CHAPTER IV

THE IMPACT OF WOMEN'S LABOR SUPPLY ON THE DISTRIBUTION OF WAGES AND THE GENDER GAP, 1960 TO 2000

From 1970 to 1990, the labor-force participation of younger, U.S. women increased by more than at any other time during the Twentieth Century. Among women ages 25 to 35 market participation at the extensive margin rose by 30 percentage points, a 70 percent increase. These women chose different careers, acquired more formal education, and worked more hours per week and more weeks per year on average than their predecessors. Over the same period, the U.S. experienced the largest increase in wage inequality of any developed country. The difference in log wages between the 90th and 10th percentile rose by roughly 60 log points among full-time, full-year men (Juhn, Murphy, Pierce 1993; Card and DiNardo 2002).[70] Because few sources of credibly exogenous variation in labor supply exist, economists have only a limited understanding of the impact of the profound changes in women's labor supply on male wage inequality in the U.S. over the last 40 years.

Understanding the relationship of these two phenomena is important for a number of theoretical and practical reasons. First, economists have proposed a variety of theoretical models that relate changes in the relative supply of workers to changing wage inequality among men (e.g., Bound and Johnson 1992; Berman, Bound, and Griliches, 1994; Autor, Katz, and Krueger, 1999). Empirical studies may be seen as a test of these models and a means to identify the magnitude of their parameters.

Second, a direct test of the impact of labor supply shifts may inform the empirical literature that has emphasized the importance of skill biased technical change (SBTC) in exacerbating earnings inequality during the 1980s (Bound and Johnson 1992; Topel 1994; Katz and Murphy 1992; Juhn, Murphy Pierce et al. 1993; Juhn and Kim 1999; Autor, Katz, and Krueger 1999). This literature has abstracted from changes in labor supply (or has used endogenous measures to control for them) in order to identify important shifts in the structure of demand for labor. By examining the same question from the perspective of labor supply, empirical studies provide insight as to the validity of the assumptions in this line of research and the generality of its conclusions. Moreover, they may provide useful insight into the gender-bias of technological change over the 1980s (Weinberg 2000;

[70] The measure of wage inequality used to generate this statistic is the variance of the log of weekly wages.

Card and DiNardo 2002).

Third, research on the link between increases in women's market work and their wage growth may shed light upon the long-awaited and abrupt decline in the gender wage gap during the 1980s. Research on the convergence in earnings between men and women has almost exclusively treated changes in the distribution of wages as exogenous to the determinants of women's labor-force participation. For instance, Blau and Kahn (2000) assume that "male wage inequality is determined by forces outside the gender pay gap and is a useful indicator of the price of skills affecting both men and women" (p. 88; see also O'Neill 2003). A test of how increases in women's labor market participation affected the distribution of wages may help gauge the degree to which these studies mis-state changes in the gender gap.[71]

While increases in women's labor-force participation may have caused changes in the distribution of wages, the complicated interrelationship of these two phenomena has limited economists' ability to quantify their effects. Women may have entered the workforce to offset absolute and relative declines in the earnings of their husbands. Another (and not unrelated) explanation emphasizes that gender-biased technical change raises the earnings power of women relative to men and may have induced more households to send women to work outside the home. Finally, an increase in women's labor supply may have actually caused, among other things, the real wages of men to decline as women competed for similar jobs (Topel 1994, 1997; Fortin and Lemieux 1998; Acemoglu et al. 2004).

This chapter explores the impact of relative shifts in women's labor-force participation on the distribution of wages after 1970 using an instrumental variable for women's labor supply. The instrument is constructed using cross-state variation in fertility rates from 1940 to 1956 and variation in the timing of legal changes that allowed younger, unmarried women to obtain oral contraceptives. Research suggests that these legal changes provide exogenous variation in access to the "pill" and that they are strongly associated with changes in women's labor-force participation (Goldin and Katz 2000, 2002; see chapter III).

Using the 1960 to 2000 Integrated Public Use Data Microdata Series (IPUMS), I find strong

[71] Studies of the gender gap in wages tend to take the distribution of wages as the exogenously given price of skill. For instance, Blau and Kahn (1997, 2000) among others (Katz and Autor 1999, Card and DiNardo 2002) suggest that the gender gap should have risen over the 1980s because the returns to skill increased. The argument is that men have more labor market experience on average, so increases in the returns to experience should have benefited men more than women across the distribution. This argument assumes that changes in the age-experience composition of the labor force did not affect the returns *per se*. While Blau and Kahn (2000) conclude that supply shifts probably dampened the growth in women's wages over the 1980s, there is little way to assess the relative magnitude of this effect.

evidence that the instrument is both valid and relevant in the context of this study. Even after allowing for a variety of fixed effects, the instrument remains strongly correlated (and statistically significant at the one percent level) with shifts in the number of weeks worked by women relative to those of men from 1960 to 2000. I analyze the impact of changes in the relative supply of women on the wages of women, men and the wages of women relative to those of men across the distribution using this instrument.

The core results are that relative shifts in the labor supply of women during the 1980s tended to *increase* the wages of the most skilled men. In contrast to research that has emphasized female labor inputs as substitutes for less skilled men, my results suggest that increases in women's work tended to increase male earnings at the 90th percentile and leave earnings at the median unaffected. This finding implies that women's skills may have been complementary to the productivity of men at the upper end of the skill distribution. Second, I find that relative increases in women's labor supply strongly mitigated the wage gains of women at the mean during the 1980s. The gender gap in wages would have been roughly half as large in 1990 in the absence of the epochal changes in women's labor supply—a particularly remarkable figure as the 1980s experienced the largest declines in the gender gap since World War II.

The Potential Effect of Labor Supply Shifts on the Gender Gap and Skill Premia

The relatively small body of literature on this topic is evidence of the difficulty in understanding the interrelationship of labor supply shifts and changes in earnings dispersion. Moreover, existing research in this area tends to reach diverging conclusions about this relationship as well. Topel's seminal article claims there is "no evidence that different regional evolutions of wages are demand driven" and suggests that "whole story [of regional changes in wage inequality] is on the supply side" (1994: 20).[72] His results imply that, in the absence of increases in women's labor-force participation, the wages of less skilled men would *not* have changed over the 1980s. Because the median of women's wage distribution falls at about the 25th percentile of men's distribution, he suggests that skilled women compete more with lower skilled men. In his 1997 article, however, Topel cautions that the patterns of substitution required for his earlier argument to hold finds limited support in the data.[73] In fact, less skilled men and women work in very different

[72] His empirical model uses cross-regional variation in the predicted relative supplies of workers at various percentiles in the male wage distribution.
[73] A similar model links the changes in the gender gap to changes in the distribution of wages. Fortin and Lemieux (2000) hypothesize that the distribution of wages has remained unchanged and explore how far observed changes in

occupations and industries.

Juhn and Kim (1999) address this concern by examining the characteristics of women in the paid market from 1970 to 1980 and find little evidence that an increase in the relative supply of women exacerbated wage inequality among men. On the contrary, they argue that increases in women's market work may have stemmed the tide of rising wage inequality over the 1980s. First, the growth in women's labor-force participation slowed over the 1980s at the same time that the economy witnessed large increases in male wage inequality. Second, over the 1980s women added to the supply of skilled labor, which tended to dampen wage growth in the higher skill percentiles of men rather than bid down the wages over lower skilled men.

While the direct evidence from the 1970s and 1980s is mixed, a recent study of the 1940s by Acemoglu, Autor and Lyle (2004) sheds new light on this debate. Using cross-state variation in World War II mobilization rates as an exogenous shifter of women's labor supply, they find that the large increase in women's work temporarily reduced both male and female wages during the 1940s. This analysis is most similar in methodology to their work, but the context of the 1940s and the nature of women's supply shifts during this decade are distinctly different from the post-1960 period in at least three respects.

First, the wage structure narrowed considerably during the 1940s (Goldin and Margo 1992). In contrast, the 1970s and 1980s recorded the largest increases in the dispersion in earnings of the Twentieth Century (Topel 1997, Katz and Autor 1999). Similar labor supply shifts in each period may, therefore, result in substantially different patterns of substitution across and between skill groups and, by extension, have different impacts on the dispersion of wages.

Second, the technological structure of the economy in the 1980s—the diffusion of micro-computers, photocopiers and other office equipment—tended to enhance worker productivity differentially and especially in occupations where women were concentrated (for instance, among clerical workers rather than manufacturing operatives).[74] As a result, if, for instance, clerical workers are complementary to the productivity of highly skilled men, women's entry may have raised the wages of men at the upper end of the distribution rather than bidding down those of the less skilled. This suggests a pattern of wage effects similar to those proposed by Grant and Hammermesh (1981),

---

the wage structure might be explained by women displacing men in the skill hierarchy. However, once the distribution is taken as fixed, the only explanation for changing dispersion is on the supply side.

[74] Card and DiNardo (2002) note that women were, on average, roughly 40 percent more likely to be using computers at work than men in 1984. Women with less than a college degree were more than twice as likely to be using a computer at work than men with less than a college degree, and this gap changed little over the course of the 1980s.

Topel (1994) and Weinberg (2000), although the underlying cause is different. A decline in the wages of unskilled relative to skilled men reflects complementarity of women with high skilled men rather than their substitutability with unskilled workers.

Third, as noted by Juhn and Kim (1999) the skills of women in the workforce in the 1980s were much different than the skills of the entrants during World War II. Wartime mobilization during the 1940s was associated with large but *temporary* increases in women's labor supply. Approximately 50 percent of the wartime entrants had exited the market by 1950 (Goldin 1991). The 1980s were associated with large and *permanent* shifts in women's work. Women not only chose to work more at a point in time, but they remained in the workforce acquiring more experience and on-the-job skills (Fortin and Lemieux 1998, Blau and Kahn 2000, O'Neill 2003).

This is the first study to attempt to quantify the impact of labor supply shifts in the 1970 to 1990 on the structure of wages using an instrumental variables approach.

A Simple Theoretical Framework

The discussion can be framed with a simple extension of a model used in the SBTC literature (Bound and Johnson 1992; Autor, Katz and Krueger 1999; Card and DiNardo 2002; Acemoglu, Autor and Lyle 2004). Aggregate output, Y, is produced using the skilled, S, and unskilled, U, labor of men and women, W. The production function can be written

$$Y = A\left\{[s\,(a_s S_m)^{\mu} + (1-s)(cW)^{\mu}]^{\rho/\mu} + u(a_u U_m)^{\rho}\right\}^{1/\rho} \tag{1}$$

where A is a total factor productivity parameter, s and u are the share [0 to 1] parameters, and a and c are labor-augmenting productivity parameters that vary across skill and over time.[75] The substitution parameters, ρ and μ, are assumed to be less than or equal to 1. The aggregate elasticity of substitution between S and U is $\sigma_{su} \equiv 1/(1-\rho)$. The elasticity of substitution between skilled men and women is $\sigma_{mw} \equiv 1/(1-\mu)$.

Assuming that workers are paid the value of their marginal products, the skill premium paid to men can be written as

$$\omega_{SmUm} \equiv w_{Sm}/\,w_{Um} = [\mathrm{s} / \mathrm{u}]\,[\,a_s^{\,\mu} / a_u^{\,\rho}]\,[S_m^{\,\mu-1}/\,U_m^{\,\rho-1}]\,[s\,(a_s S_m)^{\mu} + (1-s)(cW)^{\mu}]^{(\rho-\mu)/\mu}\,. \tag{2}$$

[75] I abstract from capital inputs into the production process for simplicity, but see Acemoglu, Autor and Lyle (2004) for a simple extension of this model. Time subscripts are omitted for notational simplicity.

By the same logic, the gender gap in pay between skilled men and women can be expressed as

$$\omega_{SmW} \equiv w_{Sm}/\, w_W = [\mathrm{s} \,/\, (1\text{-s})]\, [\, a_{\mathrm{s}} / \, c]^{\,\mu}\, [S_m/W]^{\,\mu\text{-}1}\, . \tag{3}$$

The wage premium between women and unskilled men is

$$\omega_{UmW} \equiv w_{Um}/\, w_W = [\mathrm{u} \,/(1\text{-s})]\, [\, a_{\mathrm{u}}^{\,\rho} / \, c^{\,\mu}\,]\, [U_m^{\,\rho\text{-}1}/\, W^{\mu\text{-}1}]\, [s\, (a_s\, S_m)^{\mu} + (1\text{-}s)(cW)^{\mu}]^{(\rho\text{-}\,\mu\,)/\mu} \tag{4}$$

This simplified framework makes salient the relationships of skill and gender-biased technological progress as well differential shifts in the skill and gender composition of workers. It is easy to see that a change in the skill premium may represent changes in workers' relative productivities (a_{s}, a_{u}, c, s, or u) or their relative numbers (S_m, U_m, W) or both.

Technological progress may affect the aggregate gender gap (Acemoglu, Autor and Lyle 2004; Card and DiNardo 2002, Juhn and Kim 1999), as well as the skill premium among men. Previous empirical work has largely focused on changes in technological parameters with an emphasis on the wage structure of men. The consensus emerging from this work is that expanding wage inequality is due in large part to skill biased technical change, or relative increases in a_{s} or s (Autor and Katz 1999; Card and DiNardo 2002).

Recent work also suggests that technological progress may be gender-biased as well. For instance, the introduction of micro-computers may have enhanced women's productivity relative to that of men (Weinberg 2000, Welch 2000, Card and DiNardo 2002). If personal computers were more complementary, say, for clerical workers than for machine operatives and women were more concentrated in clerical occupations, then the introduction of personal computers may have increased c relative to a_{s} and a_{u} (although not necessarily proportionately). Weinberg (2000) suggests that computers changed the productivity of women relative to blue-collar male workers more (c relative to a_{u}) than they changed the productivity of women relative to skilled men (c relative to a_{s}). While changes in these parameters are clearly important in changing the overall structure of wages, I focus on the potential impact of relative supply shifts.

Shifts in the relative supply of workers—women relative to men and skilled relative to unskilled workers—might affect the aggregate and skill-specific gender wage differentials as well (Topel 1994, 1997; Fortin and Lemieux 1998, 2000; Juhn and Kim 1999). The sign and magnitude of the impacts will reflect both the size of the relative changes and the relationships of the substitution parameters ρ and μ, which may or may not be fixed over time. The model yields simple comparative statics on the effects of an exogenous increase in the number of labor supplied by women on the gender gap in wages as well as male wage inequality. It is straightforward to show that the sign of the effect of an increase in women's labor supply on the

1. gender gap in wages is given by sign $<\partial\omega_{SmW}/\partial W>$ = sign$<\mu\text{-}1>$;
2. male skill premium is given by sign $<\partial\omega_{SmUm}/\partial W>$ = sign$<\rho\text{-}\mu>$, and
3. wage differential between unskilled men and women is given by

$$\text{sign} <\partial\omega_{UmW}/\partial W> = \text{sign}< 1 - \mu > \times \omega_{SmUm} + \text{sign}< \rho - \mu > \times \omega_{SmW}.$$

If skilled men and women are not perfect substitutes, $\mu<1$, then an increase in the number of skilled women will unambiguously increase the gender gap in wages, because more skilled women will bid down their wages in equilibrium.

The effect of this change on the male skill premium, however, depends upon the relationship of ρ and μ. For instance, if the degree of substitution between skilled and unskilled men, ρ, is smaller than the degree of interchangeability between skilled men and women, μ, then an increase in the supply of skilled women will decrease wage inequality among men. When $\rho < \mu$, skilled women compete with skilled men (rather than unskilled men) and tend to bid down their relative wages. Moreover, an increase in women will unambiguously decrease the premium they earn over unskilled men.

On the other hand, if skilled women are more competitive with unskilled men ($\rho > \mu$), an increase in skilled women will increase wage inequality among men by driving down the wages of unskilled men. The first term, $< 1 - \mu > \times \omega_{SmUm}$ is negative as long as women and skilled men are not perfect substitutes. However, the sign of the second term depends both on the relative substitutability of women with skilled and unskilled men (ρ - μ), and the magnitude of the skill premium, ω_{SmUm} relative to gender gap in wages between skilled men and women, ω_{SmW}. Therefore, the impact of a relative supply shift of women on the gap between their wages and those of unskilled men is theoretically ambiguous.

Trends in Wage Inequality and Women's Labor-Force Participation

I use the one percent samples of the U.S. population from the Integrated Public Use Microdata Series (IPUMS) (Ruggles et al. 2004). The sample includes all U.S.-born individuals between the ages of 16 and 60 years old and who are not residing group quarters. Following the empirical literature, I also restrict the sample to individuals who report plausible earnings. Because I use weekly earnings as my measure of wages, I limit the sample to the equivalent in weekly terms which I count as $40 to $4000 in weekly earnings on average annual basis (Card and DiNardo 2002).

Weekly wage estimates were obtained by dividing the respondent's earnings in the previous calendar year by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. Rather than relying on interval midpoints, I estimate actual weeks

Table 1. IPUMS Sample Statistics, 1960 to 2000

	Census year				
	1960	*1970*	*1980*	*1990*	*2000*
A. Men					
Weekly earnings	446.3	559.0	575.8	553.5	592.6
	(58.7)	(61.8)	(54.6)	(60.8)	(72.5)
Weeks worked	45.7	45.7	45.3	45.7	46.7
	(0.829)	(0.636)	(0.602)	(0.611)	(0.558)
In labor force [a]	0.899	0.868	0.864	0.872	0.834
	(0.021)	(0.022)	(0.020)	(0.021)	(0.026)
Observations	421,422	474,681	598,492	629,739	689,195
B. Women					
Weekly earnings	257.0	315.3	316.5	354.4	393.1
	(29.9)	(30.9)	(26.3)	(47.5)	(51.3)
Weeks worked	37.6	38.1	40.5	43.0	44.4
	(1.39)	(1.05)	(0.934)	(0.759)	(0.698)
In labor force [a]	0.404	0.486	0.609	0.714	0.722
	(0.035)	(0.031)	(0.036)	(0.036)	(0.036)
Observations	453,029	513,091	631,217	663,962	724,559
C. Instrument					
Early legal access [b]	---	0.048	0.169	0.152	0.126
		(0.048)	(0.169)	(0.152)	(0.126)
PUMS sample	General	Form 2 State	5% State	1% Unweighted	1% PUMS

The statistics include all native born individuals residing in the U.S. between the ages of 16 and 60, who earned on an average annual basis between $40 and $4000 per week in 1989 dollars. Wages are expressed in 1989 dollars in each census year. A one percent sample is constructed from the 1980 IPUMS five percent state sample by taking an unstratified twenty percent sample. The standard errors of the state-time cell means are in parenthesis below the means. All estimates use census population weights. [a] "In labor force" is not used in the analysis but is included for reference. It is the fraction of all workers—not the workers meeting the sample requirements—observed in the census that are coded as being in the labor force. [b] "Early legal access" represents the fraction of women born from 1940 to 1956 that would have had access to the pill if they were unmarried and younger than age 21 and were residing in their state of birth. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

worked using information on the distribution of work intensity within each interval from the 1980 census. I also apply a new methodology to correct for the top coding of earnings. Instead of applying a uniform correction across all states and years as is usual in the literature, I exploit variation in the real value of the top code across years and make state specific adjustments based on information from the 1970 census.[76] (Appendix G provides more details on the construction of variables used in the analysis and corrections for earnings top codes). Population weighted means of key variables in the analysis are reported in Table 1 by census year separately for men and women.

Because studies of wage inequality have relied on annual figures of the *Current Population Surveys* (March Supplements, 1963-1988, Bound and Johnson 1992; Katz and Autor 1992; Juhn et al. 1993; Outgoing Rotation Supplements, 1979-2000, Card and DiNardo, 2000; DiNardo, Fortin and Lemieux, 1996) and because it is well known that different data sources and different sampling restrictions can significantly alter measured inequality, I begin by documenting the similarity of the

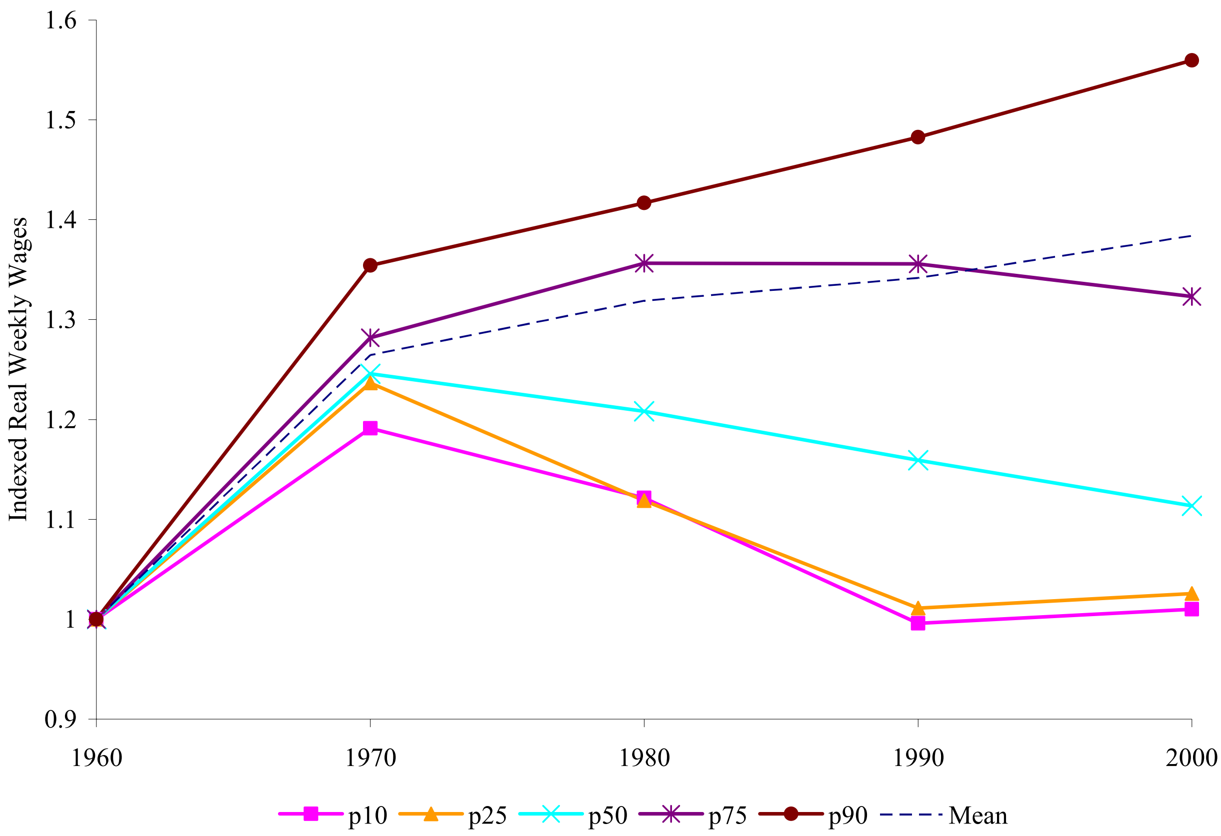

Figure 1. Indexed Real Weekly Wages of Men by Percentile, 1960 to 2000

Source: Author's calculations using the 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

---

[76] The 1970 census had the highest real value of the earnings top code relative to the other census years from 1960 to 2000. It provides information on the distribution of earnings above lower top codes in the other census years.

trends observed in my IPUMS samples with those in other studies.

Figure 1 plots the evolution of percentiles in the male wage distribution relative to their real 1960 value for each decade from 1960 to 2000. As noted in Juhn et al. (1993), wages across all percentiles rose by roughly 30 percent over the 1960s (this is consistent with studies using the CPS). Over the next two decades, the erosion of wages among less skilled male workers was so severe that wage earners below the median earned less in real terms in 1990 than they had in 1960. During the same period, the wages of highly skilled workers grew to roughly 140 percent of their 1960 levels. By the mid 1990s the U.S. had one of the highest rates of wage inequality of any developed country (Katz and Autor 1999).

As graphed in Figure 2, wage growth among women proceeded more evenly across the distribution. Indexed real weekly wages for women had grown to at least 130 percent of their real 1960 value in *every* percentile. As men's wages below the median eroded over the 1970s, the wages of less skilled women not only increased but grew at a faster rate than the wages of women at the 75th and 90th percentiles (which stagnated). During the 1980s, these trends reversed: The real wages of

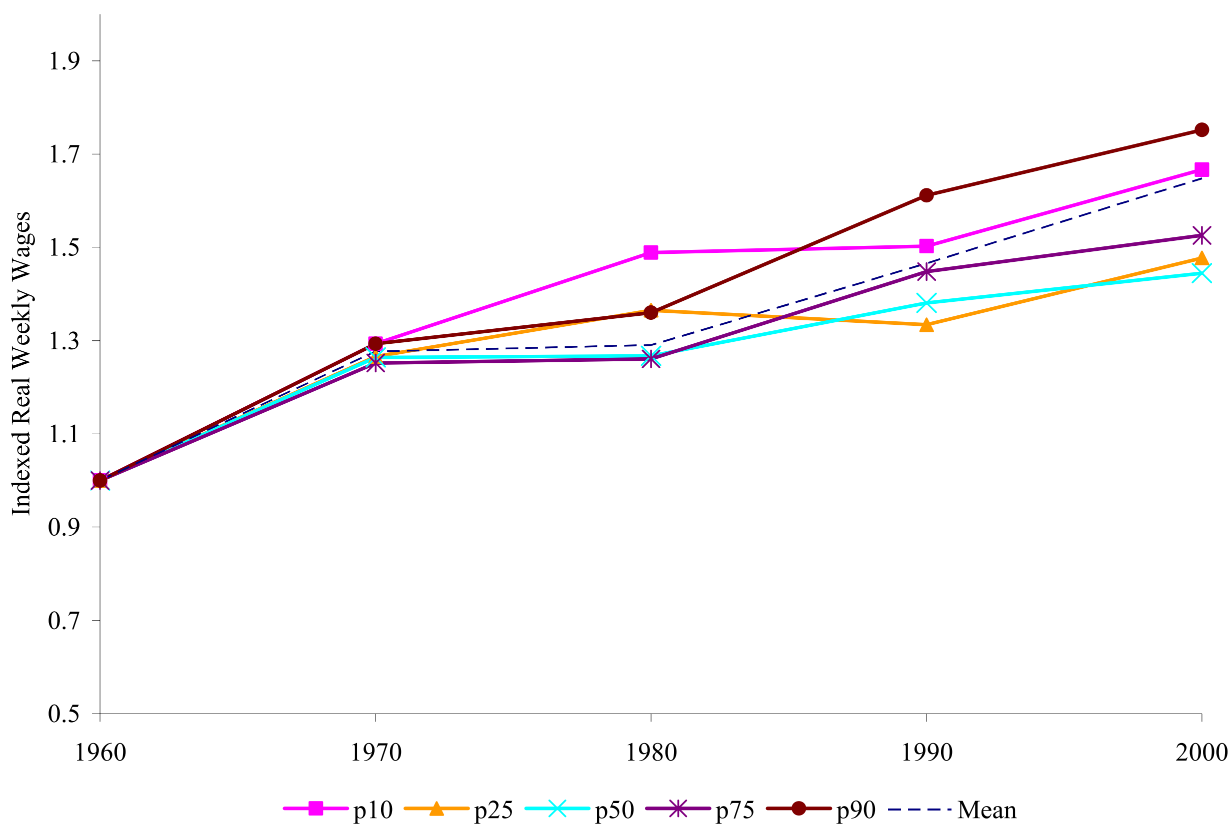

Figure 2. Indexed Real Weekly Wages of Women by Percentile, 1960 to 2000

Source: Author's calculations using the 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

more highly skilled women grew more rapidly than those of workers below the median.

The net effect of these differing trends on changes in the relative wages of women, or the gender gap, across the distribution can be seen in Figure 3. Figure 3 plots the ratio of weekly wages of women relative to men from 1960 to 2000. In 1960 women at the mean (the dashed line) earned roughly 55 percent of what men earned on an average annual basis. In 2000 the same figure was roughly 65 percent (for the adjusted figures see O'Neill 2003). It is well known that differences in the mean and the median between women and men did not begin to increase until the early 1980s (cf. Goldin 2002, Blau and Kahn 1997). However, these trends mask important changes in the lower percentiles. The wages of women below the median relative to those of men at the same percentiles increased rapidly over the 1970s. The general upward trend indicates that the gap at each percentile in the male and female wages narrowed in the last forty years of the Twentieth Century, although the rate of convergence slowed in the 1990s (cf. O'Neill 2003).

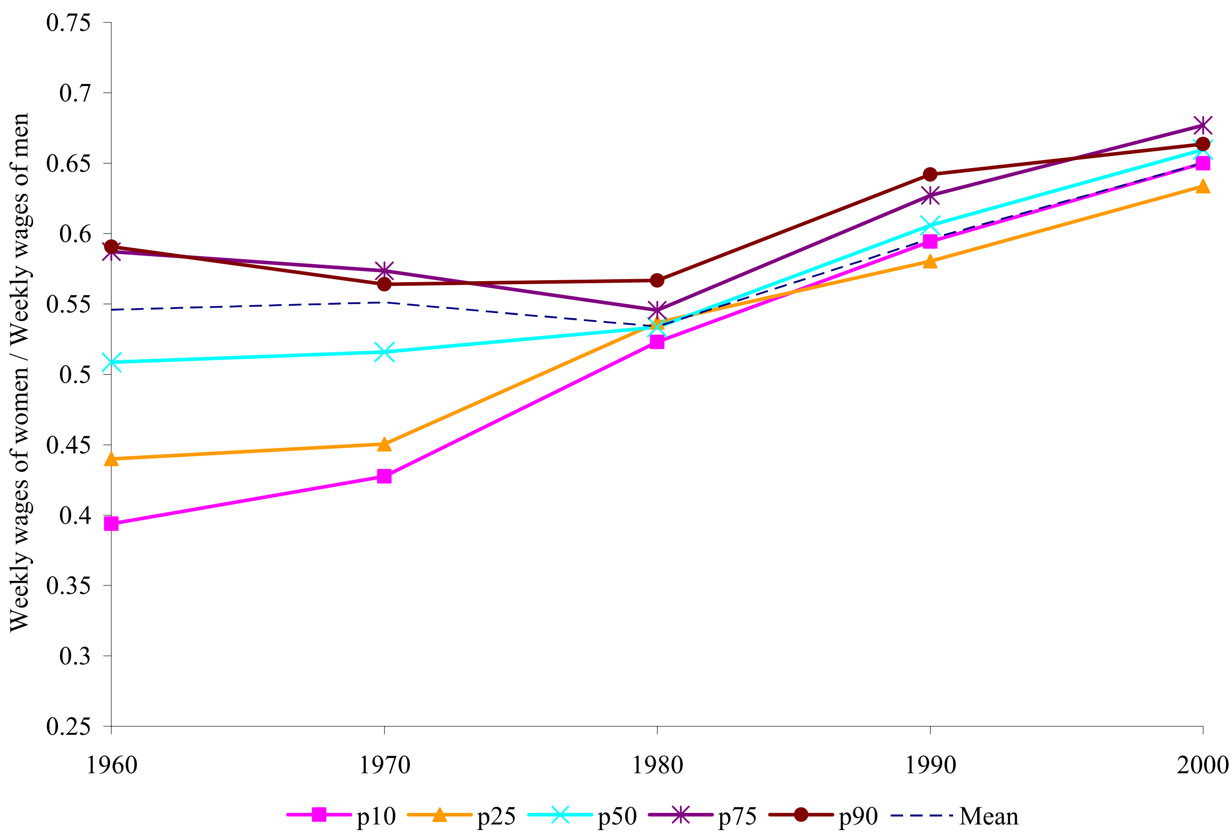

Figure 3. Changes in the Gender Gap in Weekly Wages by Percentile, 1960-2000

Source: Author's calculations using the 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

The absolute and relative gains in women's wages occurred in large part due to the entry of younger and better educated cohorts (see chapter III). Between 1940 and 1960, women had changed their labor-force participation little at ages 26 to 40. After 1960, however, women dramatically increased their market work across age groups, and the growth in relative participation was strongest among women in the peak of their childbearing years, or ages 26 to 35 (cf. Smith and Ward 1984, Goldin 2002, chapter II). Figure 4 plots the fraction of women in the labor force at the extensive margin relative to the number of men by age group.

Because the instrument I propose in the next section predicts the labor-force participation rates of women born from 1940 to 1956, I overlay the averaged lifecycle participation profiles of these cohorts (dashed line) as they age across the census years for reference. The instrumental variables analysis examines the relationship of changes in the relative labor supply of women in these cohorts with changes in the dispersion of earnings from 1960 to 2000.

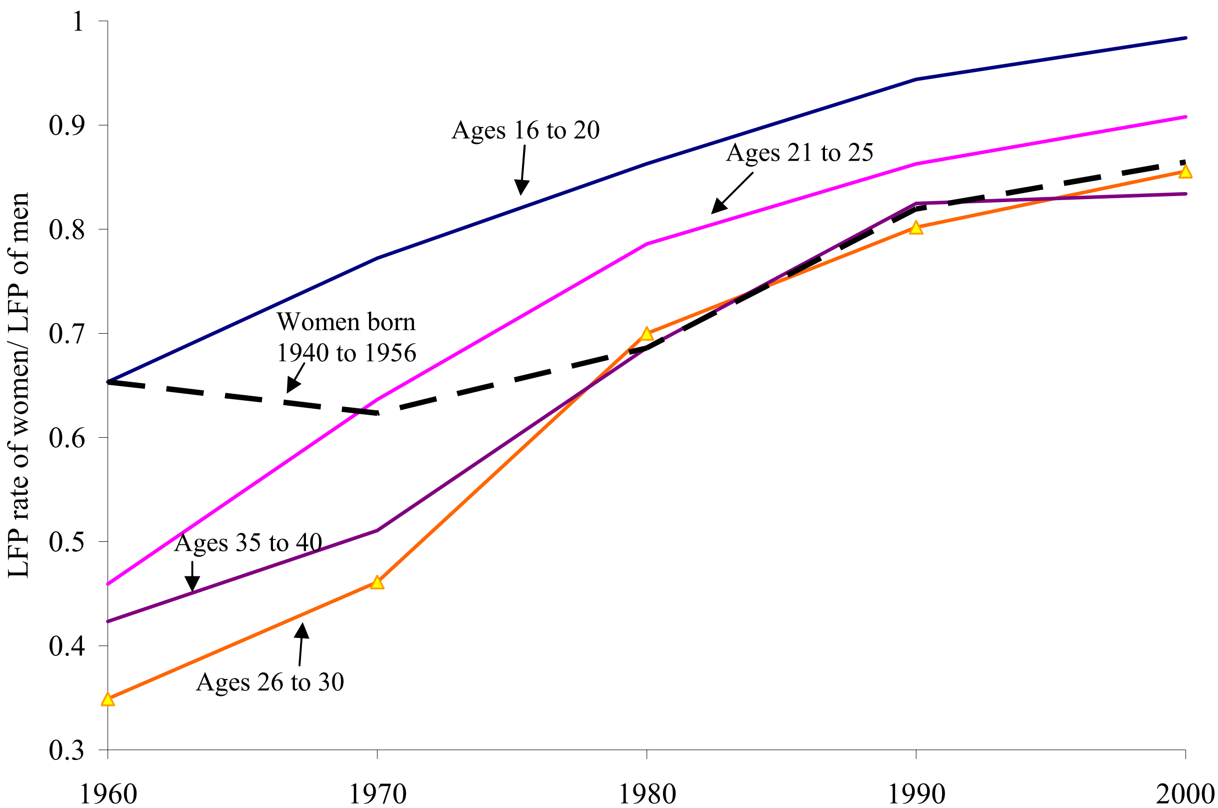

Figure 4. Labor-force participation rates of women relative to men by age group

The vertical axis plots the fraction of women in a particular age group participating in the labor-force relative to the fraction of men in the same age group who report being in the labor-force. Source: Author's calculations using the 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

The Impact of Relative Supply Shifts on the Structure of Wages

The theoretical framework suggests a simple econometric specification. A logged adaptation of equation (3) can be expressed as

$$\log \omega_{st} = \lambda_s + \lambda_s \cdot t + \lambda_t + \alpha \log(W_{st}/M_{st}) + \xi_{st} \tag{5}$$

where the dependent variable is a measure of earnings in state *s* at time *t*; the fixed effects, λ_s, $\lambda_s \cdot t$, and λ_t, denote dummy variables for fifty states, linear state-specific time trends, and dummy variables for the years 1970 to 2000, respectively; $\log(W_{st}/M_{st})$ denotes the log of a measure of the labor supply of women relative to that of men. This specification implicitly assumes that states can be treated as separate labor markets over ten year periods. It is certainly true that cross-state migration and capital mobility will tend to mitigate any effects in the longer run. Therefore, α may be regarded as a short-run elasticity of substitution parameter.

Before turning to the instrumental variables estimates, Table 2 documents the correlation between relative changes in the employment of women and weekly wages using ordinary least squares (OLS) regression. The dependent variable is the mean of weekly wages in the first column and is the difference in the log of the wages at given percentiles in the remaining three. Each of the population-weighted models controls for state, year and state-linear time trends. I omit specifications with only state and year fixed effects to save space, but the results are qualitatively identical. Among both women (Panel A) and men (Panel B) these OLS estimates provide little or no evidence that shifts in the relative supply of women impacted the dispersion of wages. None of the point estimates is even marginally statistically significant. The pattern of results suggests that relative supply shifts may have reduced earnings at the mean and increased wage inequality as measured by the range statistics. The point estimates are larger for women than for men, which is consistent with the notion that women may be better substitutes for other women than for men.

The causal effect of a relative increase in women's labor market participation, α, is only identified for exogenous shifts in the relative supply of women. The estimates in Table 2 are likely biased because observed employment rates and changes in them are functions of market wages, or

$$\log(W_{st}/M_{st}) = \eta \log\omega_{wm,\, st} + \log P_{st} \tag{6}$$

Table 2. OLS Estimates of the Impact of Relative Supply Shifts on the Distribution of Wages, 1960 to 2000

	Dependent variable: Log(Weekly Wages)			
	A. Sample: Women in years 1960-2000 (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.247	0.035	0.645	0.365
	[0.168]	[0.081]	[0.625]	[0.260]
R^2	0.900	0.800	0.630	0.580
Fixed effects	S,Y,S·Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y
	B. Sample: Men in years 1960-2000 (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.011	-0.063	0.155	0.091
	[0.154]	[0.165]	[0.174]	[0.223]
R^2	0.88	0.90	0.92	0.71
Fixed effects	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y

The dependent variable is the mean or indicated difference in the percentiles in the distribution of the log of weekly wages. Weekly wages are obtained by dividing last year's earnings by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked and correct for top-coding using the methodology described in Appendix A. Log weeks worked by W/M is the log of the number of weeks worked by women relative to the weeks worked by men in the previous calendar year. All regressions are weighted using population weights. Heteroskedasticity robust standard errors are in brackets beneath the point estimates. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included. S•Y indicates that the set of state dummies, S, is interacted with a linear time trend. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

where η is the elasticity of skilled women's labor supply with respect to the relative market returns to work, ω_{wm}, and P is the exogenously given population of women relative to men (cf. Bound and Johnson 1992). Therefore, OLS estimates α may yield misleading results about the impact of exogenous changes in employment on the structure of wages.

Exogenous Shifts in the Relative Supply of Women

In order to recover an unbiased estimate of α, I construct an instrument for relative shifts in women's labor supply, $\log(W_{st}/M_{st})$, using the product of two sources of plausibly exogenous cross-state and cross-birth cohort variation.

The first component of the instrument is the difference in legal access to newly available birth control technology—the pill. From 1960, when the first birth control pill was introduced, until 1976, when all unmarried women ages 18 or above by U.S. Supreme Court's *Danforth* decision gained legal access to abortion and contraception, states varied significantly in their legal policies governing access to birth control. As I show in chapter III and as documented in previous research, the timing of these changes appears unrelated to differences in the levels of women's wages across states shortly before the pill was introduced (Goldin and Katz 2000, 2002).

The second component of my instrument is variation in birth rates across states from 1940 to 1956. Cross-state differences in the annual number of births reflect the childbearing decisions of the parents of the Baby Boom generation and occurred at least fifteen to thirty years before the expansion of earnings inequality that emerges after 1970. Therefore, they are plausibly exogenous to skill-biased changes in market demand from 1970 to 2000. Using these two sources of variation, I construct the instrument "early legal access", or ELA, as

$$ELA_{bt} \equiv \Big[\sum_{c=40}^{56} N_{bct}\Big]^{-1} \sum_{c=40}^{56} (Law_{bc} \cdot N_{bct}) \tag{7}$$

where N_{bct} is the number of women who born in state *b* in year *c* who are observed in the census in year *t*; Law_{bc} is a binary variable equal to one if state *b* allowed early legal access to birth cohort *c* (once legalized, no state repeals these laws). ΣN_{bct} is the sum of all individuals observed in year *t* that were born in state *b* from 1940 and 1956. Women born in these years are the only cohorts for which I have within-cohort variation in legal access. Moreover, without limiting the affected group, ELA increases with every cohort born after the date of legalization. These women do not add information to the analysis and may induce spurious correlations with the dependent variable. This variable also differs from the measure in chapter III, because ELA varies across census year and state of birth rather than by year of birth and state of residence (see Table 1 for sample statistics for ELA). ELA, thus, denotes the fraction of women born between 1940 and 1956 that enjoyed early legal access to contraception in each state and year in the sample.

Previous research supports the validity of this instrument (see chapter III and Goldin and Katz 2000, 2002), but I document its relevance for this study in Table 3. I estimate models of the following general form,

$$\log(W_{st}/M_{st}) = \delta_s + \delta_s \cdot t + \delta_t + \psi\, ELA_{st} + \xi_{st}\,. \tag{8}$$

The models include a constant as well as different combinations of fixed effects for between state differences, δ_s, differences in the evolution of state-level unobservables over time, $\delta_s \cdot t$, and differences across years, δ_t. The dependent variable is constructed as the log of the sum of the number of weeks worked in the previous calendar year by women relative to those worked by men in a particular state and point in time. The analysis is run on state of birth rather than state of residence (i.e. the $ELA_{bt} = ELA_{st}$) so as to abstract from the potential effects of demand induced migration on the wage dispersion.[77]

Panel A presents the results for these regressions using a sample of state-level observations pooled across years. The specification includes no fixed effects in the first column, state fixed effects in the second model, state and year fixed effects in the third column, and linear state time trends in the final column. Across all specifications ELA appears positively and significantly related to changes in the relative number of weeks worked by women. In every case, a Wald test rejects the hypothesis that the coefficient on ELA is equal to zero with F-statistics over ten.

Panel B splits the observations into pairs of years to examine the instrument's relevance in different periods. The first column presents the results for the years 1960 and 1970; the second column for 1970 and 1980; and so forth. The purpose is to examine the specific effect of ELA as the affected cohort ages. These regressions only include state fixed effects and one year dummy, because the number of observations is reduced to one hundred when considering only two points in time. The point estimate is statistically significant at the five percent level in the 1960s and at the one percent level in the 1980s. The large average effects reported in Panel A is driven by strong effects in two decades, rather than a uniform effect across the entire period. Therefore, the instrument will only be relevant for estimating the period effects in those two decades.

Is it reasonable for ELA to have different effects across years? The shape of the life-cycle labor-force participation profile of women born from 1940 to 1956 presented in Figure 4 corresponds closely to this pattern. Early legal access to the pill allowed this cohort to delay childbearing into the 1970s, so the supply of their market time relative to that of men declined (although not significantly so) during that period. Moreover, I document in chapter III that the labor-force participation effects wane as these cohorts of women reached their 40s and 50s. As a check on the robustness of this instrument, I also report the effects of ELA on the log of the mean weeks worked by men in Panel C. Changes in relative supplies might be driven by the changes in men's behavior, rather than changes among women. The only marginally significant coefficient is estimated in the model without any fixed effects. Although the effect is negative in each case, adding controls diminishes the magnitude and statistical significance of the correlation. In three of the four of the specifications, the correlation is weak and statistically indistinguishable from zero.

[77] By running the analysis on state of residence, I risk confounding the effects of differential migration on wage dispersion. For instance, if women with early legal access tend to differentially migrate to areas with higher rates of growth in the new skill-biased economy, then the IV estimates would confound the effects of increased labor-force participation with skill-biased technical change. Running the analysis on state of birth allows me to circumvent this source of bias.

Table 3. First Stage Estimates of the Impact of Early Legal Access on Women's Market Work, 1960 to 2000

	Dependent variable: Log ($weeks^W/weeks^M$)			
	A. Pooled years, 1960-2000 (N=250)			
ELA	1.463	2.684	0.319	0.375
	[0.247]	[0.285]	[0.095]	[0.103]
Adjusted R^2	0.24	0.4	0.97	0.98
F-statistic [a]	35.1	88.91	11.27	13.23
Fixed effects		S	S,Y	S,Y,S•Y
	B. Pairs of years, 1960-2000 (N=100)			
	1960-70	1970-80	1980-90	1990-2000
ELA	0.354	-0.182	0.851	0.42
	[0.133]	[0.299]	[0.268]	[0.318]
Adjusted R^2	0.91	0.94	0.91	0.79
F-statistic [a]	7.09	0.37	10.08	1.74
Fixed effects	S,Y	S,Y	S,Y	S,Y
	Dependent variable: Log ($weeks^M$)			
	C. Men only, 1960-2000 (N=250)			
ELA	-0.141	-0.01	-0.061	-0.051
	[0.072]	[0.041]	[0.098]	[0.041]
Adjusted R^2	0.09	0.77	0.83	0.92
F-statistic [a]	3.83	0.06	0.38	1.53
Fixed effects		S	S,Y	S,Y,S•Y

The dependent variable is the log of the sum of weeks worked by women divided by the sum of weeks worked by men in state *s* at time *t*. Weeks worked is the reported number of weeks worked in the previous calendar year. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked using the methodology described in Appendix A. ELA varies by state of birth and year. Robust standard errors are in brackets beneath the point estimates. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included. S•Y indicates that the set of state dummies, S, is interacted with a linear time trend. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

[a]The F-statistic reported is from a Wald test of the coefficient on ELA being equal to zero.

One concern in using this identification strategy is that cross-state variation in early legal access may be capturing other unobserved changes that affected the employment of women only, that were not the result of a labor-supply shift. One plausible hypothesis is that the negative coefficient are capturing the effects of the Vietnam War. Because the timing in changes in the age of majority, which permitted younger women to obtain the pill, are also correlated with casualty rates in Vietnam, ELA may reflect the impact of the war on men's labor-force participation.

Regressions not reported here suggest that this is not the case. I break the sample into decades as in Panel B to explore the evolution of the correlation across periods. If Vietnam were responsible for the observed pattern, one would expect the strongest negative effects to appear in the 1970 census data. However, the strongest effects appear in the 1980s suggesting that the immediate impact of Vietnam on men and women's behavior is not responsible for the participation effects.[78]

It is also possible, however, that these laws affected couples' labor market decisions. For instance, Goldin and Katz (2000, 2002) find that early access to the pill enabled more women to continue their professional training after their college degrees. For the cohorts born from 1940 to 1956, the timing of educational investments corresponds closely to the period correlations in the 1970s and 1980s. Without the need to support their families and children, more men could stay in school and defer going into the workforce. Moreover, if delaying childbearing increased women's earnings power through college and career investments, households may have altered the allocation of time of husband and wife to home and market production both in the immediate and longer term. However, this explanation is not inconsistent with the notion that ELA cause a shift in the relative supply of weeks worked by women. Because I am interested in the shifts of the labor supply by women *relative* to those worked by men, the remainder of the analysis uses the ratio of total weeks worked by women to those worked by men.

The Impact of Changes in the Relative Supply of Women on the Structure of Wages

Using the IPUMS data, I estimate variations of equation (5) and instrument for the ratio of weeks worked by women relative to men, or $\log(W_{st}/M_{st})$. While these models do not include controls for education, marital status or other covariates often included in wage regressions, the fixed effects should capture the effects of unobservables as long as they evolve smoothly over time. In addition, abrupt nonlinear shifts in the wage structure, the technological parameters (*a* and *c* in the model), and labor market conditions will be captured in the year fixed effects, as long as they do not

[78] I also control directly for Vietnam casualty rates in my work on women's labor-force participation and find that this changes the results very little.

impact states differentially.

The instrumental variable estimates of the impact of relative supply shifts on the dispersion in wages are presented in Table 4. Among women (Panel A), the pattern of results suggests that relative supply shifts tended to increase mean wages and increase earnings dispersion below the 75th percentile. Above the 75th percentile, earnings dispersion is negatively correlated with relative supply shifts. This is consistent with Juhn and Kim's (1999) argument that women increased the supply of skill over the 1980s. For men (Panel B) the pattern of results implies that relative supply shifts tended to increase mean earnings and dispersion in both tails of the wage distribution. The negative relationship with the 75-25 range statistic indicates that relative supply shifts tended to decrease wage inequality through the middle of the distribution. These patterns, however, are only suggestive because none of the point estimates is even marginally statistically significant. Bootstrapping the standard errors and other finite sample corrections do not alter this conclusion.

Table 4. Pooled IV Estimates of the Impact of Relative Labor Supply Shifts on Wages, 1960- 2000

	Dependent variable: Log(Weekly Wages)			
	A. Sample: Women in years 1960-2000 (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	0.026	-0.130	0.480	0.315
	[0.522]	[0.302]	[2.139]	[0.802]
F-statistic	38.42	19.80	7.33	9.68
Fixed effects	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y
	B. Sample: Men in years 1960-2000 (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	0.435	0.201	0.391	-0.476
	[0.415]	[0.516]	[0.356]	[0.593]
F-statistic	17.65	27.44	76.86	9.34
Fixed effects	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y

The dependent variable is the mean or indicated difference in the percentiles in the distribution of the log of weekly wages. Weekly wages are obtained by dividing last year's earnings by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked and correct for top-coding using the methodology described in Appendix A. Log weeks worked by W/M is the log of the number of weeks worked by women relative to the weeks worked by men in the previous calendar year. All regressions are weighted using population weights. Heteroskedasticity and autocorrelation robust standard errors are in brackets. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included. S•Y indicates that the set of state dummies, S, is interacted with a linear time trend. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

While the discussion up to now has focused on the pattern and sign of the point estimates, these results at face value suggest that relative supply shifts had little effect on the distribution of wages within gender. The period effects presented in Table 5 imply a limited role for supply shifts in the 1960s and 1980s as well. These regressions are similar to those in Table 4, but limit the analysis to the two decades for which ELA is a relevant instrument. Because there are only 100 observations, I include state fixed effects and a binary variable equal to one for the later year in the period. From the estimates, it does not appear that relative supply shifts significantly affected the mean or the dispersion of wages among women (Panel A) or among men (Panel B) during the 1960s. The estimates in the 1980s for women (Panel C) are also statistically insignificant.

Statistically significant, positive impacts on earnings dispersion above the median among men in Panel D stand in contrast to these results. Relative shifts in the supply of women of approximately 17 percent over the 1980s significantly increased dispersion above the median by approximately 0.24 log points (1.45 x 0.17). Individual regressions on the 90^{th} and 50^{th} percentiles indicate that the increase in dispersion is *not* due to a decline in median male earnings. In fact, relative supply shifts tended to increase the median (b=0.093, se=0.352), although the increase is not statistically significant. Therefore, relative supply shifts appear to have *increased* male earnings at the 90^{th} percentile relative to the median during the 1980s.

Taken together these results suggest that relative increases in the supply of women had little impact on wage inequality among men and among women. Despite the strong statistical relationship between ELA and relative labor supply, virtually no evidence supports an impact of changes in women's labor supply on their own earning or those of men from 1960 to 2000. The only exception is that increases in women's labor supply tended to increase dispersion in male earnings above the median in the 1980s.

The lack of an effect may be due to several reasons. First, the model presented implicitly assumes that women workers are of a homogenous variety and explores how women as a group may affect the structure of wages. However, the greater skill accumulation of women born between 1940 and 1956 suggests that the coefficients for women may reflect a change in selection. The well-known increases in skill among these cohorts (college, occupation, more labor market experience) may have outweighed the potentially negative impact of their greater numbers in the workplace.

Table 5. Period IV Estimates of the Impact of Relative Labor Supply Shifts on Wages, 1960 to 1970 and 1980 to 1990

	Dependent variable: Log(Weekly Wages)							
	A. Sample: Women in years 1960-1970 (N=100)				C. Sample: Women in years 1980-1990 (N=100)			
	Mean	*90-50*	*50-10*	*75-25*	*Mean*	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.197	-0.18	1.05	0.615	0.202	0.108	-0.368	-0.247
	[0.286]	[0.234]	[1.874]	[0.829]	[0.459]	[0.264]	[0.302]	[0.387]
F-statistic	34.16	270.77	25.58	246.21	443.40	100.10	4370.99	461.09
Fixed effects	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y
	B. Sample: Men in years 1960-1970 (N=100)				D. Sample: Men in years 1980-1990 (N=100)			
	Mean	*90-50*	*50-10*	*75-25*	*Mean*	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	0.061	0.166	-0.434	-0.349	0.633	1.454	0.31	0.192
	[0.466]	[0.296]	[0.536]	[0.455]	[0.377]	[0.514]	[0.583]	[0.376]
F-statistic	61.43	43.26	246.06	93.75	689.17	17.36	123.97	143.58
Fixed effects	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y	S,Y

The dependent variable is the mean or indicated difference in the percentiles in the distribution of the log of weekly wages. Weekly wages are obtained by dividing last year's earnings by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked and correct for top-coding using the methodology described in Appendix A. Log weeks worked by W/M is the log of the number of weeks worked by women relative to the weeks worked by men in the previous calendar year. All regressions are weighted using population weights. Heteroskedasticity robust standard errors are in brackets. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

Second, according to the intuition provided by the model, women may have no impact on the skill premium among men if they are comparable substitutes for skilled and unskilled men, or ρ-$\mu \cong 0$. These parameters may be too similar to distinguish empirically, particularly given the small sample size. On the other hand, if aggregate relative supply shifts mask roughly equal shifts in the number of weeks supplied across skill groups, then no percentile will be differentially affected. Therefore, failing to find an effect might suggest that relative supply shifts did *not* affect the distribution of wages or that the methodology used in this study was not able to detect it.

The Impact of Changes in the Relative Supply of Women on the Gender Gap in Wages

Table 6 explores the potential impact of changes in the relative supply of women on the gender gap in wages. In separate regressions for men and women as in Table 4 and Table 5, the dependent variable is the log of weekly wages of women at a given percentile in the female wage distribution (or the difference in two percentiles) relative to the same statistic in the distribution of

Table 6. The Impact of Relative Labor Supply Shifts on the Gender Gap in Wages, 1960 to 2000

	Dependent variable: Log(Weekly Wages Women/ Weekly Wages Men)			
	A. OLS estimates (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.236	-0.121	0.555	0.395
	[0.200]	[0.186]	[0.731]	[0.335]
Adjusted R^2	0.88	0.65	0.56	0.72
Fixed effects	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y
	B. Instrumental variables estimates (N=250)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.409	-0.521	0.956	0.507
	[0.582]	[0.442]	[2.081]	[0.893]
F-statistic	48.09	10.26	7.01	16.42
Fixed effects	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y	S,Y,S•Y

Weekly wages are obtained by dividing last year's earnings by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked and correct for top-coding using the methodology described in Appendix A. Log weeks worked by W/M is the log of the number of weeks worked by women relative to the weeks worked by men in the previous calendar year. All regressions are weighted using population weights. Heteroskedasticity and autocorrelation robust standard errors are reported in brackets. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included. S•Y indicates that the set of state dummies, S, is interacted with a linear time trend. Source: 1960-2000 Public Use Micro Series (Ruggles et al. 2004).

male wages. In this framework, a positive coefficient would imply that dispersion increased among women more than among men. Panel A presents the OLS results and Panel B presents the instrumental variable estimates. In neither case do the point estimates indicate that the greater supply of women to the labor market impacted the mean wages of women relative to men or the relative their dispersion of earnings.

The period effects for the 1960s and the 1980s in Table 7 are quite similar. As before, these models include state fixed effects and a binary variable equal to one for the later year. Relative supply shifts appear to have little effect in the 1960s. However, the effect is strong and statistically significant on wage dispersion above the median during the 1980s. Increases in relative supply not only tended to increase dispersion above the median among men, but the increase in dispersion was much greater among men than women.

Table 7. Period IV Estimates of the Impact of Relative Labor Supply Shifts on the Gender Gap in Wages, 1960 to 1970 and 1980 to 1990

	Dependent variable: Log(Weekly Wages Women/Weekly Wages Men)			
	A. Sample: Years 1960-1970 (N=100)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.257	-0.347	1.484	0.964
	[0.534]	[0.306]	[2.124]	[1.100]
F-statistic	559.75	59.26	39.92	360.09
Fixed effects	S,Y	S,Y	S,Y	S,Y
	B. Sample: Years 1980-1990 (N=100)			
	Mean	*90-50*	*50-10*	*75-25*
Log weeks worked by W/M	-0.431	-1.345	-0.679	-0.439
	[0.236]	[0.589]	[0.388]	[0.550]
F-statistic	459.49	791	32.86	109.97
Fixed effects	S,Y	S,Y	S,Y	S,Y

The dependent variable is the mean or indicated difference in the percentiles in the distribution of the log of weekly wages of women minus the same statistic for men. Weekly wages are obtained by dividing last year's earnings by the number of weeks worked. For census years prior to 1980, weeks worked were recorded in intervals. I infer actual weeks worked and correct for top-coding using the methodology described in Appendix A. Log weeks worked by W/M is the log of the number of weeks worked by women relative to the weeks worked by men in the previous calendar year. All regressions are weighted using population weights. Heteroskedasticity robust standard errors are reported in brackets. S indicates that a set of dummies for state of birth is included (Alabama is the omitted category). Y indicates that a set of dummies for the census year (1960 is excluded) is included.

More interesting, however, is that the instrumental variable estimates suggest that relative supply shifts contributed to a stagnating gender gap at the mean during the 1980s. The point estimate implies that in the absence of the 17 percent increase in the number of weeks worked by women relative to men the gender gap would have been approximately 56 smaller (.17x-.43/.13). This figure is particularly remarkable given that the 1980s witnessed the most rapid convergence in the gender gap in the second half of the Twentieth Century. Moreover, it provides concrete support for claims made in partial equilibrium wage decompositions that increasing labor supply depressed the wage growth of women in the 1980s (Blau and Kahn 1997, 2002).

Conclusions

Dramatic increases in young women's labor market participation and wage inequality have transformed the U.S. economic landscape since 1960. This study seeks to quantify the impact of changes in the relative supply of women on the structure of wages using a new instrumental variable.

I provide evidence of a strong correlation of the instrument with the intensity of labor-force participation. Moreover, the evidence suggests that the instrument is uncorrelated with changes in market demand. Despite the strong relationship in the first stage, the results suggest that the causal impact running from greater labor supply to the growth in wage inequality is weak at best. Therefore, relative changes in women's labor supply appear to have played a minor role, if any, in changing earnings inequality from 1960 to 2000.

There are two surprising exceptions to this conclusion. First, relative supply shifts tended to increase the wages of men at the 90th percentile relative to the median during the 1980s. This provides an alternate and, perhaps, more plausible interpretation of changes in wage dispersion during the 1980s. Instead of bidding down the wages of low skilled men, women workers tended to enhance the wages of highly skilled men. Relative supply shifts may have caused changes in the distribution of wages because women's skills—perhaps clerical or administrative—were more complements, rather than substitutes, to the skills of men in certain skilled occupations. This is a novel hypothesis for the trends witnessed in the 1980s. This explanation reconciles the observation that men and women tend to be concentrated in different occupations and industries with research that suggests that increases in the supply of women have increased earnings inequality among men (cf. Topel 1994, Acemoglu et al. 2004). Moreover, this hypothesis suggests that women's entry into the labor force may have *accelerated* skill-biased technical change, rather than stemmed the growth in earnings above the median as suggested by Juhn and Kim 1999.

Second, relative changes in the number of weeks worked functioned to increase gender gap in wages at the mean over the 1980s even as the wages of men and women converged. While changes in women's work were certainly one of the most important factors contributing to women's earnings growth over the 1980s, changes in their labor supply tended to mitigate their wage gains. In contrast to the studies that have focused on the demand determinants of growth in women's labor-force participation, these results suggest that changes on the supply side in the 1960s might have had far reaching and durable effects on women's earnings growth well into the 1990s.