2003.

[19] The assumption was made that the wage structure remained stable during these early years, i.e. that the gender gap for a 39 year old in 1960 was a good proxy for a 39 year old in 1962. In sensitivity testing, the later CPS observations (2001-2004) are also included.

[20] To maintain comparable units between specifications, we estimate potential experience*cohort interactions using the mean of (age/10) in the cell, minus the mean of (age/10) at the first observation of the cohort. In

Table 4 reports the results of the second-stage regressions. In column 1, the Census data are subjected to this procedure, with the same 10-year age and cohort groups as before. In column 2, the CPS data from the five Census years are used, with very similar results. Column 3 still uses 10-year age and cohort groups, but includes data from between-Census years.[21] For a visual comparison to the Census estimates, Figures 5 and 6 display the results of the Column 3 specification. Column 4 does the parallel analysis using 5-year cohorts and age groups. Column 5 is similar to columns 1 and 2, but with 1-year cohort and age groups, for a total of 1640 1-year cohort*age cells. This final specification includes 40 age levels and 80 cohort-specific fixed effects. The estimated potential experience-cohort (*pe*cohort*) interaction effects are similar in all five specifications.

All specifications produce comparable estimates of the *pe*cohort* interactions, and find U-shaped age-relative earnings profiles with upswing after age 47.[22] To help assess the magnitudes of the experience effects estimated from each of the 5 specifications, the analog of the Table 3, column 3 counterfactual log wage differential is included at the bottom of each column.[23] In every case, the CPS estimates of the importance of changing slopes to the overall change in the gender wage gap are very close to the corresponding Census estimate (based on male coefficients). These estimates do not change much as the within-cell age-range shrinks, indicating that the larger groupings do not obscure important changes in slope. All estimates, both Census and CPS, attribute 30-40 percent of the change in the log wage differential to changing slopes.

---

specifications where cohorts do not always fit neatly into a single age group, the age group indicator is replaced with a vector indicating the proportion of the cell in that age group.

[21] In column 3, the age-group controls indicate the proportion of the cell in the age group.

[22] The column 5 specification contains 39 age coefficients, too many to report in Table 4. These are monotonically increasing from age 47 through 62.

[23] Here, we subtract the *pe*cohort* interactions estimated in the corresponding specification from women's actual 1999 log earnings, then estimate the 1999 log wage differential using this counterfactual as the dependent variable.

Additional sensitivity testing finds that these results are robust to variations in specification. For example, the estimated importance of potential experience*cohort effects is virtually unchanged if we use an estimate of age minus education (rather than simply age) as our measure of of potential experience.[24] Similarly, there is very little change if the oldest workers are eliminated from the sample.[25] If the analysis is restricted to college graduates only, changes in slopes play a somewhat larger role when the Census data are used, but a somewhat smaller role when the CPS data are used. Although we expected work experience to play a much greater role in the closing of the college graduate gender gap, our data do not provide clear support for this hypothesis.

Only one experiment led to a substantial change in the results. When the window of observation was shifted by two or four years (1962-2002, or 1964-2004), the estimated potential experience*cohort effects fell substantially, implying an even smaller role for rising post-schooling investments.[26] For example, when the Table 4, column 5 specification was run with the window shifted by two years, the estimated importance of changing slopes fell from 39 percent of the closing to only 24 percent of the closing. When the window was shifted by four years, it fell even further, to only 12 percent. This sensitivity to the time frame chosen suggests that a more complete model should incorporate a time trend-- a possibility we explore in the following section.

[24] Since Census questions do not always allow us to estimate total years of education received, and since men and women had similar education levels throughout our sample period, we simply used age-22 as our potential experience measure. When this is changed to age-education-6, estimated *pe*cohort* effects were virtually unchanged, but slightly smaller.

[25] Estimated *pe*cohort* effects were similar when the age range was 23-52 rather than 23-62.

[26] The two- or four-year shift included the start and end dates of the sample, and the birthdates of cohorts used to calculate the *pe*cohort* interactions.

5. Interpreting the Changes in Slope

In our discussion so far, we have made two assumptions that considerably simplify the interpretation of our empirical results. One of these is the (implicit) assumption that the population of women employed at any point in time (for whom we have wage data to analyze) is representative of all women who ever participate in the labor force. How can we interpret our estimates if this is not true? To address this question, we note first that --since our interest is in changes over time-- nonrandom selection of women into work will not affect our estimates of the importance of PSIs if the nature of selection is constant over time. But what if the nature of women's selection into work has become increasingly positive (with respect to unobserved ability) over the past 40 years, as both Blau and Kahn (2006) and Mulligan and Rubinstein (2005) have argued? Of course, if changing selection into work has been purely a cohort-specific phenomenon, its effects would be absorbed by our cohort coefficients; in this case our estimates of cohort-age interactions, and thus the importance of PSIs would remain unbiased.

In consequence, changing selection of women into work will affect our estimates of the contribution of changing PSIs to the declining gender wage gap only to the extent that the nature of within-cohort selection into work over the life cycle changed between cohorts. For example, if the within-cohort trend towards positive selection into work was stronger within earlier cohorts of women than within more recent ones, our estimates will understate the contribution of PSIs to the decline in the gender gap. While this is certainly possible, the opposite seems more likely to us: that the recent time trend towards positive selection into work has had both a between- and a within-cohort component. If that is the case, the changes in slope that we estimate here would include some positive within-cohort selection effects, and would thus overstate the true contribution of changing post-schooling investments to the narrowing of the gender wage gap.[27]

A second key assumption in our analysis so far is the absence of any pure year effects in our models of the gender wage gap. In addition to some well-known reasons related to model identification (discussed below), we formulated our baseline model in terms of cohort and age effects (and their interactions) because in a certain sense year effects are the hardest to interpret in this context: cohort effects naturally capture pre-market investments and cultural differences across generations, while age effects (and the changes therein across cohorts) naturally capture post-schooling investments. Aside from a decline in discrimination, it is harder to think of reasons why the relative price employers are willing to pay for female versus male labor of given age and experience would rise purely as a function of calendar time across all cohorts.

That said, however, suppose now that women's relative wages are also affected by a simple time trend in the relative price of female labor (caused by an increase in relative demand for female labor, such as a decline in discrimination). How would the presence of such a time trend affect our conclusions? Of course it is impossible to estimate a regression that adds a complete set of year effects to equation 1: year, cohort and age effects are not separately identified because year equals cohort plus age (see for example Deaton and Paxson 1994 for an excellent exposition).[28] Still, it is important to ask whether the presence of such a trend would

---

[27] One exception we can think of is that welfare reform during the 1990s may have pushed an additional group of less-able women into the labor market. This might account for the deceleration in year effects during that decade, discussed later in the paper. Additional insights on changing patterns of selection into work could of course be gained from examination of long, nationally representative panel data sets. While certainly of interest, this is beyond the scope of the current paper.

[28] As is very well known, when an outome variable (say RW) is modelled as a linear function of age, cohort and year, then (because year=age+cohort) a regression of RW on any two of these three variables is numerically equivalent to a regression on any other two. Less well known is the fact that this observational equivalence breaks down when age, cohort and year are each represented as a vector of fixed effects. To see this most quickly, suppose (as is the case in this paper) the the number of years defining a time period is the same as the number

affect our quantitative assessment of the contribution of cohort-age interactions to the decline in the gender wage gap. We address this question formally in Appendix 2; the results are as follows. First, if there is a true, positive linear time trend in women's relative wages, both the age effects (β's) and the cohort effects (γ's) presented in all our Tables so far will be overestimates, because any true year effects will be absorbed into these age and cohort coefficients. That said, our estimates of cohort-year interactions (the θ's) will remain unbiased. More importantly, it follows that the share of the year-to-year decline in the gender gap attributed by our model to cohort-experience interactions (versus cohort and year effects together) is invariant to the magnitude of any linear time trend affecting women's wages.

Unfortunately, as Appendix 2 goes on to show, the above reasoning does not extend to nonlinear time trends; in such cases our estimates of the relative importance of cohort-year interactions (versus all other factors together) will be biased. In particular, when the time trend is accelerating, our estimated changes in earnings slopes will overstate the true trend in cohort-year interactions; when the time trend is decelerating, changes in slope understate the trend in cohort-year interactions. Given this information, is it possible to place bounds on the true contribution of PSIs to the declining gender gap?

To do this, we first argue that our best empirical evidence about time trends, net of PSI effects, comes from observation of gender gaps among very young workers. These are depicted graphically in Figure 7, using both Census and CPS data for ages 23-27. These series are flat between 1960 and 1970, then increase at a constant rate until 1990 or 1995 before flattening out again. If any part of this nonlinear trend reflects year rather than cohort effects,

---

defining a cohort. Then in any data set, the number of time periods cannot equal the number of cohorts. In data (like ours) collected as a series of cross sections, there will be more cohorts than years (8 versus 4 in our case). In data sets that follow a group of cohorts over time (such as the NLS series) there will be more years than cohorts.

then the changes in earnings slopes we estimate in Tables 2 and 4 incorporate both PSI effects and these year effects.

Next, we perform a series of simulations, all of them variants of the following experiment: Assume that a fraction α of the time trend in Figure 7 is caused by year effects, while the remainder is due to cohort effects. Use this assumption to create a trend-adjusted version of the relative wage variable, taking out both the assumed year-specific and cohort-specific trends.[29] Then estimate the *pe*cohort* interactions from the de-trended dependent variable. The results of this exercise are reported in Table 5, with the fraction of the trend attributed to year effects (α) varying from zero to one.

In column 1 of Table 5, the entire time trend is attributed to cohort effects.[30] In columns 2-5, a gradually increasing proportion of the time trend is attributed to year effects. As expected (recall that year effects were accelerating in the early part of our sample period), the estimated *pe*cohort* interaction falls for the 1937-1946 birth cohort as increasing importance is assigned to year effects. The effect is substantial. A significant drop is also apparent in the next entering cohort. For the youngest (1957-1967) cohort, which entered the labor market as the time trend began decelerating, the estimated *pe*cohort* interaction grows as year effects are introduced, but only slightly. Overall, as we simulate a larger demand shift towards female labor, the contribution of PSI effects (i.e. cohort-age interactions) to the narrowing of the gender wage gap tends to fall, but not dramatically (see Figure 8, or the last row of Table 5).

---

[29] The exact specification can be found in Appendix 4.
[30] This differs slightly from the Table 4, Column 1 specification because the cohort fixed effects are replaced by the more constrained functional form described in Appendix 4. Similarly, the Column 5 result differs only slightly from that obtained with unconstrained year fixed effects.

For example, if only one-quarter of the trend is attributed to year effects, the contribution of cohort-age interactions falls from .32 to .29 and the fit improves (see Figure 9).[31]

We conclude our discussion of Table 5 by noting an important sense in which the changing slopes and the changing levels models are complementary, rather than substitutes. Specifically, the introduction of even a small time trend induces the set of *pe*cohort* interaction terms to increase monotonically, as the changing slopes model predicts they should (Table 5). Put another way, a model without any year effects suggests that there was a nonmonotonic pattern in the rates of post-schooling investment across successive cohorts of women, which to us seems unlikely. Finally, we note that in the presence of a positive time trend, the estimated contribution of women's rising post-schooling investments to the narrowing of the gender wage gap is unambiguously less than our "baseline" estimate; i.e. less than one-third.

6. Discussion

A widely cited explanation of the recent decline in the gender wage gap focuses on changes across cohorts in the rate at which women make post-schooling investments in their earnings capacity, such as accumulating work experience. In this paper we have assessed the contribution of changing post-school investments to the recent decline in the gender wage gap by decomposing the decline into components associated with the slopes versus levels of women's relative wage profiles across eight cohorts of women observed over the past 40 years. Compared to existing studies, this decomposition offers a complementary (and to our knowledge unique) perspective: rather than attempting to exhaustively measure all PSIs in

[31] This is not an exhaustive set of simulations. For example, suppose relative demand grows more quickly for women at the entry level, while older women remain on previously established career tracks, at least for a while. Then a portion of the time trend will be captured by cohort effects, rather than by changes in slope, and the original PSI estimates will contain less bias than they do in the cases we simulate.

successive cross-sections, we search for the “footprint” of all PSIs together in each cohort’s overall rate of age-related wage growth over several decades of Census or CPS data. Aside from providing an alternative measure, this estimate may help bracket the total contribution of PSIs because --under fairly reasonable conditions-- it places an upper bound on the role of PSIs.

Our analysis focuses on the gender gap that remains after controlling for differences in educational attainment and other factors. We find that (as predicted by a simple PSI-based model) the gender gap does tend to widen during the earliest years of the career, but then actually narrows substantially during much of the life cycle for all cohorts of U.S. workers in our data. Second (and related), for every cohort in our data a surprising share of its lifetime maximum gender-wage gap is present very early in life. Third, some cross-cohort increases in women’s relative rate of age-related wage growth are observed. These increases in slope can account for about one-third of the narrowing of the gender wage gap in a simple baseline model. When that model is expanded to allow for pure time effects on women’s relative wages, or when changing patterns of selection into the labor market are taken into account, the likely contribution of rising post-school investments to the narrowing gender wage gap falls to under one third. Thus, while our analysis provides some support for the popular PSI-based explanation of the decline in the U.S. gender wage gap, it also documents the existence of large, unexplained wage differences across cohorts that are already present at the start of women’s working lives. Further, the reduction in these entry-level differentials accounts for the majority of the narrowing of the overall gender wage gap. What underlying factors, already present at labor market entry, might explain this larger portion of the change in the earnings of women relative to men with similar levels of educational attainment?

Obviously, one set of factors that might account for the large cohort effects in our data is unmeasured changes in pre-market investments. In other words, while our wage gap estimates hold years of education constant, trends in the type or quality of human capital women bring to the labor market (Polachek 1978; Brown and Corcoran 1997; Weinberger 1998, 1999, 2001) could account for the large cohort effects we estimate here. Using information on the detailed college majors of a panel of college-educated workers of all ages from 1989 to 1999, Weinberger (2005) finds that controlling for detailed college major does little to attenuate the large cohort effects.[32] Datcher Loury (1997) finds similar results in younger panels. If no effect is found within panels of college graduates, it seems unlikely that trends in unobserved pre-market human capital investments can account for much of the "unexplained" decline in the gender wage gap in the population of women as a whole.[33]

A second possible explanation of the decline in the gender gap at labor market entry extends the human capital model of equation (1) by allowing cross-cohort differences in *expected* labor force attachment to affect women's entry-level earnings. Of course, in the standard general training model (e.g. Blau et al. 1997, chapter 6), this makes it even harder to explain the cohort effects we estimate here: controlling for pre-market investments, the standard model predicts that early-career wages should actually be lower for persons who expect to be more committed to the labor market, while we observe rising entry-level wages over time. For inter-cohort differences in expected labor force attachment to explain the large cohort effects in our data, one would thus need early career investments to take a different form from what is usually assumed. For example, suppose that rather than taking time away from

[32] Another relevant finding of this paper is that the U-shape is even more pronounced, and swings upward earlier, when a fixed sample of college graduates is followed over time, compared to the corresponding synthetic cohort analysis.

[33] As mentioned earlier, the importance of cohort effects is surprisingly similar at different levels of educational attainment.

production, training investments take the form of increased hours or effort (beyond the level that would be optimal based on the worker's current productivity alone). In contrast to the standard model, these factors should raise earnings during the training period. That said, our main findings include controls for work hours; thus we are skeptical that a model based on changing hours or effort across cohorts of young women can explain the large cohort effects in our data.[34]

Finally, of course, as noted there may have simply been a time trend in the relative price employers are willing to pay for female labor of a given level of education, training and expected future work attachment; a possible cause of such a trend is declining discrimination. As shown in Section 5, such a time trend would appear in our baseline model as upward bias in our cohort effects. Thus, declining discrimination could also account for the large and declining estimated cohort effects in our baseline model. In this regard, we also note that, in at least one key sense, introducing some pure year effects actually improves the performance of the "changing slopes" model in our data: when year effects are introduced, our estimated trend in age-cohort interactions becomes monotonically positive --as predicted by the PSI model. In this sense, then, the "PSI model" and the hypothesis that discrimination against women has declined over the past 40 years may be complementary, rather than substitutes, as possible explanations of the recent decline in the U.S. gender wage gap.

[34] An alternative modification to the basic model would be to make training firm-specific. For example, suppose, as in Kuhn (1993), that returns to specific training are shared between workers and firms, and that entry-level wages are determined by a zero-expected-profit condition for firms given each demographic group's probability of remaining with the firm after training is complete. Now, because workers are paid some of their expected post-training productivity "up front," an increase in the expected labor force attachment of a cohort of women can, under reasonable conditions, *raise* the starting wages of that cohort. While this is an important possibility, we note that it can only apply to firm-specific components of on-the-job training.

References

Becker, Gary S. “Human Capital, Effort, and the Sexual Division of Labor” Journal of Labor Economics Part 2 January 1985; 3(1): S33-58

Blau, Francine D., Marianne Ferber, and Anne Winkler. The Economics of Women, Work and Pay (3rd edition) Upper Saddle River, NJ: Prentice-Hall, 1997.

Blau, Francine D. and Lawrence M. Kahn. “Gender Differences in Pay”. Journal of Economic Perspectives 14(4) (Fall 2000): 75-100.

Blau, Francine D. and Lawrence M. Kahn . “The U.S. Gender Pay Gap in the 1990s: Slowing Convergence”. Industrial and Labor Relations Review 60(1) (October 2006):45-66.

Borjas, George. "Assimilation, Changes in Cohort Quality, and the Earnings of Immigrants", Journal of Labor Economics 3 (1985): 463-489

Brown, Charles and Mary Corcoran. “Sex-Based Differences in School Content and the Male-Female Wage Gap.” Journal of Labor Economics 15 (1997):431-465.

Datcher Loury, Linda. “The Gender Earnings Gap among College-Educated Workers.” Industrial and Labor Relations Review, 1997, 50(4):580-593.

Deaton, Angus S. and Christina H. Paxson. “Saving, Growth and Aging in Taiwan”. In David A. Wise, ed. Studies in the Economics of Aging. Chicago: University of Chicago Press, 1994, pp. 331-361.

Goldin, Claudia. “Life-Cycle Labor Force Participation of Married Women: Historical Evidence and Implications”. Journal of Labor Economics 7(1) (January 1989): 20-47.

Kuhn, Peter. "Demographic Groups and Personnel Policy", Labour Economics 1 (June 1993): 49-70.

Madrian, Brigitte and Lars Lefgren. 1999. “A Note on Longitudinally Matching Current Population Survey (CPS) Respondents.” NBER Technical Working Paper #247, www.nber.org/papers/T0247.

Mulligan, C., and Y. Rubinstein. “Selection, Investment and Women’s Relative Wages Since 1975”. NBER working paper no. 11159, February 2005.

Mincer, Jacob and Polachek, “Family Investments in Human Capital: Earnings of Women,” Journal of Political Economy 82(2,part 2)(March-April 1974): s76-S108.

O’Neill, June. “The Gender Gap in Wages, circa 2000”. American Economic Review 93(2) (May 2003): 309-314.

O'Neill, June and Solomon Polachek. "Why the Gender Gap in Wages Narrowed in the 1980's". Journal of Labor Economics 11(1) (Jan. 1993): 205-228.

Polachek, S. "Potential Biases in Measuring Male-Female Discrimination," Journal of Human Resources, 10(2) (Spring 1975): 205-229.

Polachek, S. "Sex Differences in College Major" Industrial and Labor Relations Review July 1978; 31(4): 498-508.

Weinberger, Catherine J. "Race and Gender Wage Gaps in the Market for Recent College Graduates." Industrial Relations 37(January, 1998):67-84.

Weinberger, Catherine J. 1999. "Mathematical College Majors and the Gender Gap in Wages." Industrial Relations 38 (July, 1999):407-13.

Weinberger, Catherine J. "Is Teaching More Girls More Math the Key to Higher Wages?" in Squaring Up: Policy Strategies to Raise Women's Incomes in the U.S., edited by Mary C. King. Forthcoming June 2001, University of Michigan Press.

Weinberger, Catherine J. 2005. "In Search of the Glass Ceiling: Cohort Effects among U.S. College Graduates in the 1990's" UCSB Working Paper.

Weiss, Y. and R. Gronau. "Expected Interruptions in Labour Force Participation and Sex-Related Differences in Earnings Growth" Review of Economic Studies 48(4) (October 1981): 607-619.

Welch, Finis. "Black-White Differences in Returns to Schooling" The American Economic Review, Vol. 63, No. 5. (Dec., 1973), pp. 893-907.

Table 1: Gender Earnings Gaps by Cohort and Year, Cross-sectional Census Data Regressions

	(1)	(2)	(3)	(4)	(5)
Year	1959	1969	1979	1989	1999
female*(age 23-32)	-0.432	-0.424	-0.329	-0.216	-0.198
	(0.004)**	(0.004)**	(0.002)**	(0.002)**	(0.003)**
female*(age 33-42)	-0.558	-0.571	-0.509	-0.350	-0.263
	(0.003)**	(0.004)**	(0.003)**	(0.002)**	(0.002)**
female*(age 43-52)	-0.595	-0.600	-0.587	-0.460	-0.334
	(0.003)**	(0.003)**	(0.003)**	(0.003)**	(0.003)**
female*(age 53-62)	-0.560	-0.539	-0.537	-0.458	-0.336
	(0.005)**	(0.004)**	(0.004)**	(0.004)**	(0.004)**
Age	0.026	0.029	0.040	0.041	0.039
	(0.000)**	(0.000)**	(0.000)**	(0.000)**	(0.000)**
(age-22) squared	-0.000	-0.000	-0.001	-0.001	-0.001
	(0.000)**	(0.000)**	(0.000)**	(0.000)**	(0.000)**
35-39 hours/week	0.018	0.007	-0.069	-0.121	-0.158
	(0.003)**	(0.003)*	(0.003)**	(0.003)**	(0.004)**
41-48 hours/week	-0.015	0.050	0.089	0.121	0.133
	(0.002)**	(0.002)**	(0.002)**	(0.002)**	(0.002)**
49+ hours/week	-0.059	0.049	0.115	0.188	0.245
	(0.003)**	(0.003)**	(0.002)**	(0.002)**	(0.002)**
Observations	238937	296103	393118	497703	538469
R-squared	0.33	0.35	0.32	0.34	0.32

Robust standard errors in parentheses
* significant at 5%; ** significant at 1%
Other controls: census region, education level

Table 2- Testing alternative models of women's relative earnings

	(1)	(2)	(3)
Cohort 2 (born 1907-1916)	0.034 (0.031)	0.037 (0.049)	
Cohort 3 (born 1917-1926)	0.042 (0.030)	0.054 (0.048)	
Cohort 4 (born 1927-1936)	0.069 (0.031)	0.100 (0.047)	
Cohort 5 (born 1937-1946)	0.073 (0.042)	0.179 (0.047)**	
Cohort 6 (born 1947-1956)	0.183 (0.043)**	0.291 (0.051)**	
Cohort 7 (born 1957-1966)	0.287 (0.045)**	0.380 (0.055)**	
Cohort 8 (born 1967-1976)	0.305 (0.045)**	0.390 (0.063)**	
Potential Experience * Cohort 5	0.051 (0.014)*		0.063 (0.021)**
Potential Experience * Cohort 6	0.067 (0.022)*		0.146 (0.035)**
Potential Experience * Cohort 7	0.068 (0.041)		0.286 (0.077)**
Age 33-42	-0.115 (0.022)**	-0.063 (0.027)*	-0.229 (0.047)**
Age 43-52	-0.139 (0.027)**	-0.060 (0.029)	-0.279 (0.046)**
Age 53-62	-0.057 (0.029)	0.028 (0.031)	-0.204 (0.045)**
R squared	0.99	0.96	0.81
Adjusted R squared	0.96	0.91	0.72

Sample size for all regressions is 20 age-year cells estimated in Table 1. Omitted categories are Age 23-32, and Cohort 1 (born 1897-1906). Potential experience is measured as decades elapsed since the cohort was aged 23-32. Cohort-specific earnings growth rates are estimated relative to women born before 1937, and cannot be estimated for Cohort 8 since we have only one year of data for this cohort.

Table 3: Contribution of Changing Slopes to Changes in Gender Earnings Gaps (from Cross-sectional Census Data Regressions)

	(1)	(2)	(3)
Dependent Variable	1959 Earnings	1999 Earnings	1999 Earnings MINUS Estimated Effects of Changing Slopes[35]
Female	-0.540	-0.278	-0.361
	(0.002)**	(0.002)**	(0.002)**
Age	0.023	0.035	0.031
	(0.000)**	(0.000)**	(0.000)**
(age-22) squared	-0.000	-0.001	-0.001
	(0.000)**	(0.000)**	(0.000)**
35-39 hours/week	0.019	-0.161	-0.163
	(0.003)**	(0.004)**	(0.004)**
41-48 hours/week	-0.017	0.133	0.133
	(0.002)**	(0.002)**	(0.002)**
49+ hours/week	-0.061	0.244	0.243
	(0.003)**	(0.002)**	(0.002)**
Observations	238937	538469	538469
R-squared	0.33	0.32	0.34

Robust standard errors in parentheses
* significant at 5%; ** significant at 1%
Other controls: census region, education level

[35] Estimated Effects of Changing Slopes are based on Table 2, column 1 estimates of *pe*cohort* interactions: zero for men and for the youngest women, .068 for women age 33-42, .067*2 for women age 43-52, and .051*3 for women age 53-62.

Table 4—Robustness to Changes in Data Source and Cell-Size

	(1)	(2)	(3)	(4)	(5)
Data Source	Census	CPS	CPS	CPS	CPS
Number of Cohort Controls	8	8	8	16	80
Number of Age Group Controls	4	4	4	8	40
Potential Experience * (born 1937-1946)	0.055 (0.016)*	0.062 (0.020)*	0.066 (0.007)**	0.051 (0.005)**	0.044 (0.003)**
Potential Experience * (born 1947-1956)	0.073 (0.024)*	0.065 (0.030)	0.068 (0.011)**	0.066 (0.008)**	0.065 (0.005)**
Potential Experience * (born 1957-1966)	0.073 (0.044)	0.038 (0.055)	0.046 (0.025)	0.053 (0.016)**	0.068 (0.009)**
Potential Experience *(born 1967-1976)				0.003 (0.077)	0.015 (0.030)
Age 28-32				-0.047 (0.012)**	
Age 33-42 (33-37 in Col 4)	-0.131 (0.025)**	-0.118 (0.031)**	-0.133 (0.012)**	-0.121 (0.012)**	
Age 38-42				-0.150 (0.013)**	
Age 43-52 (43-37 in Col 4)	-0.158 (0.030)**	-0.142 (0.037)**	-0.148 (0.013)**	-0.169 (0.014)**	
Age 48-52				-0.133 (0.014)**	
Age 53-62 (53-57 in Col 4)	-0.069 (0.031)	-0.068 (0.039)	-0.061 (0.015)**	-0.085 (0.014)**	
Age 58-62				-0.011 (0.015)	
Observations	20	20	128	296	1640
R squared	0.99	0.98	0.97	0.95	0.90
Adjusted R squared	0.96	0.93	0.96	0.95	0.89
Counterfactual Log Wage Differential (LWD)	-.365	-.352	-.358	-.364	-.378
Proportion of Change in Gender Gap Due to Changing Slopes	.34	.29	.31	.33	.39

Column 1 & 2 regressions use data from the 5 Census years, columns 3-5 include annual data spanning the same four decades. In columns 1-4 the youngest age group is the omitted category. Potential experience is measured in decades. Cohort-specific earnings growth rates are estimated relative to women born before 1937. Last row computed relative to actual mean LWD from male coefficients: -.277 in 1999 and -.539 in 1959.

Table 5—Simulated Experience Effects Under Alternative Assumptions about Year Effects

	(1)	(2)	(3)	(4)	(5)
Experimental Treatment: Proportion of Entry-Level Narrowing Attributed to Pure Year Effects	**α=0.00**	**α=0.25**	**α=0.50**	**α=0.75**	**α=1.00**
Potential Experience	0.053	0.040	0.028	0.016	0.004
* (born 1937-1946)	(0.007)**	(0.005)**	(0.006)**	(0.008)	(0.011)
Potential Experience	0.070	0.062	0.054	0.046	0.038
* (born 1947-1956)	(0.011)**	(0.008)**	(0.009)**	(0.013)**	(0.017)*
Potential Experience	0.071	0.074	0.077	0.080	0.083
* (born 1957-1966)	(0.020)**	(0.015)**	(0.017)**	(0.024)**	(0.032)*
Age 33-42	-0.141	-0.147	-0.154	-0.160	-0.166
	(0.018)**	(0.014)**	(0.015)**	(0.021)**	(0.029)**
Age 43-52	-0.170	-0.182	-0.193	-0.205	-0.216
	(0.018)**	(0.014)**	(0.015)**	(0.021)**	(0.029)**
Age 53-62	-0.091	-0.109	-0.127	-0.145	-0.163
	(0.018)**	(0.014)**	(0.015)**	(0.021)**	(0.028)**
R-squared	0.92	0.95	0.94	0.90	0.85
Adjusted R squared	0.89	0.93	0.91	0.86	0.78
Counterfactual Log Wage Differential (LWD)	-.362	-.353	-.344	-.335	-.327
Simulated Proportion of Change in Gender Gap Due to Changing Slopes (net of demand shifts)	.32	.29	.26	.22	.19

Data: 20 Census observations, trend-adjusted as described in Appendix 3.

Figure 1: Hypothetical Female Relative Wage Profiles: Pure PSI Model (Changing Slopes)

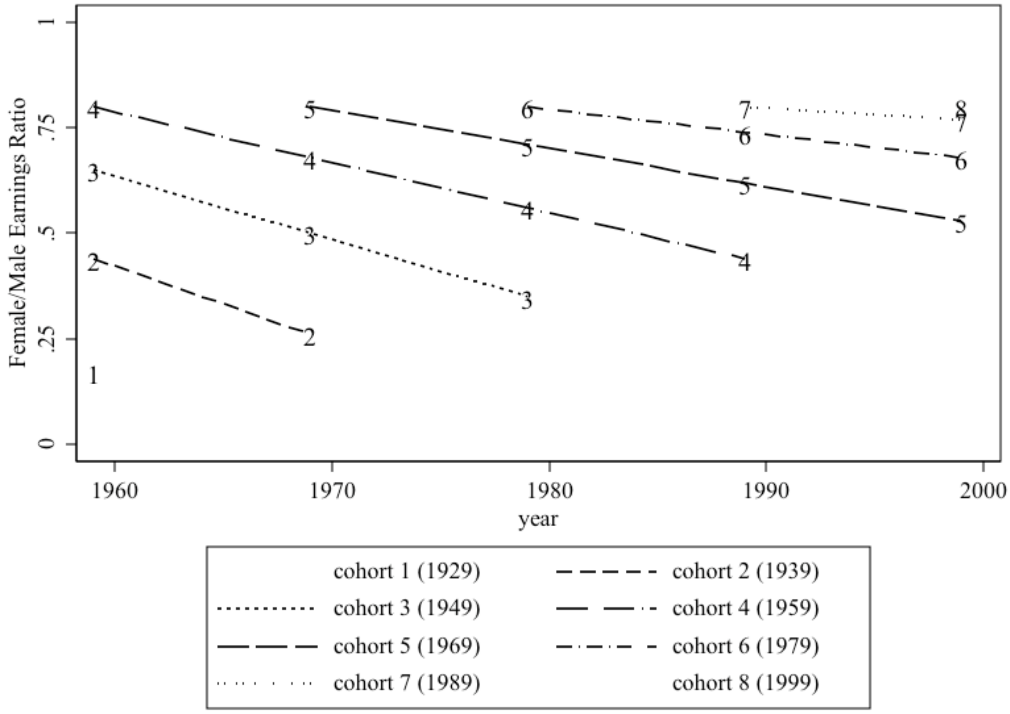

Figure 2: Hypothetical Female Relative Wage Profiles: Pure Cohort Model (Changing Levels)

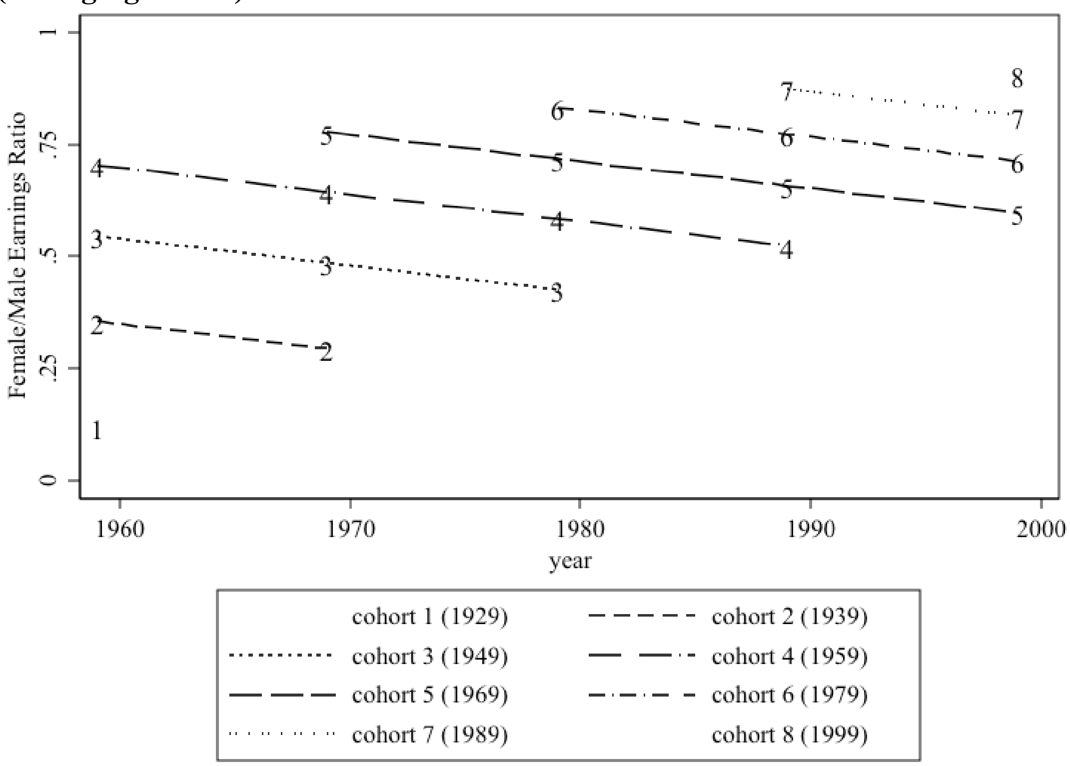

Note: Cohorts are labeled by the year in which their median age was 27

Figure 3: Actual Female Relative Wage Profiles, by year and cohort (Census data)

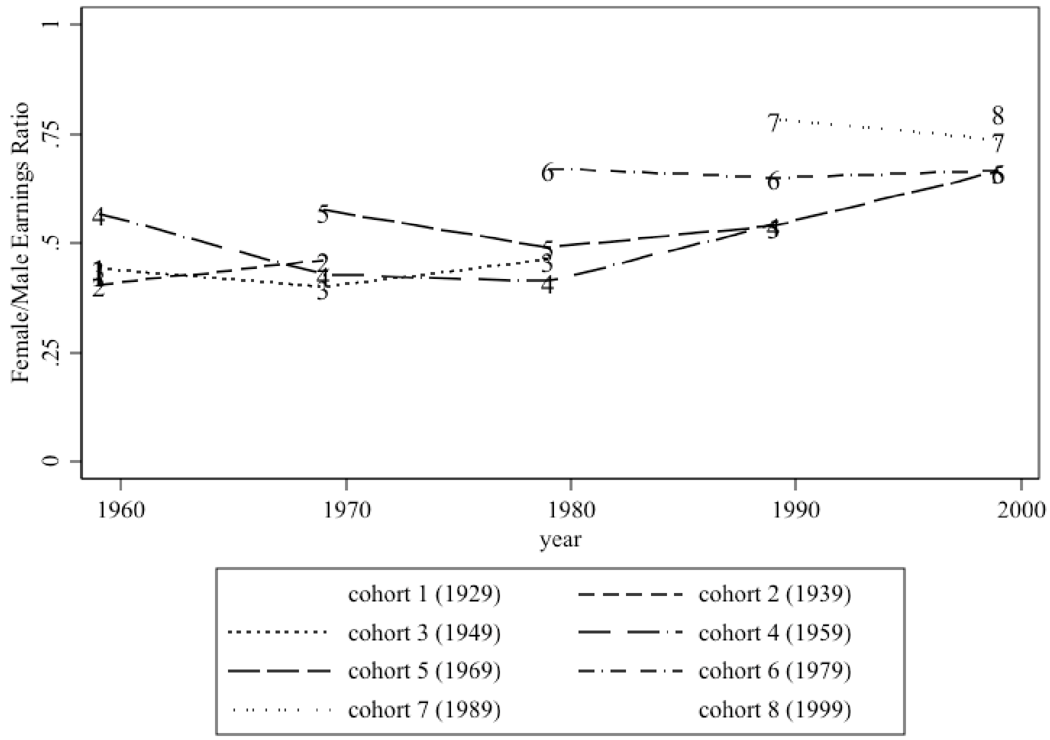

Figure 4: Actual Female Relative Wage Profiles, by age and cohort (Census data)

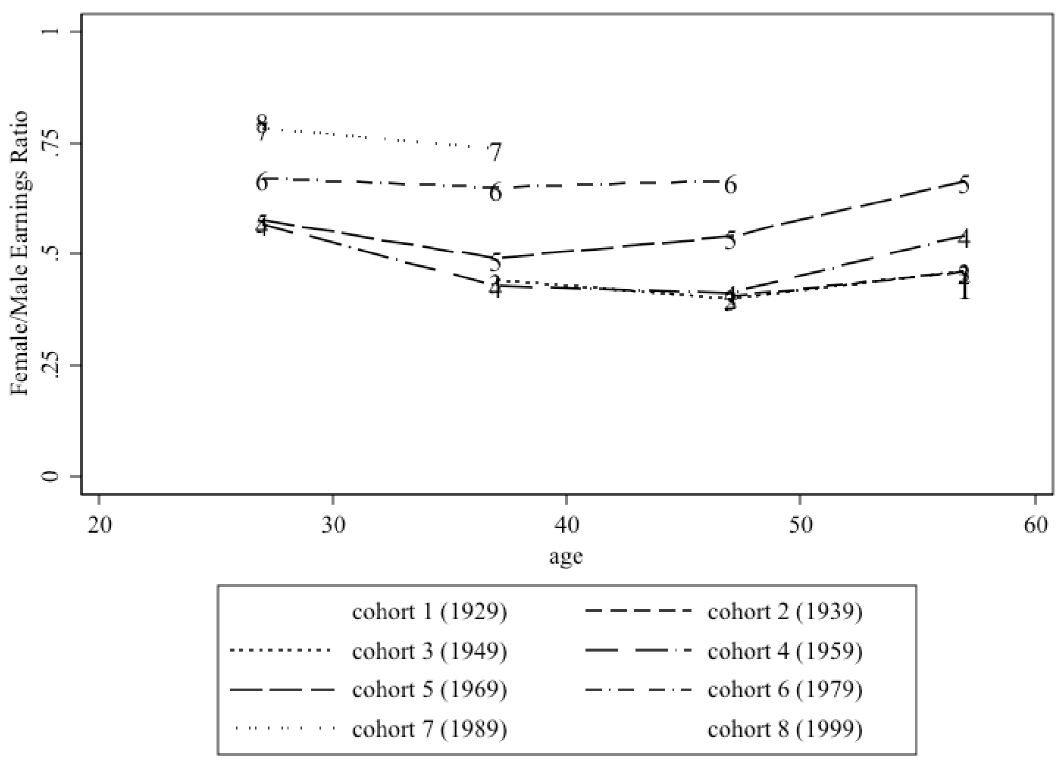

Note:Cohorts are labeled by the year in which their median age was 27.

Figure 5: Actual Female Relative Wage Profiles, by year and cohort (CPS data)

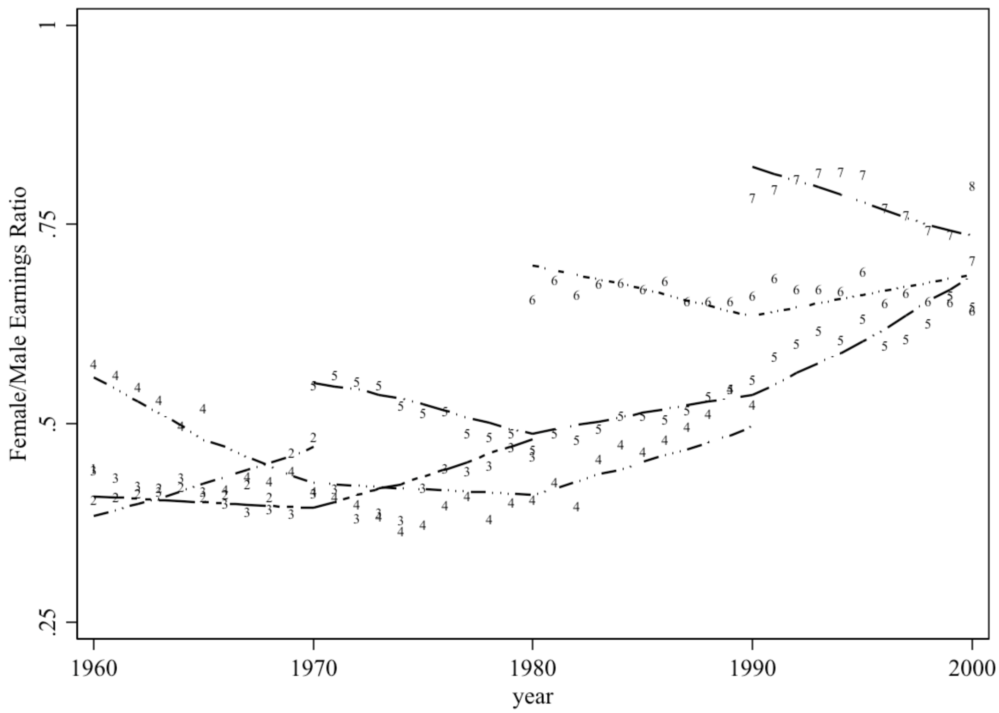

Figure 6: Actual Female Relative Wage Profiles, by age and cohort (CPS data)

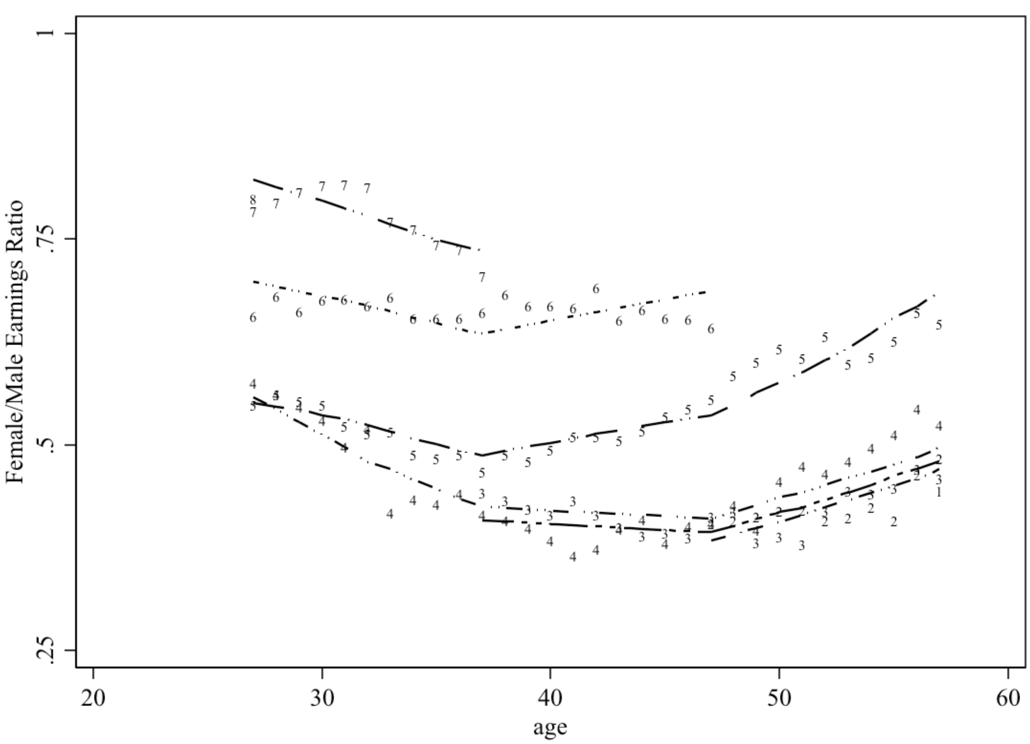

Note: Fitted lines are estimated in Table 4, Column 3. Numeric symbols indicate first stage estimates of the gender gap for each cohort and year. Cohorts 1-8 are defined as in the Figure 3 and 4 legends.

Figure 7: Relative Wages at Labor Market Entry, and Simple Fitted Time Trend.

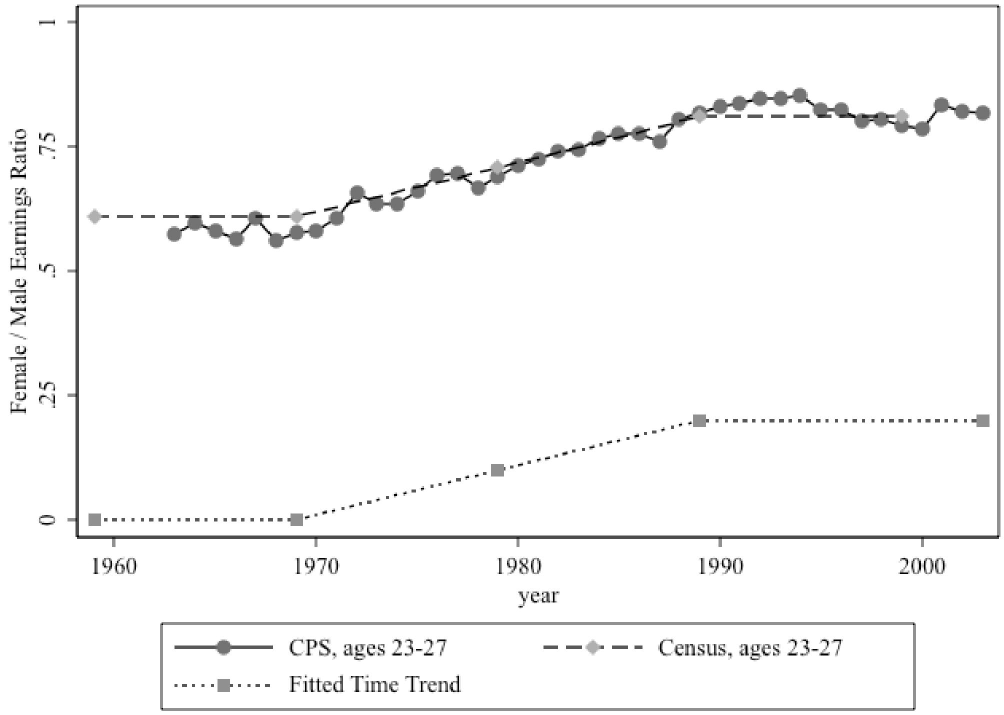

Note: The fitted time trend is a piecewise linear function, with kinks at 1969 and 1989.

Figure 8: Sensitivity of Estimated Cohort and Cohort*Age Effects to Simulated Time Trend in the Demand for Women

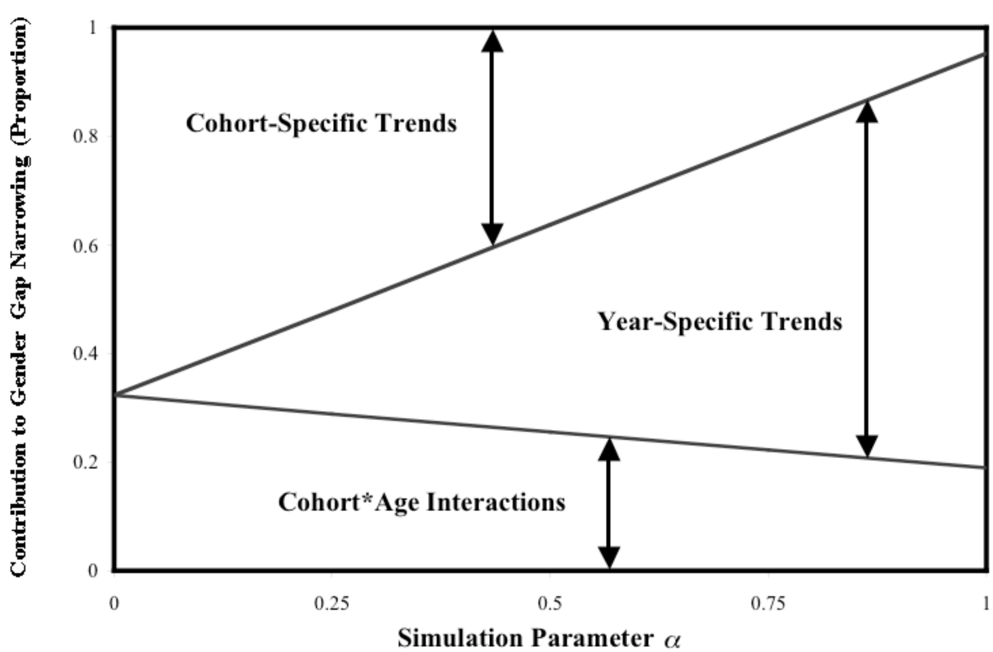

Note: Simulations are described, and α is defined, in Appendix 3.

Figure 9: Actual vs. Fitted Profiles in Simulation Models

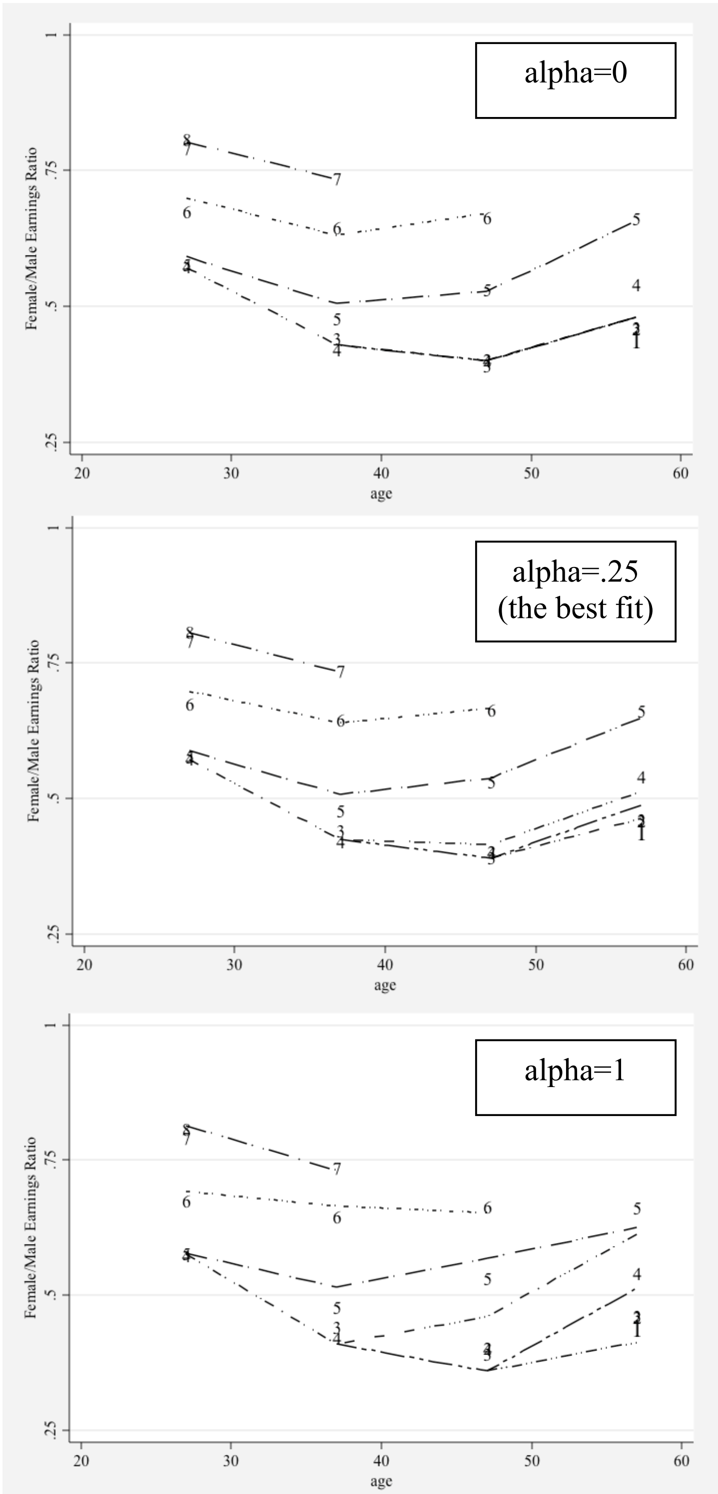

Note: Fitted lines are estimated in Table 5, Columns 1, 2 and 5. Numeric symbols indicate first stage estimates of the gender gap for each cohort and year, using Census data. Cohorts 1-8 are defined as in the Figure 3 and 4 legends.

APPENDIX 1:

Sample means of selected Census variables, by gender (standard deviation in parentheses)

Year	1959	1969	1979	1989	1999
All Men (age 23-62):					
Proportion employed full time, full year	0.64	0.67	0.67	0.67	0.67
N	342443	365602	433780	500352	519909
All Women (age 23-62):					
Proportion employed full time, full year	0.17	0.21	0.29	0.39	0.43
N	355566	385058	448945	516213	531392
Men employed full time, full year:					
Income	6200 (3200)	10100 (5800)	20600 (11500)	35100 (22900)	48800 (32800)
Age	40 (10)	41 (11)	39 (11)	39 (10)	41 (10)
Hours < 40	.04	.05	.04	.03	.02
Hours > 50	.22	.23	.22	.29	.34
College graduate	.13	.16	.26	.29	.31
N	183797	220156	267601	309713	321517
Women employed full time, full year:					
Income	3600 (1500)	5600 (2800)	11700 (5800)	22300 (13200)	33700 (22000)
Age	42 (10)	43 (11)	39 (12)	39 (10)	41 (10)
Hours < 40	.16	.16	.15	.12	.10
Hours > 50	.05	.05	.06	.11	.16
College graduate	.07	.09	.16	.23	.29
N	55140	75947	125517	188010	217073
Census Gender Gap from Male Coefficients	-0.54	-0.54	-0.47	-0.34	-0.28
CPS Gender Gap from Male Coefficients	(1963): -0.54	-0.54	-0.48	-0.34	-0.30

APPENDIX 2: Introducing Year Effects

In order to examine the effects of introducing a time trend into our estimated model, we first formalize post-schooling investments (PSIs) as a cumulative, non-linear, gender- and cohort-specific function of potential labor market experience. To that end, assume that the mean earnings of men after *e* years of potential experience are given by:

$$y(e) = Y^0 + \int_0^e PSI^M(a)da\,, \tag{A1}$$

where Y^0 gives the entry wage and $PSI^M(a)$ gives the rate of post-schooling investment at age a. Similarly, assume that for women of cohort c:

$$y_c(e) = Y^0 - G_c + \int_0^e PSI_c^F(a)da\,, \tag{A2}$$

where G_c gives women's earnings disadvantage at labor market entry. Together, (A1) and (A2) imply that women's relative wages are given by:

$$RW(e,c) = 1 - G_c - \int_0^e \left(PSI^M(a) - PSI_c^F(a)\right)da. \tag{A3}$$

Letting c=0 be the reference cohort, it follows that:

$$RW(e,c) = (1-G_0) + (G_0 - G_c) - \int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da + \int_0^e \left(PSI_c^F(a) - PSI_0^F(a)\right)da \tag{A4}$$

The first term, (1-G_0), in equation A4 corresponds to the intercept, α, in our estimating equation, (1). The second term, (G_0- G_c), corresponds to the cohort effects (γ_c) in (1); the third to the "baseline" age effects (β_a) on women's relative wages. Finally, the θ_c parameters in equation (1) approximate the average value (over all ages) of $(PSI_c^F(a) - PSI_0^F(a))$, i.e. the cohort-age interactions at the heart of the pure PSI model. This allows us to test the argument underlying the PSI model, i.e. that $PSI_c^F(a)$ is higher, on average, for recent cohorts of women, relative to earlier cohorts at the same age.

We are now in a position to add a time trend, trend(t), to equation (A4), and to explore its impact on our estimates of PSI effects. To do this, we normalize units of measurement so that time (t) is measured in decades, and $e+c=t$. After adding the trend, equation A4 becomes:

$$\begin{aligned} RW(e,c) = (1-G_0) + (G_0 - G_c) - \int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da \\ + \int_0^e \left(PSI_c^F(a) - PSI_0^F(a)\right)da + trend(e+c) \end{aligned} \tag{A5}$$

Suppose first that the trend is a linear function, i.e. ($trend(t)=bt=be+bc$), with $b > 0$. In this case, (A5) can be rewritten:

$$\begin{aligned} RW(e,c) = (1-G_0) + (G_0 - G_c + bc) + \left(be - \left[\int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da\right]\right) \\ + \int_0^e \left(PSI_c^F(a) - PSI_0^F(a)\right)da \end{aligned} \tag{A6}$$

Thus, our estimated cohort effects, (G_0- G_c + bc), will overstate the true cohort effects, (G_0- G_c). Likewise, the estimated baseline age effect, $\left(be - \left[\int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da\right]\right)$, will understate the extent to which women's PSI accumulation tends to fall behind men's, $-\left[\int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da\right]$. Critically, however, a linear trend will not affect our estimated PSI effects (*pe*cohort* interactions) at all. In sum, the presence of a true, positive linear time trend (obviously) affects the relative role of time versus cohort effects in explaining the narrowing of the gender wage gap (raising the importance of time and reducing that of cohort). It also biases our estimated age coefficients upward relative their true values. That said, the presence of such a time trend leaves the total relative wage change attributable to pure time and cohort effects together unchanged, and leaves the contribution of cohort-year interactions relative to these two alternative mechanisms unchanged as well.

Finally, assume a more general case where the trend is not necessarily linear. The effect of any monotonically increasing trend can be seen most clearly if we decompose the trend in the following way:

$$trend(e+c) = trend(c) + trend(e) + \left[trend(e+c) - trend(c) - trend(e)\right] \qquad \text{(A7)}$$

Where $trend(c)$ represents the value of the trend at cohort c entry, and $trend(e)$ represents the value of the trend when the baseline cohort reaches age e, relative to the value at t=0.

Then:

$$\begin{aligned} RW(e,c) = (1-G_0) + (G_0 - G_c + trend(c)) + \left(trend(e) - \left[\int_0^e \left(PSI^M(a) - PSI_0^F(a) \right) da \right] \right) \\ + \left(\left[trend(e+c) - trend(c) - trend(e) \right] + \int_0^e \left(PSI_c^F(a) - PSI_0^F(a) \right) da \right) \end{aligned} \qquad \text{(A8)}$$

Again, both cohort and baseline age effects are upward biased. But here, the estimated PSI effects (*pe*cohort* interactions) will be biased when the trend is not linear —upward when the trend is accelerating, and downward when the trend is decelerating.[36]

[36] To see this simply recall that acceleration means $f(x+\varepsilon) > f(x) + f(\varepsilon)$.

APPENDIX 3: Simulations

According to Appendix 2, our estimates of *pe*cohort* interactions in Tables 2 and 4 are biased upward when there is an accelerating time trend in women's relative wages, and downward when the trend is decelerating. Further, the data on entry-level wages in Figure 5 shows acceleration in the early years of our data and deceleration later. Unless this trend is purely cohort-specific, it follows that our *pe*cohort* interactions are biased upward for old cohorts, and downward for more recent cohorts.

The simulations presented in Table 5 estimate the size of this bias under a range of hypotheses about the nature of the true time trend. Based on the empirically observed time trend among entry-level workers (depicted in Figure 5), define:

$$T(t)= \begin{cases} 0.0 | year \in \{1959,1969\} \\ 0.1 | year = 1979 \\ 0.2 | year \in \{1989,1999\} \end{cases}$$

Based on the relative wage equation (A5), evaluated at e=0, $T(c)$ approximates the time trend in relative wages at labor market entry:

$$T(c) \approx RW(0,c) - RW(0,0) = (G_0 - G_c) + trend(c) \qquad \text{(A9)}$$

Since it is impossible to identify the relative contributions of true cohort-specific effects $(G_0 - G_c)$ and year-specific time trends $(trend(c))$ to $T(t)$, we vary the relative importance of each in a series of simulations. Suppose that (for some value of $\alpha \in [0,1]$):

$$(G_0 - G_c) = (1-\alpha)T(c) \qquad \text{(A10)}$$

Then:

$$trend(c) \approx \alpha T(c) \qquad \text{(A11)}$$

In this specification, α describes the proportion of $T(c)$ that is due to year effects, rather than cohort effects.

Returning to relative wage equation (A5), the dependent variable for observations at all ages can be "de-trended" using the following transformation:

$$
\begin{aligned}
RW(e,c) - trend(e+c) - (G_0 - G_c) &= RW(e,c) - \alpha T(e+c) - (1-\alpha)T(c) \\
&= (1-G_0) - \int_0^e \left(PSI^M(a) - PSI_0^F(a)\right)da + \int_0^e \left(PSI_c^F(a) - PSI_0^F(a)\right)da
\end{aligned}
\qquad \text{(A12)}
$$

If the correct value of α is known, regression of this de-trended dependent variable on age controls and pe*cohort interactions will yield unbiased estimates of age and PSI effects. Of course, we do not know the true value of α. The best we can do is to see how the PSI estimates (and age-earnings profiles) change as different values of α are simulated. Columns 1-5 of Table 5 report the results of simulations where α takes values 0, 0.25, 0.50, 0.75, and 1.