

Life Cycle Effects of Fertility on Parents'

Labor Supply

James P. Vere*

University of Hong Kong

May 23, 2005

*School of Economics and Finance, University of Hong Kong,

Pokfulam Road, Hong Kong. Email: jpvere@hku.hk.

Abstract

This paper uses U.S. Census data from 1980, 1990 and 2000 to estimate synthetic-cohort life cycle effects of fertility on women's and couples' labor supply. Multiple births are used as an instrument to control for unobserved heterogeneity. For single women, the causal effect of fertility has declined significantly over time. Couples, however, have become more specialized along traditional lines, with married men tending to increase labor earnings rather than reduce hours worked.

JEL codes: J12, J13, J16, J22

1 Introduction

Economists have devoted much time and effort to understanding the effects of children on their parents' labor supply decisions. Moreover, though the negative relationship between fertility and female labor supply has long been recognized, efforts to minimize these effects have been the subject of numerous policy initiatives over recent decades. For instance, a major goal of welfare and tax reform in the 1990s was to reduce low-income mothers' disincentives to work; other reforms, such as the Family and Medical Leave Act of 1993, were targeted toward increasing the compatibility of raising children and work in the formal sector. Understanding how the causal effects of fertility have changed over time, and particularly the 1990s, is vital to assessing the effects

of these reforms and pinpointing where future change may be required.

This paper uses a twins-based instrumental variables strategy and U.S. Census data from 1980, 1990 and 2000 to identify the causal effects of second and third children on their parents' labor supply. An important characteristic of the strategy is that, since the multiple births in the data occur at different times in the past, the causal effects of such children can be estimated separately by age. This permits holding the age distribution of children identifying the instrumental variables estimates constant over time, which is necessary for the results to be comparable across years. In addition, sums of the age-specific causal effects over a fixed time horizon can be interpreted as synthetic-cohort life cycle effects. These are useful because they measure the total impact of an unexpected birth on household labor supply.

Several studies examining the effects of fertility on parents' labor supply have focused on panel data as a way to correct for endogeneity problems when estimating a causal relationship. Lundberg and Rose (2000) use data from the 1968-1992 Panel Study of Income Dynamics (PSID) files to find the effects of fertility on husbands' and wives' earnings and hours worked; Lundberg and Rose (2002) extend this analysis, examining the effects of additional children on men's labor supply by sex and parity. Therefore, a useful aspect of instrumental variables estimates with U.S. Census data is that they offer a robustness check of the PSID results. Further, since data from 2000 are now available, changes in the causal effects of second and third children from the mid- to late 1990s can be observed. ³This is an interesting time because of

the implementation of numerous policy reforms, many of which were targeted toward changing parents' incentives to work.

A number of instrumental variables studies also estimate the causal effects of fertility on household labor supply. For instance, Bronars and Grogger (1994) use multiple births and U.S. Census data from 1970 and 1980 to identify the effects of second children on women's earnings and labor force participation. In a similar vein, Angrist and Evans (1998) use 1980 and 1990 U.S. Census data and the sex mix of a woman's first two children as an instrument to estimate the effects of a third child on her and her husband's labor supply outcomes.

This paper adds to the above results by estimating the effects of second and third children on marital stability, household labor supply, and the degree of husbands' and wives' specialization in market work (for the latter two outcomes, the measures are those introduced by Lundberg and Rose (1999)). So far, the literature has examined causal links between fertility, marital stability, and women's income (Angrist 2004, Bedard and Deschenes 2005), but the effects of fertility on couples' joint outcomes remain largely unexplored with instrumental variables methods. Nevertheless, since most children are raised by two parents, the effects of fertility on household income and both parents' labor supply are important and of wide general relevance.

The empirical results show that the largest shifts in behavior in the 1980s and 1990s occur for single women and married men. For single women, the causal effects of fertility decline markedly over time—the synthetic-cohort life

cycle effect of a second child over that child's first thirteen years declines from $-6,956$ labor hours in 1980 to $-2,522$ hours in 2000. This suggests that incentives for single mothers to work have improved considerably over time. For married men, the change is even more dramatic—though in 1980 a second child causes an average total decline of $\$23,874$ in men's market earnings (in 2004 dollars), in 2000 the second child causes an increase of $\$78,856$. The implication is that men have become more likely to respond to unexpected births with income rather than time. Therefore, even though husbands' and wives' market wages have become more equal over time (Pencavel 1998), specialization within married couples in response to fertility has increased.

It should be noted that since these are synthetic-cohort effects, they are best viewed as summaries of age-specific causal effects for a particular year rather than descriptions of the experience of any one birth cohort. In this way they are similar to life expectancy statistics, which summarize a year's age-specific mortality rates. Because of this, and also because they are identified from a fixed age distribution, they are most useful for analyzing how these effects have changed over time.

This paper is organized as follows. The second section discusses the data, instruments, and first-stage relationships. The third section presents results on fertility and female labor supply, and the fourth section extends the analysis to marital stability, couples' joint labor supply and specialization between market and household work. The fifth section concludes.

2 Data, Instruments and First-Stage Relationships

2.1 Data

The empirical work in this paper uses data on fertility, demographic characteristics and labor supply from the one percent and five percent Census Public Use Micro Samples (PUMS) from 1980, 1990 and 2000. The one percent and five percent data sets for each year are non-overlapping random samples from the same population. They contain different geographic descriptors but are similar in all other respects. Standard demographic information and labor supply measures are available in each of these years.

To motivate later analysis, Panel A of Table 1 summarizes these characteristics for women aged 21-35 in each data set. For 1980, there are three samples: a five percent sample (the A sample) and two one percent samples (the B and C samples). For 1990 and 2000, there are two samples per year—a one percent sample and a five percent sample.

Separating the results by sample allows a brief empirical check of the assertion that they are, in fact, drawn from the same population. Within years, differences in mean demographic and labor supply characteristics are small and no more than would be expected due to ordinary sampling error. Therefore, the data sets can be merged to create a seven percent sample for 1980 and six percent samples for 1990 and 2000. While this is unnecessary for most

purposes, since the twin births on which the estimation strategy relies are rare events, it is useful to assemble as much data as possible.

Panel B of Table 1 summarizes characteristics of women ages 21-35 with at least one child ever born. This population is important because instrumental variables estimates are often interpreted as local average treatment effects (Angrist and Imbens 1995). The multiple births on which the estimation strategy relies are, of course, impossible without a single intended birth in the first place. As a consequence, when indicators for multiple births are used as instruments, they identify the average causal effect of second and higher-parity children for women with at least one child already.

Because the PUMS data sets include only limited information on fertility history, it is necessary to impute the number and timing of a woman's children from data on household composition. The method of doing so is very similar to that employed by Angrist and Evans (1998). A set of women with imputable birth histories was constructed consisting of all women reported as being either household heads, the spouse of the household head, parents in a parent-child subfamily, or wives in a married-couple subfamily. Relationship codes and subfamily identifiers were then used to match children to mothers within households and subfamilies. In addition, for 1980 and 1990, mothers were deleted from the sample if their reported number of children ever born did not match the number of children present.¹ Women were also deleted if they were less than 16 years older than their first child or if any of the ages or

¹This question was not asked in 2000.

sexes of their children were allocated by the U.S. Census Bureau.

A possible concern is that limiting the analysis to women with imputable birth histories introduces a degree of selection into the sample. To address this issue, Panel C of Table 1 summarizes characteristics of women with at least one child for whom birth histories can be imputed. These women tend to be slightly more educated and less likely to have multiple children than the women in Panel B. A possible explanation is that more educated women fill in the long census forms more accurately (or at least more consistently), making it more likely that viable birth histories can be constructed from their responses. However, though these differences are statistically significant, they are not substantial in economic terms, and differences in other characteristics are negligible. Therefore, the imputation method does not seem to create an appreciable selection problem. Alternatively, the estimation results can be interpreted with the caveat that they represent a population that is slightly more educated than average.

2.2 Instruments

The estimation strategy in this paper uses indicators of multiple births as instruments for fertility. Since women are not asked directly whether they have given birth to twins, triplets or other multiple births, these indicators must be imputed from data on children's dates of birth. For example, the indicator TWINS2 is set equal to one if a woman has a second child born at

the same time as her first. Variables on quarter and year of birth are available for the 1980 data; for 1990 and 2000, only data on year of birth are available.

Table 2 shows the sample means of the instruments for women with imputable birth histories and at least one child in the merged data sets for 1980, 1990 and 2000. TWINS3 indicates whether a woman has a third child born at the same time as her second. For this instrument, the sample is further limited to women with at least two children. For the 1980 data, the suffix Q indicates that multiple births were imputed based on both quarter and year of birth data; where the instrument names do not have this suffix, they were constructed with year of birth data only.

A potential concern is that, since it is possible for children to be born in the same year and not be twins, constructing the instruments with only year of birth data introduces some degree of measurement error. From an instrumental variables standpoint, however, this measurement error is only problematic when it is correlated with the error term in the second stage (that is, the unobserved propensity to supply labor). Overidentification tests with the 1980 data can be used to determine whether such a correlation exists. The results from these tests are presented later, in section 3.1. For now, it suffices to say that there is no evidence that IV estimates of the effects of fertility on female labor supply have different probability limits depending on whether quarter of birth data are used to create the instruments. Therefore, the correlation is either nonexistent or too small to be detected even in a very large data set.

This issue aside, the means of the multiple birth indicators are close to what one would expect given historical rates of twin births in the United States, which have been increasing since 1981. In 1980, there were 18.9 twins per thousand live births. Since both twins are counted in this statistic, it should be divided by two to be comparable with the means in Table 2. The twin birth rate increased steadily over the 1980s, rising to 22.6 per thousand births by 1990. After 1990 it rose more quickly, reaching 31.1 by 2002 (Martin, Hamilton, Sutton, Ventura, Menacker and Munson 2003). Most of this trend can be attributed to increased use of fertility drugs and other forms of assisted reproduction technology. Martin and Park (1999) provide a discussion of this trend and detailed statistics of twinning rates over time.

2.3 First-Stage Relationships

The intuition behind the use of multiple births as a natural experiment is clear and easily grasped. Every year, several million women in the United States give birth. Out of these, biological processes cause a few out of every thousand to be multiples, resulting in two or three children where only one was intended (certainly for the moment, and in some cases indefinitely). While twins are more likely in certain demographic groups—black women and older women, for example—the biological processes themselves are random and not related to a woman’s unobserved propensity to supply labor. Hence, the incidence of multiple births can be used as an instrument to estimate the effects of

unexpected births on women’s labor supply decisions.

In the simplest case, multiple births would be assigned with equal probability among the population of women choosing to have one child in a given year. Then the average effect of multiple births on fertility could be found simply by differencing the average number of children among both groups of women. However, since twinning rates depend on race and age, it is important to control for these characteristics when taking this difference. Therefore, a linear regression is useful for calculating the average effect of multiple births on fertility when these characteristics are held constant.

Along these lines, Table 3 reports ordinary least squares estimates of the equation

$$C = \pi_0 + \boldsymbol{\pi}_1 \cdot \mathbf{X} + \pi_2 \cdot Z + \eta \tag{1}$$

where C is an indicator equal to one if a woman has at least n children, \mathbf{X} is a vector of exogenous covariates and Z is an indicator equal to one if the $n - 1$ th and n th children are part of the same multiple birth set. The numerals in the instrument names denote n . Specifying C in this way is useful because it emphasizes the fact that the instrument causes fertility to move from $n - 1$ to n children; it also permits comparing the results to Angrist and Evans’s (1998), who make the same choice for similar reasons. In addition, for each instrument, the sample is restricted to women with at least $n - 1$ children as in Table 2. The covariates in \mathbf{X} are age, age squared, years of education, and dummy variables for being black and being married.

Table 3 confirms that multiple births have significant, measurable effects on whether an n th child is present in the household. For example, the results for TWINS2Q indicate that only about two thirds of women with a twin second child in 1980 would have had a second child on their own by the date of the survey. For TWINS3Q, the effects of multiple births are more pronounced because even fewer women would have had a third. These effects are also more pronounced for 1990 and 2000 because fertility has been declining generally. This causes the chance that a woman with twins would not have had an additional child by the survey date to increase over time.

When the instruments are constructed with only year of birth data, measurement error subjects estimates of these effects to a slight degree of attenuation bias. This does not affect the final instrumental variables results because estimates of the instruments' effect on labor supply outcomes are attenuated by the same factor, cancelling out the bias when the ratio is taken. This ratio is relevant because, as Angrist, Imbens and Rubin (1996) show, instrumental variables estimates can be written as a ratio of the causal effects of a binary instrument on two outcomes of interest (in this case, fertility and female labor supply). As for the first-stage results, inspection of the estimates for 1980 suggests that the magnitude of the attenuation bias is less than five percent. Additionally, since fewer women have chosen to have additional children over time, the attenuation bias is less for the 1990 and 2000 results because more children born the same year in these data sets are actually twins.

3 Fertility and Female Labor Supply

3.1 Two-Stage Least Squares Estimates

As Angrist and Imbens (1995) have shown, two-stage least squares estimates can be interpreted as average per-unit causal effects of a treatment variable on an outcome of interest. In the case of fertility, the meaning of a “per-unit” increase is clear: instrumental variables estimates calculated using the TWINS instruments identify the causal effects of second or third children on female labor supply. An extensive literature supports the notion that this relationship is negative (for a survey, see Browning (1992)).

Accordingly, Table 4 reports two-stage least squares estimates of δ , the local average treatment effect of the n th child, in the equation

$$Y = \beta_0 + \beta_1 \cdot \mathbf{X} + \delta \cdot C + \varepsilon \tag{2}$$

where Y is a labor supply outcome, \mathbf{X} is the vector of covariates, and C is the endogenous fertility indicator, equal to one if a woman has at least n children. The first stage is given by equation (1). As in Table 3, the sample is limited to women with at least $n - 1$ children for each choice of C , and the covariates in \mathbf{X} are age, age squared, years of education, and dummy variables for being black and being married. For labor outcomes relating to the year prior to the survey, the sample is limited to women with at least $n - 1$ children in that year.

A distinct observation that can be made from Table 4 is that, without exception, the causal effect of fertility declines markedly with child parity. For example, in 2000, the second child causes an average loss of \$3357 in annual labor income (in 2004 dollars), while the third causes an average loss of only \$1390. Considering that the average labor income for women with one child in 2000 was \$16,430, these effects are substantial. Similarly, the effect of the third child on usual hours worked per week is only about half that of the second child.

Where comparable, the estimates in Table 4 are quite close to others found in the literature. For example, Carrasco (2001) analyzes 1986-1989 data from the U.S. Panel Study of Income Dynamics with a switching probit model that endogenizes the fertility decision. Her estimate of the average causal effect of an additional child on female labor force participation is -12.9 percent. This is quite close to the second-child effect on female labor force participation of -13.8 percent in the 1990 data.

The effects of the third child are also in consensus with the literature. For instance, Angrist and Evans (1998), using 1980 U.S. census data and the TWINS2 instrument, calculate the effect of the third child on employment to be -7.9 percent; the analogous figure in Table 4 is -8.6 percent. Angrist and Evans's (1998) other calculations with the 1980 census data and this instrument are also quite comparable.

Finally, Table 4 also reports, for the 1980 data, overidentification test statistics testing the exogeneity of multiple birth indicators constructed with only

year of birth data. For effects relating to the second and third children, exogeneity of the instruments is never rejected at the one percent level and rejected only once at the five percent level. Since there are twelve of these statistics, that one should be significant at the five percent level is not that surprising. Moreover, given the size of the data set, the overidentification test rules out even a very small correlation between the year of birth instruments and the second stage error term. Consequently, the absence of the quarter of birth variable for the 1990 and 2000 data does not seem to pose any particular problem for these estimates.

3.2 Age-Specific and Life Cycle Effects

Up to this point, fertility has been characterized as a simple, constant-intensity treatment—either there is a second or third child in the household or there is not. However, it is more realistic to imagine fertility as a variable-intensity treatment, where the effects on female labor supply are greatest at young ages and diminish as the child gets older. When this is the case, each result in Table 4 can be interpreted as the average causal effect of a second or third child drawn from a particular age distribution.

As Angrist and Imbens (1995) have shown, instrumental variables estimates are weighted causal effects of incremental changes in a variable treatment, where the weight assigned to each incremental change is proportional to the probability that the instrument induces that change. Therefore, the weight

of each age identifying the causal effects in Table 4 is proportional to the probability that the instrument induces an n th child of that age.

These probabilities can be estimated by the following method. First, define an age-specific fertility indicator $C(a)$ to be equal to one if there is an n th child of age a in the household and zero otherwise, so that $\sum_A C(a) = C$. Then the weight of each age identifying the causal effects in Table 4 is proportional to

$$E[C(a) | Z = 1] - E[C(a) | Z = 0] \tag{3}$$

where Z is the relevant instrument. Expanding expression (3) to include covariates would complicate the notation but not change the intuition. The first-stage estimates in Table 3 result when expression (3) is summed over all child ages in the sample.

Figure 1 graphs these probabilities by year, instrument, and age of the potential second or third child. The instrument TWINS2 is used to calculate the second-child probabilities and TWINS3 is used to calculate the third-child probabilities. In both cases, the sample is limited to women with at least one or two children, respectively.

A noticeable feature of the distributions in Figure 1 is that third children identifying the estimates in Table 4 tend to be drawn from younger age distributions than second children. This occurs for two reasons. First, because women have second children after first children, twins of a second child are always younger than twins of a first. Second, multiple births sometimes ac-

celerate the births of previously-planned children, in effect replacing a younger child in the household with an older one. Since more women intend to have two children than three, this factor is more pronounced for the TWINS2 instrument than it is for TWINS3. Moreover, since fewer women have intended to have second and third children over time, this is also why the distributions are younger in 1990 and 2000.

One might wonder whether the trends in Table 4 reflect changes in the age distributions identifying these estimates. This is not the case for child parity—third children have less impact on their mothers' labor supply decisions in spite of being younger and requiring more care. Nevertheless, the fact that the children are drawn from younger age distributions in later years may explain why some of the causal effects have risen over time.

The variable-intensity character of fertility can be explored more fully by estimating the causal effects of second and third children separately by age. The intuition behind doing so is to realize that, since the natural experiment brought about by occurrences of multiple births is ongoing, the PUMS data sets contain snapshots of the outcomes of many such experiments, all initiated at different points in time. Therefore, multiple births will have varying effects on a set of age-specific fertility indicators depending on how long ago the multiple birth occurred. For instance, a twin birth occurring seven years ago induces the presence of a seven-year-old (and possibly has a negative effect on the presence of a younger child). This means that, by interacting the multiple birth indicator with a set of dummy variables for time elapsed

since the multiple birth occurred, a separate age-specific causal effect can be estimated for each time dummy variable.

More specifically, age-specific causal effects can be estimated by two-stage least squares on the equation

$$Y = \beta_0 + \boldsymbol{\beta}_1 \cdot \mathbf{X} + \boldsymbol{\delta} \cdot \mathbf{C} + \varepsilon \quad (4)$$

where \mathbf{C} is a vector of endogenous age-specific fertility indicators and the other variables are defined as earlier. The age-specific causal effects are the elements of $\boldsymbol{\delta}$. The first stage consists of k equations of the form

$$C(a) = \pi_{0a} + \boldsymbol{\pi}_{1a} \cdot \mathbf{X} + \boldsymbol{\pi}_{2a} \cdot \mathbf{Z} + \eta_a \quad (5)$$

where k is the number of age-specific causal effects to be identified and $a = 0, 1, \dots, k - 1$. \mathbf{Z} is a vector created by interacting the multiple birth indicator with a set of k time dummy variables, T_0, T_1, \dots, T_{k-1} , where T_i is equal to one if the multiple birth occurred i years ago and zero otherwise. Since the number of age-specific causal effects to be estimated is equal to the number of elements in \mathbf{Z} , the system is exactly identified.

Figures 2 and 3 show estimates of these causal effects for two labor supply outcomes, labor force participation and usual hours worked last year, over a thirteen-year horizon. For these estimates, the sample is limited to women with an $n - 1$ th birth occurring within the past thirteen years. Analogous

graphs for other outcomes are qualitatively very similar, differing only in the scale of the vertical axis.

Thirteen years is chosen as the horizon for a variety of reasons. First, children start junior high school at thirteen and no longer need constant supervision. Perhaps not coincidentally, age thirteen is also when child care expenses cease to qualify for favorable tax treatment. Finally, prior studies (e.g., Angrist and Evans 1998) have found the effect of fertility on female labor supply to reach zero at this age. Therefore, the interval from age zero to age thirteen seems to be the most sensible for calculating the total effect of an unplanned birth.

Figures 2 and 3 confirm that the effects of fertility on female labor supply are greatest when the child is born and then rapidly decline. For instance, in Figure 2, a second child under twelve months old in 1990 reduces labor force participation by 25 percent, but for a two-year-old the reduction is only 17 percent. At six, the effect diminishes to 13 percent, and by age twelve it is no longer significantly different from zero. For the third child, the decline is even more rapid—Figure 3 shows that in 2000, a child under one reduces usual hours worked per week by five, but the effect of a six-year-old is only half that; the causal effects cease to be statistically significant at age eight.

Figures 2 and 3 also show that, as in Table 4, the causal effects of fertility decline sharply with child parity. Interestingly, though third children have less impact than second children at all ages, the divergence becomes more pronounced at age six, or the start of first grade. This suggests that women

realize more economies of scale when raising school-age children than when looking after infants and preschoolers.

The total effect of an unplanned birth over a woman's career can be measured by summing the age-specific effects over the thirteen-year interval. This is the synthetic-cohort life cycle effect of an unplanned birth on female labor supply. Table 5 presents these sums and their standard errors. Since the sums are by year, the weekly hours measures are multiplied by 52 to convert them to annualized measures.

For the time-based measures, the results in Table 5 all point in very similar directions. In 1980, the second child causes a total decline in female labor supply equivalent, on average, to about two years of full-time work (or 104 weeks or 4,160 hours); the causal effect of the third child is about half that. Further, both the effects of the second and the third child diminish over time. By 2000, the average causal effect of the second child declines to about 1.7 years of full-time work, and the effect of the third child falls by approximately the same proportion.

The results for wage and salary income highlight the opportunity costs women face when leaving the labor force. The effect of the second child on total market income falls from \$47,052 (in 2004 dollars) in 1980 to \$35,302 in 1990 and then remains relatively stable. The effect of the third child, however, climbs from \$20,317 in 1980 to \$31,822 in 1990 before falling back to \$19,131 in 2000. An explanation for the peaks is that they represent points at which women tend to leave full-time jobs, thus incurring higher opportunity costs

than if they simply work fewer hours in a part-time position.

Comparing wages and hours reduced in each case lends support for this hypothesis. For instance, in 1980, dividing the wage effect by the effect on usual hours per week for the second child yields an average loss of \$11.47 per hour of market work. For the third child, the average opportunity cost per lost hour of market work is only \$9.30, which suggests that women with an unexpected second child were leaving better jobs than women with an unexpected third child. On the other hand, the opposite is true in 1990: these ratios are \$10.07 per hour for the second child and \$15.49 for the third.

There is no obvious spike in 2000 as there are in earlier years. There are two possible reasons for this. One is that the peak has simply moved beyond the third child and is no longer visible. The other is that, due to institutional changes in the labor market, the decision to work full-time or raise children is no longer the either-or proposition it once was. For example, the 1993 Family and Medical Leave Act required employers to provide twelve weeks of unpaid maternity leave; this legislation was part of a general movement in the 1990s toward making employment more compatible with family obligations (Waldfogel 1999). Finally, it is worth noting that these explanations are not mutually exclusive—the increasing compatibility over time may well explain why more women have continued in full-time jobs even with ever-higher numbers of children at home.

4 Fertility and Household Labor Supply

4.1 Marital Stability and the Labor Supply of Single and Married Women

In addition to direct effects on women's labor supply decisions, unplanned births can also affect female labor supply by inducing changes in household structure. Moreover, these changes often have profound consequences. For instance, if the unexpected birth causes a marriage to dissolve, the newly single mother will make her time allocation decisions under vastly different conditions than if her marriage had remained in place.

The literature has been divided on whether such children tend to induce stable marriages or push them to the breaking point. According to established theory, children should promote marital stability if dissolution of the marriage reduces the utility parents derive from them. On the other hand, unanticipated events, whether positive or negative, tend to be destabilizing when the initial marriage decision is made under uncertainty (Becker, Landes and Michael 1977).

Panels A and B of Table 6 examine both of these hypotheses. The entries in these panels are two-stage least squares estimates of the *average* element of λ in the equation

$$M = \alpha_0 + \alpha_1 \cdot \mathbf{X} + \lambda \cdot \mathbf{C} + \varepsilon \tag{6}$$

where M is an indicator for being currently married, \mathbf{C} is a vector of en-

ogenous age-specific fertility indicators, and \mathbf{X} is a vector of covariates (the covariates are age, age squared, years of education, and a dummy variable for being black). As in section 3.2, the first stage is given by equation (5) and age-specific causal effects are estimated until age thirteen. The entries can be interpreted as local average treatment effects of unplanned children on the probability of being married, where the children are drawn from a distribution that weights each age from zero to thirteen equally. Adopting this method permits comparing the estimated effects to one another while holding the age distribution of the children identifying these effects constant.

These results in these panels are generally consistent with others in the literature. For instance, Heaton (1990), using data on fertility and marital histories from the 1985 June Current Population Survey, finds that children up to the third improve marital stability but are destabilizing after this point. A decline like this is sensible if additional children yield diminishing marginal utility to their parents. In Panels A and B, the effects of later children are also more negative than the effects of earlier children, though the point at which they fall below zero seems to occur at the third child rather than beyond this point.

More recently, Angrist (2004) uses 1990 census data to estimate the causal effect of the third child on marital stability, instrumenting for fertility with the sex mix of the first two children. His estimate of the effect of the third child on M in a model with covariates is -0.052 , which is similar in magnitude to the estimate of -0.043 given in Panel B. Both of these estimates are statistically

significant.

If a marriage does, in fact, dissolve, the environment in which a woman allocates her time fundamentally changes. First, though she may receive some form of child support, in most cases she has substantially less household income than she did earlier (Duncan and Hoffman 1985, Page and Stevens 2004). This raises the marginal utility of such income and hence the value of an additional hour spent working in the formal sector.

On the other hand, though there is little evidence that marital status *per se* has a causal effect on earnings (Krashinsky 2004, Bedard and Deschenes 2005), a loss of family income can reduce a woman's incentive for market work if she becomes eligible for welfare benefits or other means-tested social assistance. Moreover, if her former husband had been assisting with household work, dissolution of the marriage will also raise the marginal value of an hour spent raising children. Whether single women reduce their labor supply by more or less than married women in response to unexpected births depends on which of these factors outweighs the other.

Panels C and D of Table 6 examine this question, reporting life cycle effects of the second child on labor supply outcomes for single and married women. Except for being estimated separately by marital status, these effects are calculated by the same method as those in Table 5. The effects for the third child are not shown but ultimately lead to similar conclusions. In 1980 and 1990, single women reduced their labor supply far more than married women did in response to unexpected children, suggesting either that husbands played a

valuable role in freeing up their wives' time to work in the formal sector or that single women faced very strong disincentives to doing so. In 2000, married women spent more time out of the labor force in response to a second child but both groups of women incurred similar reductions in labor income.

A striking characteristic of these results is that while the effects of fertility on married women's labor supply have remained fairly stable over time, the effects for single women have fallen swiftly over the past two decades. Meyer and Rosenbaum (2001) provide a compelling explanation for this decline, noting that changes in welfare and tax policy over this time greatly increased incentives of single mothers to work. The most important of these were the expansion of the Earned Income Tax Credit (EITC) and other tax reforms in the late 1980s and early 1990s. Meyer and Rosenbaum attribute a large portion of increases in single mothers' employment between 1984 and 1996 to these policy shifts. Grogger (2003) examines labor earnings and the EITC and arrives at similar conclusions.

The implication of these findings is that, since single mothers have historically faced strong disincentives to work, the potential for later children to destabilize marriages tended to magnify their negative effects on female labor supply. This was especially true in the 1980s and early 1990s, but less so after welfare and tax reforms improved single mothers' incentives to pursue market work. By contrast, the effects of earlier children, which tend to promote marital stability, were mitigated by this mechanism.

4.2 Effects on Men's and Couples' Labor Supply

In stable marriages, unexpected births affect the labor supply decisions of not just mothers but fathers as well. Ever since Becker (1981), a number of theories have been proposed to explain how fertility affects married couples' time allocation decisions. Of these, the most consistent with empirical evidence have been those that model children as a jointly produced public good (e.g., Lam 1988, Lundberg 1988).

In this type of model, though additional children have an unambiguously negative effect on female labor supply, the implications for male labor supply are less clear. If the wife allocates more time to home production, the husband may allocate more time to market work so the couple can realize gains from specialization. On the other hand, since the value of time spent raising children increases for both parents, the couple may be better off if he allocates more time to home production instead. Lundberg and Rose (2002) call these effects the specialization effect and the home-intensity effect, respectively.

Panel A of Table 7 examines the response of married men to unexpected second children. The estimates in this panel are synthetic-cohort life cycle effects calculated by the same method as in Table 5. In 1980 and 1990, the effects of the second child on employment, labor force participation, hours worked and weeks worked are negative, but much smaller in magnitude than the effects on married women's labor supply in Table 6. In 2000, the estimates vary in sign but none are statistically significant. Therefore, for the time-based

measures, the home-intensity effect outweighs the specialization effect in 1980 and 1990, but neither factor appears to dominate the other in 2000.

Where labor earnings are concerned, the effect of the second child turns from negative to strongly positive over time. By 2000, this effect is highly significant—an unexpected second child causes husbands to earn, on average, \$78,856 more over the first thirteen years than they otherwise would have (in 2004 dollars). The specialization effect is more pronounced for labor earnings than it is for the time-based measures because the former measure implicitly puts more weight on the decisions of high-wage men (who are, in turn, more likely to have a comparative advantage in market work). Similar observations are made by Lundberg and Rose (2002), who, using data from the Panel Study of Income Dynamics (PSID), also find that men’s labor earnings rise in response to fertility.

An important difference between these results and Lundberg and Rose’s (2002) is that, while Lundberg and Rose find that men increase their hours worked in response to additional children, the results in Table 7 show that men decrease their hours worked, particularly in earlier survey years. However, it should be noted that the instrumental variables estimates identify the effects of *unexpected* children on men’s labor supply. It is quite likely that couples consider their options in the labor market before deciding to have a child; if so, fathers of twins may have fewer options to increase their labor supply than fathers of second children born under more ordinary circumstances. Therefore, it is sensible for the instrumental variables estimates to be more negative than

similar panel data estimates.

Since the time allocation decisions of a married couple are made jointly, it is also useful to consider the effects of fertility on the labor supply of the couple as a unit. Lundberg and Rose (1999) propose two measures of a couple's joint labor supply decision: market intensity, which is the husband's and wife's labor supply measures added together; and specialization, which is the difference between the husband's and the wife's labor supply. Panels B and C tabulate life cycle effects of the second child on both of these measures. The methodology is the same as in Panel A, except that the wife's characteristics are included as covariates in addition to the husband's.

Comparing the results in Panel A to the effects on couples' market intensity in Panel B shows that, on average, married women bear far more of the time costs of unexpected second children than their husbands do. In 1980 and 1990, the husband's reduction in labor supply comprises, on average, less than a fifth of the couple's total response; in 2000, this proportion is even lower. This decline is surprising because, since the comparative advantage of husbands in market work has fallen over time (Gray 1997, Pencavel 1998), one would have expected the opposite to occur.

The results for wage and salary income in Panel B reinforce this finding. While in 1980 unexpected second children caused an average net loss to couples of \$56,714 in wage income over thirteen years, in 1990 their effects on wage income were negligible and by 2000 such children actually caused couples to gain \$31,545. Since these changes are largely due to changes in the behavior

of married men, the implication is that, despite declines in men's comparative advantage in market work, specialization in couples' work behavior has increased over time.

Panel C of Table 7 examines the effects of unexpected second children on couples' specialization in more detail. Although there is a slight decline in time specialization between 1980 and 1990, there is a strong increase in both types of specialization between 1980 and 2000. This increase is consistent with Lundberg and Rose's (2002) observation that men from later cohorts tend to increase their labor supply by more in response to second children than men from earlier cohorts. These results suggest that, over the 1980s and 1990s, married men—and especially high-wage married men—have become more likely to respond to unexpected births with income rather than time.

In summary, though the effects of fertility on married women's labor supply have remained relatively stable, husbands' responses have shifted toward specialization in market work. As a result, the traditional division of labor in households with children has been reinforced, despite the fact that husbands and wives increasingly have similar market wages. However, since husbands have become more likely to increase their labor income in response to unexpected children, the income shock to the household caused by such children has become positive over time.

5 Conclusion

This paper has estimated the causal effects of second and third children on women's and couples' labor supply in 1980, 1990 and 2000. Although the behavior of married women has been relatively stable over time, the effects of fertility on single women's labor supply have declined considerably. Married men have become more likely to respond to additional children by increasing their earnings from market work. Also, while second children tend to promote marital stability, third children tend to destabilize marriages. This reduces the effect of second children on female labor supply but magnifies it for third children.

The results for single women coincide with a number of policy initiatives designed to increase their incentives to work. The evidence in this paper, combined with the weight of past studies of the issue, makes it reasonable to surmise that these policy reforms have had at least some success in doing so. For married women, changes in the effects of fertility were much less dramatic, which may be because many of these policies (e.g., welfare reform and expansion of the Earned Income Tax Credit) were targeted at people with lower incomes.

Finally, the results for couples are striking because, contrary to expectations, specialization in husbands' and wives' roles when raising children has become more pronounced over time. This may be beneficial because they increase their joint labor earnings considerably in the process. Nevertheless,

it is not clear why this shift has taken place and further research into its causes is warranted.

References

Angrist, Joshua D., “Treatment Effect Heterogeneity in Theory and Practice,” *Economic Journal*, March 2004, 114 (494), C52–C83.

— and **Guido W. Imbens**, “Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity,” *Journal of the American Statistical Association*, June 1995, 90 (430), 431–442.

— and **William N. Evans**, “Children and Their Parents’ Labor Supply: Evidence from Exogenous Variation in Family Size,” *American Economic Review*, June 1998, 88 (3), 450–77.

—, **Guido W. Imbens**, and **Donald B. Rubin**, “Identification of Causal Effects Using Instrumental Variables,” *Journal of the American Statistical Association*, June 1996, 91 (434), 444–455.

Becker, Gary S., *A Treatise on the Family*, Cambridge, MA: Harvard University Press, 1981.

—, “Human Capital, Effort, and the Sexual Division of Labor,” *Journal of Labor Economics*, January 1985, 3 (1, Part 2: Trends in Women’s Work, Education, and Family Building), S33–S58.

— , **Elisabeth M. Landes, and Robert T. Michael**, “An Economic Analysis of Marital Instability,” *Journal of Political Economy*, December 1977, *85* (6), 1141–1188.

Bedard, Kelly and Olivier Deschenes, “Sex Preferences, Marital Dissolution and the Economic Status of Women,” *Journal of Human Resources*, forthcoming, 2005.

Bronars, Stephen G. and Jeff Grogger, “The Economic Consequences of Unwed Motherhood: Using Twin Births as a Natural Experiment,” *American Economic Review*, 1994, *84* (5), 1141–1156.

Browning, Martin, “Children and Household Economic Behavior,” *Journal of Economic Literature*, September 1992, *30* (3), 1434–1475.

Carrasco, Raquel, “Binary Choice with Binary Endogenous Regressors in Panel Data: Estimating the Effect of Fertility on Female Labor Force Participation,” *Journal of Business and Economic Statistics*, October 2001, *19* (4), 385–394.

Duncan, Greg J. and Saul D. Hoffman, “A Reconsideration of the Economic Consequences of Marital Dissolution,” *Demography*, November 1985, *22* (4), 485–497.

Gray, Jeffrey S., “The Fall in Men’s Return to Marriage: Declining Productivity Effects or Changing Selection?,” *Journal of Human Resources*, Summer 1997, *32* (3), 481–504.

- Grogger, Jeffrey**, “The Effects of Time Limits, the EITC, and Other Policy Changes on Welfare Use, Work, and Income Among Female-Headed Families,” *Review of Economics and Statistics*, May 2003, 85 (2), 394–408.
- Heaton, Tim B.**, “Marital Stability Throughout the Child-Rearing Years,” *Demography*, February 1990, 27 (1), 55–63.
- Krashinsky, Harry A.**, “Do Marital Status and Computer Usage Really Change the Wage Structure?,” *Journal of Human Resources*, Summer 2004, 39 (3), 774–91.
- Lam, David**, “Marriage Markets and Assortative Mating with Household Public Goods: Theoretical Results and Empirical Implications,” *Journal of Human Resources*, Autumn 1988, 23 (4), 462–487.
- Lundberg, Shelly**, “Labor Supply of Husbands and Wives: A Simultaneous Equations Approach,” *Review of Economics and Statistics*, May 1988, 70 (2), 224–235.
- and **Elaina Rose**, “The Determinants of Specialization Within Marriage,” June 1999. Working paper, University of Washington.
- and —, “Parenthood and the Earnings of Married Men and Women,” *Labour Economics*, November 2000, 7 (6), 689–710.
- and —, “The Effects of Sons and Daughters on Men’s Labor Supply and Wages,” *Review of Economics and Statistics*, May 2002, 84 (2), 251–68.

Martin, Joyce A. and Melissa M. Park, “Trends in Twin and Triplet Births: 1980-97,” *National Vital Statistics Reports*, September 1999, 47 (24), 1–17.

—, **Brady E. Hamilton, Paul D. Sutton, Stephanie J. Ventura, Fay Menacker, and Martha L. Munson**, “Births: Final Data for 2002,” *National Vital Statistics Reports*, December 2003, 52 (10), 1–114. Revised June 2004.

Meyer, Bruce D. and Dan T. Rosenbaum, “Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers,” *Quarterly Journal of Economics*, August 2001, 116 (3), 1063–1114.

Page, Marianne E. and Ann Huff Stevens, “The Economic Consequences of Absent Parents,” *Journal of Human Resources*, Winter 2004, 39 (1), 80–107.

Pencavel, John, “Assortative Mating by Schooling and the Work Behavior of Wives and Husbands,” *American Economic Review*, May 1998, 88 (2), 326–329.

Waldfogel, Jane, “Family Leave Coverage in the 1990s,” *Monthly Labor Review*, October 1999, 122 (10), 13–21.

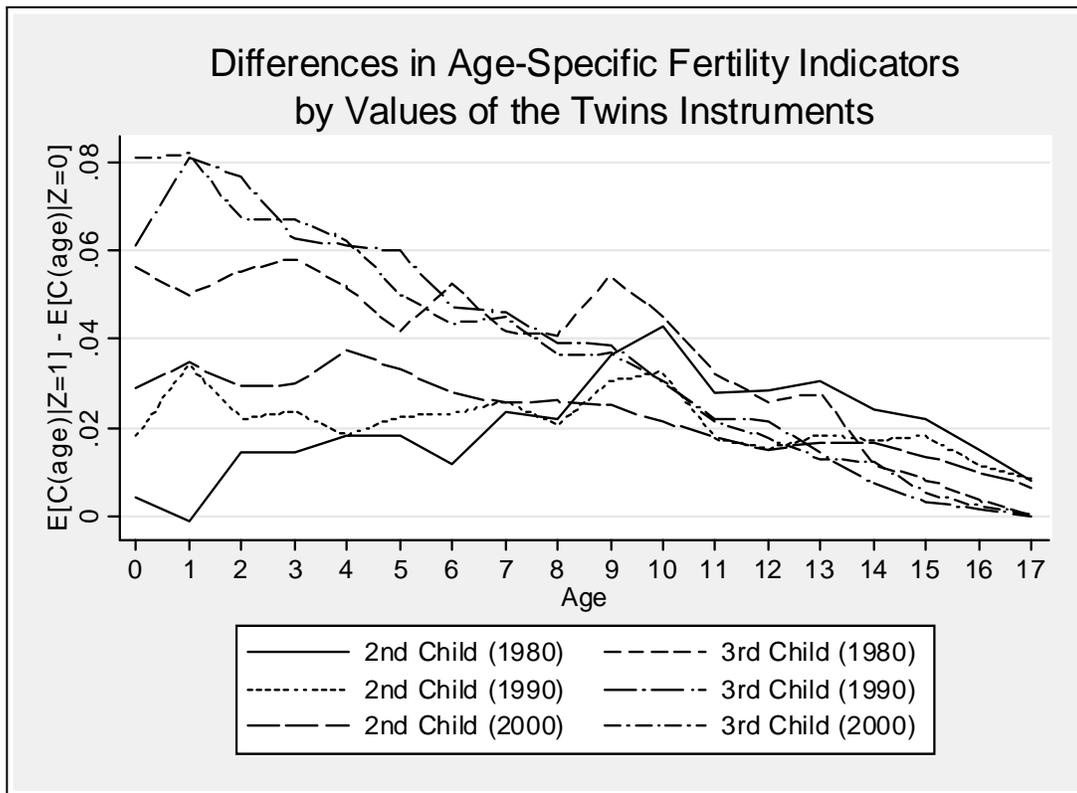


Figure 1:

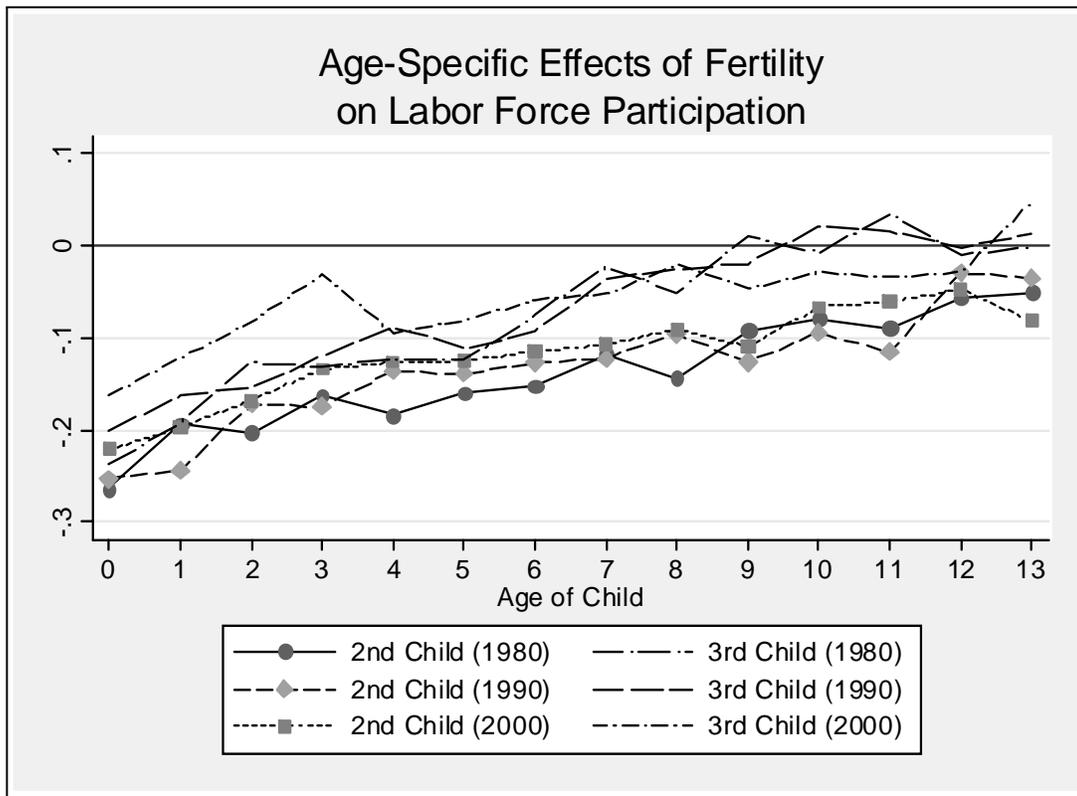


Figure 2:

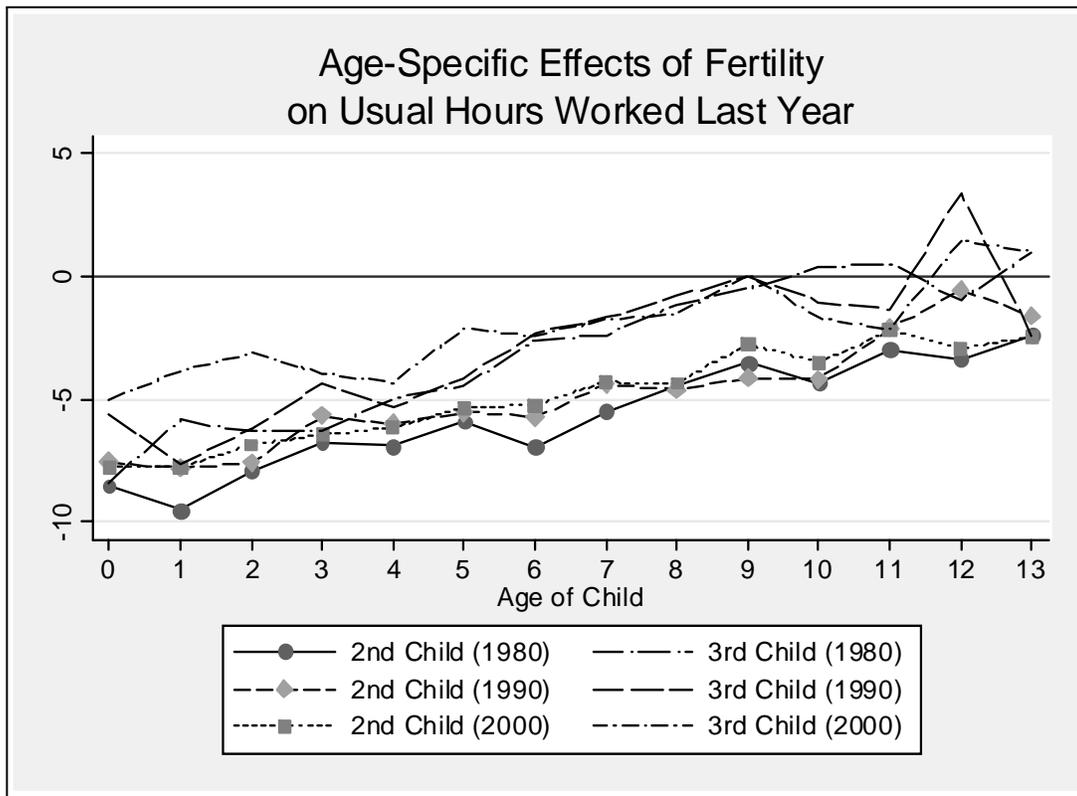


Figure 3: