

DO HIGHLY EDUCATED WOMEN CHOOSE SMALLER FAMILIES?*

Moshe Hazan and Hosny Zoabi

We present evidence that the cross-sectional relationship between fertility and women's education in the US has recently become U-shaped. The number of hours women work has concurrently increased with their education. In our model, raising children and homemaking require parents' time, which could be substituted by services such as childcare and housekeeping. By substituting their own time for market services to raise children and run their households, highly educated women are able to have more children and work longer hours. We find that the change in the relative cost of childcare accounts for the emergence of this new pattern.

Ever since the demographic transition, conventional wisdom suggests that income and fertility are negatively correlated. This has been documented at the aggregate level in a cross-section of countries (Weil, 2005); over time within countries and regions (Galor, 2011) and in cross-sections of households in virtually all developing and developed countries (Kremer and Chen, 2002). Jones and Tertilt (2008) document the relationship between fertility choice and key economic indicators at the individual level for American women born between 1826 and 1960. They found for all cohorts a strong negative cross-sectional relationship between fertility, on the one hand, and income and education of both husbands and wives, on the other. Finally, Preston and Hartnett (2008) and Isen and Stevenson (2010) found similar patterns for cohorts born through the late 1950s.¹

In this article, we present evidence that the cross-sectional relationship between fertility and women's education in the US between 2001 and 2011 is U-shaped. Specifically, we classify women into five educational groups: no high school degree, high school degree, some college, college degree and advanced degree. We start by estimating the total fertility rate (TFR) and show that this measure exhibits a U-shaped pattern. However, estimating TFR by educational group has a drawback in that women are assigned to an educational group according to their educational attainment at the

* Corresponding author: Moshe Hazan, Eitan Berglas School of Economics, Tel Aviv University, P.O. Box 39040, Tel Aviv 69978, Israel. Email: moshehaz@post.tau.ac.il.

We thank Frederic Vermeulen (the editor) and three anonymous referees. We also thank Ghazala Azmat, Alma Cohen, David de la Croix, Alon Eizenberg, Oded Galor, Cecilia Garcia Peñalosa, Jeremy Greenwood, Nezhil Guner, Zvi Hercowitz, Oksana Leukhina, Marco Manacorda, Guy Michaels, Yishay Maoz, Stelios Michalopoulos, Omer Moav, Steve Pischke, Tali Regev, Yona Rubinstein, Analia Schlosser, Christian Siegel, David Weil, Marios Zachariadis, seminar participants at the University of Cyprus, London School of Economics, Paris Seminar in Demographic Economics, Tel Aviv University, and conference participants in the II Workshop on 'Towards Sustained Economic Growth', Barcelona 2011, the Society for Economic Dynamics, Limassol 2012 and the Economic Workshop at IDC, Herzliya 2012. Maor Milgrom provided excellent research assistance.

¹ This inability to establish a positive correlation between income or education and fertility has led some scholars to doubt the assumption that children are a normal good (Jones and Tertilt, 2008; Guinnane, 2011). Black *et al.* (2013) used the exogenous increase in the price of coal during the energy crisis in the mid-1970s to document that men's income in the Appalachian coal-mining region increased and that this led to an increase in fertility.

time of the survey, which may differ a great deal from their completed schooling, especially for young women who are, by and large, still in their schooling period. We circumvent this problem by estimating ‘hybrid fertility rate’ (Shang and Weinberg, 2013). This measure combines children ever born at a specified age and current age-specific-fertility-rates from that age till the end of the fecundity period. We show that this measure also exhibits a U-shaped pattern with respect to education.²

The importance of this pattern depends on the likelihood that the observed U-shaped pattern will be translated into completed fertility rates for cohorts that have not yet completed their fertility. To address this issue, we begin by showing that the U-shaped pattern is a new phenomenon. If it is not, then there is no obvious reason to expect that this pattern will be translated into completed fertility. Indeed, we find that hybrid fertility monotonically decreases in education in 1980 and that this is also true in 1990, although the differential fertility among women with only a college degree and women with advanced degrees declines. In 2000, in contrast, we find that fertility among women with advanced degrees is slightly higher than for women with only a college degree.

Since the U-shaped pattern is indeed new, it is not surprising that it is not yet reflected in completed fertility, even for the youngest cohort for which this measure is available. Nevertheless, it is instructive to look at the fertility of cohorts that have recently completed their fertility. We show that while completed fertility monotonically declines with education for all cohorts, the changes in the cross-sectional relationship across cohorts closely follows changes in the hybrid fertility rates. In particular, the completed fertility of women with an advanced degree increases monotonically across recent cohorts, closing the gap between this group and any other group. This suggests that what we see in hybrid fertility today is likely to be translated into completed fertility in the future.

Standard models of household economics suggest that there is a negative relationship between female labour supply and fertility: women who work more have less time to raise children (Gronau, 1977; Galor and Weil, 1996). In our data, better educated women supply more hours to the labour market. Thus, our findings regarding the patterns of fertility and labour supply, raise two questions:

- (i) what can account for the U-shaped pattern in fertility? and
- (ii) what can account for the positive correlation between fertility and labour supply for highly educated women?

Our explanation relies on the marketisation hypothesis (Freeman and Schettkat, 2005). We argue that highly educated women find it optimal to purchase services such as baby-sitting and day-care as well as housekeeping services to help them run their homes. This enables these women to have more children and work more hours in the labour market. Indeed, Cortes and Tessada (2011) found that

² Shang and Weinberg (2013) studied in detail the fertility of women college graduates. They show that since the late 1990s, the fertility of college graduates has increased over time. The authors do not, however, discuss the cross-sectional relationship between fertility and female education, which is the focus of our article.

- (i) low-skilled immigration has led to an increase in hours worked by women with advanced degrees and that the effects on the labour supply are significantly larger for those with young children;
- (ii) hours spent on household chores declines quite dramatically along the educational gradient; and
- (iii) the fraction of women who use housekeeping services increase sharply with education.

Similarly, Furtado and Hock (2010) found that college-educated women living in metropolitan areas with larger inflows of low-skilled immigrants experience a much smaller tradeoff between work and fertility. Further support for the marketisation hypothesis is provided in Manning (2004) and Mazzolari and Ragusa (2013). Manning (2004) showed that the employment opportunities of unskilled labour depend on physical proximity to skilled workers and Mazzolari and Ragusa (2013) found that growth in a city top wage bill share is associated with significant low-skilled employment growth in the sector of services that substitute for home production activities.

To illustrate our argument, we use a standard model in which a mother derives utility from consumption and the full income of children.³ On the children side, parents decide upon the quantity of children (fertility) and their quality (education). We follow the standard models along two assumptions. First, we assume that education is bought in the market, as in de la Croix and Doepke (2003) and Moav (2005), and show that for highly educated women, education is relatively cheaper than for less educated women, which allows them to purchase more education for their children, even if they allocate the same share of income for quality. Second, as in Hazan and Berdugo (2002) and de la Croix and Doepke (2003), we assume that nature equips children with basic skills. These basic skills imply that as the parents' human capital increases, the share of income they allocate to the quality of each child increases at the expense of the share of income allocated to quantity. This happens because the value of the basic skills in terms of income is relatively high for low income parents. As a result, low income parents find it optimal to spend a relatively large share of income on quantity and a relatively low share on quality. In contrast, for high income parents, the value of the basic skills is relatively small. This induces parents to allocate a higher share of income for quality at the expense of quantity.

To emphasise the reliance on market substitutes for parental time, we deviate from the existing models (Galor and Weil, 1996) by allowing parents to substitute other people's time for their own time by purchasing baby-sitting services in the market.⁴ This marketisation process is an essential element in our mechanism that yields a U-shaped fertility pattern. To see this, ignore for the moment this marketisation channel and assume that quantity requires parental time only. In such a case, with an increase in the parent's human capital, there is an increase by the same proportion in both the parent's income as well as the price for quantity. However, since high income

³ We consider that a household comprises one female parent. Thus, throughout the article, we refer to female parents only, except in subsection 2.3, in which we discuss a two-parent household.

⁴ Aiyagari *et al.* (2002) also allow parents to substitute childcare for their own time. However, in their model, fertility is exogenous and, therefore, they do not study the effect of such services on fertility choice.

parents allocate a lower share of their income to quantity, the optimal number of children monotonically declines.

Marketisation, however, affects the price for quantity that parents face. For parents with low levels of human capital (i.e. low income), marketisation is low and thus the parents themselves engage in most of the child-raising. Thus, the intuition explained above holds. In contrast, parents with high levels of human capital optimally outsource a major part of their child-raising, which, in turn, reduces the cost of children from the parents' point of view. We show that this reduction can be sufficiently large to induce an increase in fertility above a certain level of human capital.

In our basic model, parental time spent on raising children decreases with parents' human capital. This occurs because the fraction of income allocated to raising children decreases with the parents' human capital while parental reliance on market substitutes increases with human capital. However, Guryan *et al.* (2008) found that a mother's time allocated to childcare increases with the mother's education. However, Guryan *et al.* (2008) defined childcare as the sum of four primary time use components: 'basic', 'educational', 'recreational' and 'travel'. Clearly, the educational and recreational components and part of the travel component are an investment in the children's quality. We show that extending the model such that the production of children's quality requires not only education bought in schools but also parental time reconciles our model's predictions with this evidence.⁵

One may suggest an alternative hypothesis to explain the positive association between fertility and female labour supply for highly educated women: spouses of highly educated women work less to compensate for their wives' extra hours in the labour market. To examine this aspect, we discuss in subsection 2.3 an extension of our model to include husbands and allow them to work and raise children. Consistent with Cherchye *et al.* (2012), who studied the allocation of time between labour supply, leisure, home production and childcare in a collective model, we find that the time the wife (husband) allocates to childcare decreases with her (his) human capital. When comparing households, however, one should consider how the human capital of both the spouses varies across households. We argue that assortative matching is sufficient to preserve all of the results found in the basic model. The formal analysis is presented in the online Appendix.

Our theory suggests that the relative price of unskilled labour intensive services, such as childcare and housekeeping, is a key explanatory variable in shaping the relationship between fertility and women's education. Specifically, the marketisation mechanism is more effective when the relative price of these services is lower. To test this empirically, we estimated the cost of childcare services relative to a woman's wage for the period 1983–2012. We found that childcare has become relatively more expensive to women with less than a college degree, but relatively cheaper for women with a college or advanced degree. We then study the association between fertility and the cost of childcare, and find it negative, highly significant, and robust to the inclusion of various controls and different specifications that correct for endogeneity of women's wages and selection bias in the labour market.

⁵ We discuss this extension in subsection 2.3.

Moreover, we show that this structural relationship has been highly stable over the past 30 years.

While these results are important in their own right, we are mostly interested in using them to explain the change over time of the cross-sectional relationship between fertility and women's education. To this end, we estimated a counterfactual cross-sectional relationship between fertility and women's education for the last decade by holding the relative cost of childcare at its early 1980s level. Interestingly, this counterfactual relationship is almost monotonically declining.⁶

The rest of the article is organised as follows. Section 1 presents evidence about the U-shaped fertility pattern. In Section 2, we lay out the model and present the main results of the theory. In Section 3 we study the relationship between fertility and the relative cost of childcare, and explore the implication of the change in the relative cost of childcare for the change in the cross-sectional relationship between fertility and education. In Section 4, we provide evidence about labour supply and marriage rates and rule out alternative hypotheses. Finally, Section 5 provides concluding remarks.

1. Patterns of American Fertility by Education

We used the American Community Survey (ACS) to document basic facts about the fertility behaviour of American women and the correlation between fertility and women's education (Ruggles *et al.*, 2010). The ACS is a suitable survey to study current trends in the fertility of American women since it explicitly asked each respondent whether she gave birth to any children in the past 12 months.

We pooled data from the ACS for the years 2001–11 and restricted our sample to white, non-Hispanic women who live in households under the 1970 definition.⁷ Using these data, we estimated age-specific fertility rates by the five educational groups presented above: no high school diploma, high school diploma, some college, college and advanced degrees.⁸ Figure 1 shows these estimates.

The pattern of these estimates is not surprising. Fertility rates of women who did not complete high school or have a high school diploma peak at ages 20–24. They peak at ages 25–29 for women with some college education and at ages 30–34 for women with college or advanced degrees.⁹

Next, we sum up these age-specific fertility rates to obtain estimates of the TFR. Figure 2 shows that TFR declines for women up to those with some college but then increases for women with college and advanced degrees. Specifically, TFR among

⁶ Our findings are related to Apps and Rees (2004) and Attanasio *et al.* (2008). Attanasio *et al.* (2008) studied the life-cycle labour supply of three cohorts of American women, born in the 1930s, 1940s and 1950s respectively. Their main finding is that the increase in participation early in life for the youngest cohort is the result of a decrease in the childcare cost. Apps and Rees (2004) argued that the cross-country relationship between the female labour supply and fertility, which was negative in 1970, turned positive in 1990 and that tax and child support policies contributed to this reversal.

⁷ Our results are unchanged if we include women of all races but we want to avoid compositional effects coming from changes in the fraction of each race and ethnic group over the period.

⁸ We assign women into educational groups according to their current highest year of school or degree completed. In subsection 1.1 we discuss the potential bias this creates and correct for it.

⁹ We do not report the standard errors of these estimates. Given the sample size, the standard errors on these estimates are essentially zero.

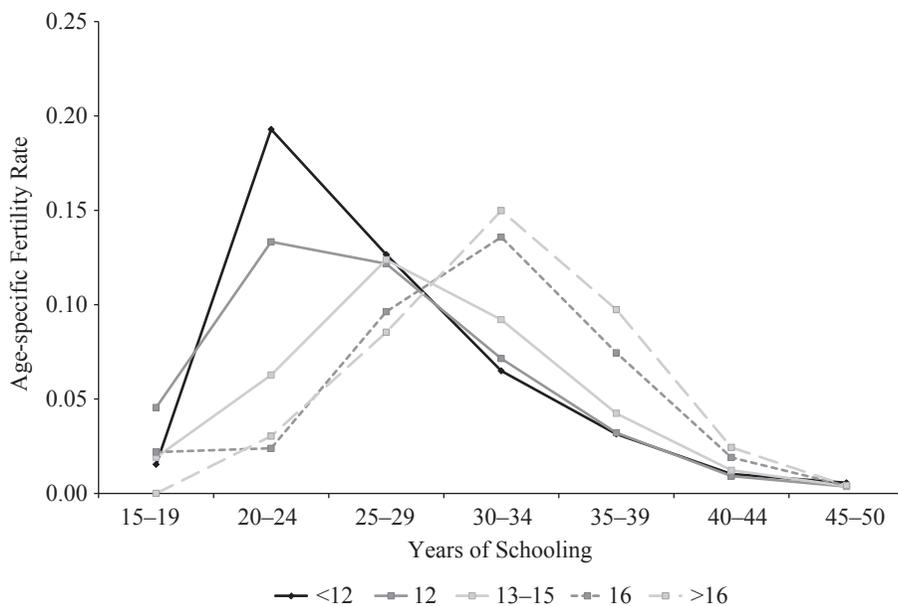


Fig. 1. Age-specific Fertility-rates by Educational Groups, 2001–11
 Source. Authors' calculations using data from the American Community Survey.

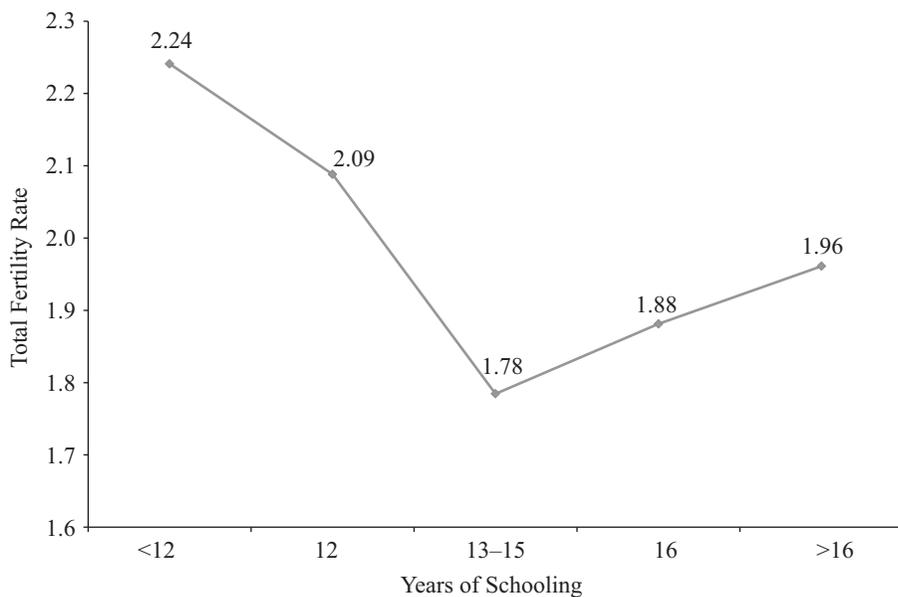


Fig. 2. Total Fertility Rate, 2001–11
 Source. Authors' calculations using data from the American Community Survey.

women with no high school diploma is 2.24; among women with a high school diploma it is 2.09; and 1.78 for women with some college. However, the TFR among women with college degrees is 1.88 and among women with advanced degrees it is 1.96.

This U-shaped fertility pattern raises a few issues. First, how does one deal with the assignment of women into educational groups which is based on current rather than complete schooling? Second, is this pattern robust to differences in the age structure, marital status and family income across women in different educational groups? Third, is the U-shaped pattern a new phenomenon or one that has been overlooked? Finally, and most importantly, will these measures of fertility be translated into completed fertility? In what follows we address each of these questions. We show that our overall analysis paints a picture of an emerging new pattern of fertility by education.

1.1. *The Assignment of Women into Educational Groups*

One concern in our analysis so far is the assignment of women into educational groups. Given the structure of our data, we observe each woman only once and assign women into educational groups according to their educational attainment at the time of the survey, as measured by the highest year of schooling completed or degree attained. While this might not be an issue for relatively older women, it creates strong biases among young women. For example, almost all women aged 15 are currently in high school. This implies that we are assigning all these women to the no high school diploma group even though many of them will undoubtedly end up with a higher level of education. The degree of misassignment, however, declines with the educational group as does the bias. Assuming that women who were mistakenly assigned have a lower fertility rate than those who were properly assigned to their group, then the misassignment may bias the estimated TFR towards a U-shaped pattern even if the true relationship between TFR and education is decreasing.

To address this concern, we estimate a 'hybrid' measure of fertility (Shang and Weinberg, 2013). As noted, the bias may be strong for young women, but is less of a concern for older women. Our hybrid measure uses actual fertility experienced by young women, combined with a period measure of fertility for older women. Specifically, we sum up the number of children ever born to women at age a and the age-specific-fertility rates from age $a + 1$ to age 49. To the extent that women complete their education by age a , all women are assigned to their true educational group. This consideration suggests that we should choose a relatively large a . Such a choice, however, comes with a cost. The higher the a , the larger the weight we put on past fertility compared to current fertility rates. Thus, if fertility rates changed differentially across the educational groups in the 2000s, choosing a relatively large a might prevent us from finding the new pattern, even if it exists.¹⁰ As a compromise, we set $a = 24$.¹¹

Figure 3, which presents this hybrid measure, shows that the U-shaped pattern is still present, albeit the lowest fertility is now attained by women with exactly a college degree. As a check of robustness, we gradually increase a from 24 to 30. We find that the lowest fertility is attained by women with a college degree up to $a = 29$, although

¹⁰ Clearly, choosing a in the 40s, coincides with completed fertility, a measure we discuss in detail in subsection 1.4.

¹¹ The average number of own children in the household at age 24 equals 1.079, 0.77, 0.486, 0.088 and 0.079 for women with no high school degree, exactly a high school diploma, some college, exactly a college degree and an advanced degree respectively.

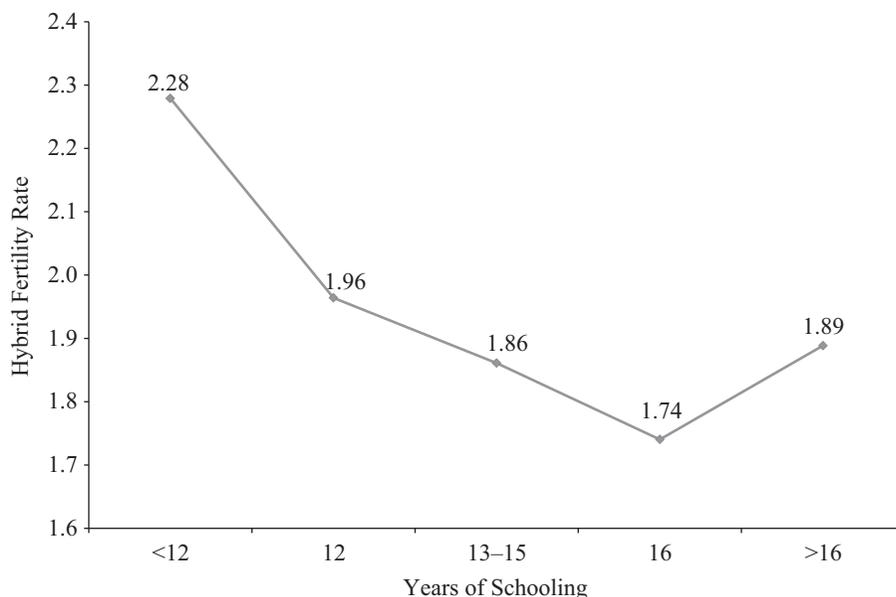


Fig. 3. *Hybrid Fertility Rate, 2001–11*

Notes. The hybrid fertility rate sums up the number of children ever born to women at age a and the age-specific fertility rates from age $a + 1$ to age 49. We assume $a = 24$.

Source. Authors' calculations using data from the American Community Survey.

the difference in fertility between this group and the group of women with an advanced degree declines monotonically. At $a = 30$, the fertility of women with exactly a college degree is larger than that of women with an advanced degree.

One noticeable difference between our estimated TFR (Figure 2) and our estimated hybrid fertility (Figure 3) is that the minimum level of fertility is attained by the some college group and the exactly college degree group respectively. Given the limitations of the data, however, we are unable to determine whether the cross-sectional relationship between completed fertility and women's education will resemble Figures 2 or 3.

1.2. *The Partial Association Between Fertility and Women's Education*

Regression models provide a different means of presenting the association between fertility and women's education. The advantage of this approach is that we can control for various characteristics such as age, marital status, family income, year and state effects that may be responsible for the relationship between fertility and women's education. Table 1 shows the results from linear probability models that take the following structure:

$$b_{ist} = \alpha + e'_{ist}\pi + \kappa N_{ist} + \mathbf{X}'_{ist}\gamma + \delta_a + \delta_m + \delta_t + \delta_s + \epsilon_{ist},$$

where b_{ist} is a dummy variable equal to 1 if woman i living in state s gave birth in year t and 0 otherwise. e'_{ist} is a set of dummy variables that correspond to the five

Table 1
The Association Between Giving a Birth and Women's Education: 2001–11

	(1)	(2)	(3)	(4)	(5)	(6)
High school graduates	0.015*** (0.001)	-0.003*** (0.001)	-0.012*** (0.001)	-0.013*** (0.001)	-0.013*** (0.001)	-0.013*** (0.001)
Some college	0.018*** (0.001)	0.001 (0.001)	-0.015*** (0.001)	-0.015*** (0.001)	-0.016*** (0.001)	-0.016*** (0.001)
College graduates	0.031*** (0.001)	0.008*** (0.001)	-0.007*** (0.002)	-0.007*** (0.002)	-0.008*** (0.002)	-0.008*** (0.002)
Advanced degrees	0.038*** (0.001)	0.012*** (0.001)	0.006*** (0.001)	0.006*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Number of children	0.001** (0.000)	-0.012*** (0.001)	-0.010*** (0.000)	-0.010*** (0.000)	-0.011*** (0.000)	-0.011*** (0.000)
Female earnings: Q1						-0.018*** (0.000)
Female earnings: Q2						-0.019*** (0.001)
Female earnings: Q3						-0.039*** (0.001)
Female earnings: Q4						-0.041*** (0.001)
Spouse earnings						-0.029*** (0.001)
Other income						0.013*** (0.001)
Marital status	No	Yes	Yes	Yes	Yes	-
Age	No	No	Yes	Yes	Yes	Yes
Year	No	No	No	Yes	Yes	Yes
State	No	No	No	No	Yes	Yes
Observations	4,046,532	4,046,532	4,046,532	4,046,532	4,046,532	2,166,054
R ²	0.003	0.022	0.071	0.071	0.071	0.089

Notes: *p < 0.05, **p < 0.01, ***p < 0.001. Robust standard errors adjusted for heteroscedasticity are reported in parentheses. Linear probability models. Women aged 15–50. All models are weighted by American Community Survey sampling weights. The main regressors in columns (1)–(5) are education dummies and the omitted group is high school dropouts. Column 6 focuses instead on female earnings. The omitted group is women without labour income and Q1–Q4 corresponds to the four quartiles of the earnings distribution.

educational groups described above and the coefficients of interest are π . N_{ist} is the number of children woman i has, not including the current birth.¹² X'_{ist} includes four dummies that split women according to their earnings, spouse's wage and other family income.¹³ δ_a are age dummies; δ_m are marital status dummies; δ_t are year dummies; and δ_s are state dummies. The educational group of high school dropouts is the omitted category, so the coefficients on the other educational groups can be interpreted as the difference in the probability of giving a birth relative to that group.

In column (1) we regress b_{ist} only on the educational dummies. Thus, the coefficients in this column are the unconditional differences in the probability of giving birth, namely the general fertility rates (GFR) relative to the GFR among women who do not have a high school diploma. As can be seen, the GFR monotonically increases with education.¹⁴ Column (2) adds dummies for marital status. Since the fraction of currently married women is the lowest for women lacking a high school diploma (see Figure 11 below) and one expects to find higher fertility rates among married women, controlling for marital status should lower the coefficients on education in column (2). Indeed, the coefficients are substantially lower in column (2) than in (1) and in particular, those in the high school diploma and some college groups are now negative rather than positive. The positive coefficients of the college and advanced degree groups imply a U-shaped pattern in fertility rates.

In column (3), we add age dummies. Since age is not monotonically related to fertility rates, the effect on the educational dummies is not predictable. As can be seen in column (3), though, adding age dummies substantially reduces the coefficients of the educational dummies. Now the coefficients of the high school diploma, some college and college graduate groups are negative and significant, while on the advanced degrees it is positive. In column (4) we add year dummies and in column (5) we also add state dummies. Neither the year dummies nor the state dummies change the results of column (3).

Finally, in column (6) we look at the association between female earnings, spouse earnings and fertility. As explained above, the omitted group is that of women without labour income and the coefficients reported in the Table give the difference in the birth rates between women whose labour income is in each of the four quartiles and the omitted group. In this specification, we also control for spouse earnings as well as all other sources of family income. As can be seen from the Table, while the fertility rate is the highest among non-working women, there is

¹² N_{ist} equals the number of own children in the household minus b_{ist} .

¹³ We use female earnings and not wage rate because, Baum-Snow and Neal (2009) argue that in the census and ACS surveys, reports concerning the usual hours worked the past year contain errors that imply incredible wages for part-time workers. The distribution of earnings has a large mass at zero and is then spread over positive values. To account for this, we assign women to five groups. Women without earnings are the omitted groups. Women with positive earnings are assigned into four quartiles.

¹⁴ This may seem at odds with the reported TFR in Figure 2, where TFR is the highest for women without high school diplomas. Notice, however, that TFR sums up age-specific fertility-rates, which are mean births rates within educational-age groups. If women were uniformly distributed across age groups, then the GFR would equal the TFR up to a multiplicative constant. In such a case, both measures would exhibit similar patterns with respect to educational groups.

a clear U-shaped pattern in fertility, where the minimum level of fertility rate prevails at the third quartile earnings group. Notice also that as predicted by economic theory, spouse earnings and other sources of income are positively associated with fertility rates.¹⁵

1.3. *Is the U-shaped Fertility Pattern New?*

As mentioned in the Introduction, many studies have shown that in cross-sections of households, fertility decreases with education. However, since the educational classifications used in these studies are different from ours, we are unable to directly compare our results with those from the literature. For example, had we classified women into three groups of education (no high school diploma, high school graduates and more than high school), we would have found a monotonically decreasing relationship between women's education and fertility as well. Accordingly, in this subsection we use earlier data to show that the U-shaped fertility pattern is indeed only a recent phenomenon.

To demonstrate this observation, we used data from the US Census in 1980, 1990 and 2000 (Ruggles *et al.*, 2010). Unlike the ACS, the census questionnaire does not contain a direct question about the occurrence of a birth over the past 12 months. The census as well as the ACS contain a related question about the age of the youngest own child in the household. One might expect, therefore, that any woman who reported giving a birth during the previous 12 months would respond that the age of the youngest own child in her household is 0.¹⁶ Hence, we construct a variable for a birth during the past 12 months if a woman reports having a child aged 0 years old.¹⁷

Figure 4 presents estimates for hybrid fertility rate for the years 1980, 1990 and 2000. The Figure shows that fertility monotonically decreases in education in 1980. This is also true in 1990, although the slope of the curve decreases substantially (in absolute terms) when moving from women with exactly a college degree to women with an advanced degree. Finally, in 2000, this is no longer true. While fertility decreases up to women with exactly a college degree, it slightly increases for women with an advanced degree. In sum, the evolution of the cross-sectional relationship between fertility rates and women's education over time shows a clear and monotonic increase in the fertility of women with an advanced degree, relative to women with lower levels of education.

¹⁵ The results of these six models are essentially the same if we use a probit instead of a linear probability model. These results are shown in Table B1 in the online Appendix.

¹⁶ Multiple births, infant mortality and handing over a child for adoption or to relatives could create some differences between these two measures, although we conjecture that in practice these occurrences are quantitatively unimportant. Consequently, we assume that discrepancies between the two measures are related to measurement errors.

¹⁷ In the online Appendix we check the reliability of this measure in the ACS data, which contains the response to both questions. We show that while the estimates of hybrid fertility based on the age of youngest child are systematically smaller, the gap between the two series is almost constant across the educational groups.

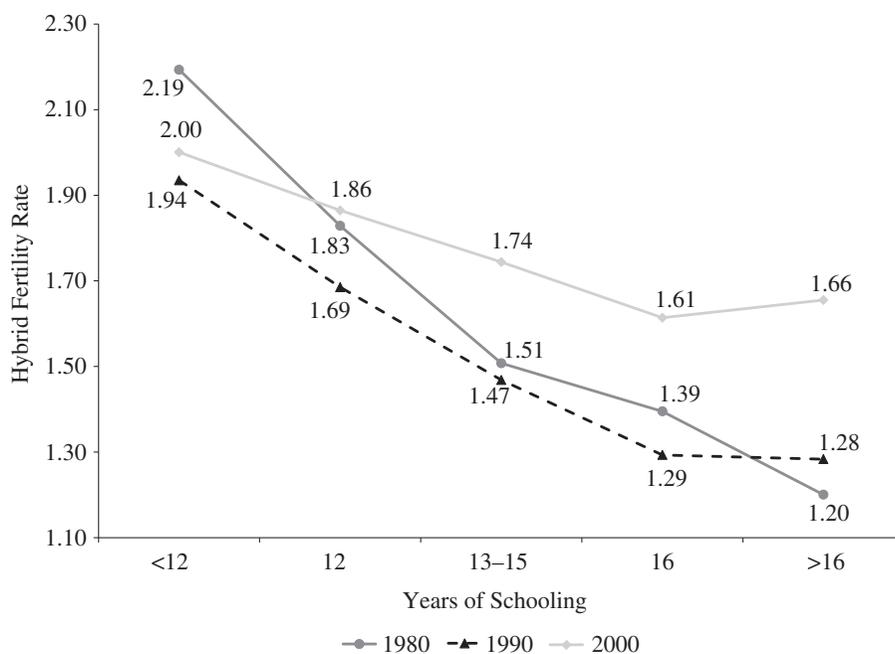


Fig. 4. *Hybrid Fertility Rate, 1980, 1990 & 2000*

Notes. The hybrid fertility rate sums up the number of children ever born to women at age a and the age-specific fertility rates from age $a + 1$ to 49. We assume $a = 24$.

Source. Authors' calculations using US Census data.

1.4. *Hybrid and Completed Fertility Rates*

Although our analysis is mostly concerned with hybrid fertility rates, our objective is to argue that the current patterns in hybrid fertility rates are likely to be translated into completed fertility rates for cohorts that have not yet completed their fertility.¹⁸ Since completed fertility is estimated for women approaching the end of their fertile period, usually taken to be 40–44 years of age, the new patterns exhibited in Figures 2, 3 and 4 are still not reflected in the completed fertility rate even for the youngest cohorts that have reached this age.

It is constructive, however, to look at the pattern of the completed fertility rate by education for cohorts who have recently reached the end of their fertile period. Using data from the 1990 Census as well as from the Fertility Supplement of the June Current Population Survey for the years 1995, 2000, 2004 and 2008, we estimate completed fertility by education for women aged 40–44. This covers the cohorts born between 1946 and 1968. These estimates are shown in Figure 5.

Two features in Figure 5 are worth mentioning. First, for all cohorts, completed fertility monotonically declines across the educational groups. Second, across cohorts, the curves shift counter-clockwise around the some college group. This feature

¹⁸ Preston and Hartnett (2008) showed that with the exception of the baby-boom period, TFR and completed fertility rates in the US almost coincide during the twentieth century.

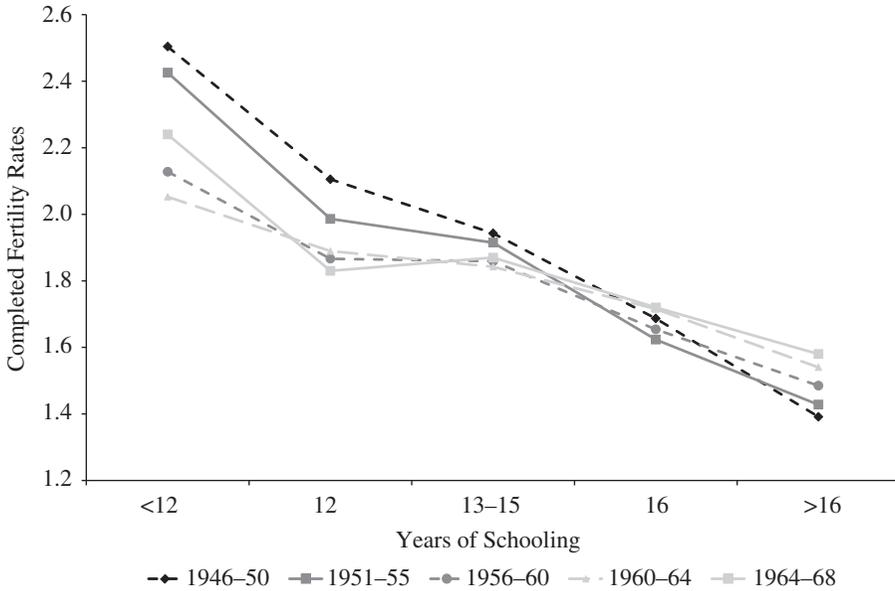


Fig. 5. Completed Fertility Rates by Education for Cohorts Born Between 1946 and 1968
 Source. Authors' calculations using data from the 1990 Census and the Fertility Supplement of the June Current Population Survey for the years 1995, 2000, 2004 and 2008.

supports our conjecture as differential fertility between the least and the most educated groups of women contracts and the level of fertility for women with advanced degrees monotonically increases across cohorts. Thus, even if we never see the U-shape in completed fertility, the marketisation hypothesis proposed below may explain the flattening of the relationship between education and fertility.

2. The Model

2.1. Structure

There is a continuum of mass one of adult individuals that differ by their level of human capital. Each individual forms a household, works, and chooses consumption and her number of children. Children are being raised and educated. Education is provided by the market through schools. To raise children, households combine the parent's time and time purchased in the market. Likewise, households combine parent's time, time purchased in the market along with a market good to produce the consumption good. This market good serves as the numeraire. Finally, the remaining time is allocated to labour market participation.

Let h_i denote the human capital of individual i , which also equals her market productivity. The preferences of household i are defined over consumption, c_i , and total full income of the children, $n_i h'_i$. They are represented by the utility function:

$$u_i = \ln(c_i) + \ln(n_i h'_i). \tag{1}$$

The budget constraint is:

$$h_i = p_{ci}c_i + p_{ni}n_i + n_i p_{ei}e_i, \quad (2)$$

where p_{ci} , p_{ni} and p_{ei} are the prices of consumption, quantity of children and children's education, e_i , faced by parent i , respectively.

Children's human capital, h'_i , is determined by their level of education, e_i , and basic skills with which nature equips each child, $\eta > 0$, regardless of her parent's characteristics. The human capital production function is:

$$h'_i = (e_i + \eta)^\theta, \quad \theta \in (0, 1). \quad (3)$$

Education is provided by schools. We assume that the average level of human capital among teachers is \bar{h} . As all parents face the same market price for education, $p_{ei} = p_e = \bar{h}$ the cost of educating n_i children at the level e_i is given by:

$$TC_i^e = n_i p_e e_i = n_i \bar{h} e_i. \quad (4)$$

Raising children requires time independent of education. The time required to raise n children can be supplied by the parent or bought in the market, e.g., childcare or baby sitters. The production function of raising n children is:

$$n = (t_M^n)^\phi (t_B^n)^{1-\phi}, \quad \phi \in (0, 1) \quad (5)$$

where t_M^n is the time devoted by the mother and t_B^n is the time bought in the market, e.g., a babysitter.¹⁹ We assume that the price of one unit of time bought in the market is some level of human capital denoted by \underline{h} . This implies that \underline{h} is the average human capital among babysitters.

The cost of raising n children is, therefore, given by the cost function,

$$TC^n(n, \underline{h}, h^i) = \min_{t_M^n, t_B^n} \{t_M^n h_i + t_B^n \underline{h} : n = (t_M^n)^\phi (t_B^n)^{1-\phi}\}.$$

The optimal t_M^n and t_B^n are:

$$t_M^n = \left(\frac{\phi}{1-\phi} \frac{\underline{h}}{h_i} \right)^{1-\phi} n \quad (6)$$

and

$$t_B^n = \left(\frac{1-\phi}{\phi} \frac{h_i}{\underline{h}} \right)^\phi n. \quad (7)$$

Using these optimal levels, we obtain the cost function:

$$TC^n(n, \underline{h}, h^i) = p_{ni} n = \varphi \underline{h}^{1-\phi} h_i^\phi n, \quad (8)$$

where $\varphi \equiv [\phi^\phi (1-\phi)^{1-\phi}]^{-1}$.

It should be noted from (8) that the marginal cost of raising children is constant. Moreover, this marginal cost increases with the mother's human capital, although its elasticity with respect to the mother's human capital is $\phi < 1$.

¹⁹ This modelling approach is similar to Greenwood *et al.* (2005).

Following Becker (1965), the consumption good that directly enters the utility function is produced by combining time and a market good. The time allocated to this production can be either supplied by the mother or purchased in the market. The production function is:

$$c = m^{1-\alpha} [(t_M^c)^\sigma + (t_H^c)^\sigma]^{\alpha/\sigma}, \quad \sigma \in (0, 1)$$

where m is the market good and $1/(1-\sigma) > 1$ is the elasticity of substitution. That is, t_M^c and t_H^c are assumed to be gross substitutes. This assumption captures the idea that a mother's time and the time of a housekeeper are highly substitutable.²⁰ We assume that the price of one unit of time bought in the market is \hat{h} . This implies that \hat{h} is the average human capital among housekeepers.

The cost of c units of consumption is, thus, given by the cost function

$$TC^c(c, \hat{h}, h^i) = \min_{m, t_M^c, t_H^c} \{m + t_M^c h_i + t_H^c \hat{h} : c = m^{1-\alpha} [(t_M^c)^\sigma + (t_H^c)^\sigma]^{\alpha/\sigma}\}.$$

The optimal t_M^c and t_H^c are:

$$t_M^c = \frac{\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha}}{h_i^{1-\alpha} \left[1 + \left(\frac{h_i}{\hat{h}}\right)^{\frac{\sigma}{1-\sigma}}\right]^{1+\alpha\left(\frac{1}{\sigma}-1\right)}} c \tag{9}$$

and

$$t_H^c = \frac{\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} h_i^{\alpha+\frac{\sigma}{1-\sigma}}}{\hat{h}^{\frac{1}{1-\sigma}} \left[1 + \left(\frac{h_i}{\hat{h}}\right)^{\frac{\sigma}{1-\sigma}}\right]^{1+\alpha\left(\frac{1}{\sigma}-1\right)}} c. \tag{10}$$

Substituting these optimal factors into the cost function yields:

$$TC^c(c, \hat{h}, h^i) = p_c c = \frac{h_i^\alpha}{\omega \left[1 + \left(\frac{h_i}{\hat{h}}\right)^{\frac{\sigma}{1-\sigma}}\right]^{\alpha\left(\frac{1}{\sigma}-1\right)}} c, \tag{11}$$

where $\omega = \alpha^\alpha (1-\alpha)^{1-\alpha}$.

2.2. *Equilibrium*

Given the prices for quality of children, quantity of children and consumption in (4), (8) and (11) respectively, the solution to maximising

²⁰ Note that we assume that in producing the consumption good, the mother's time and the housekeeper's are more substitutable than the mother's time and the baby-sitter's time in raising children. This assumption can be justified by noting that pregnancy and breastfeeding are less substitutable than cleaning and cooking. For example, Sacks and Stevenson (2010) reporting that during the 2000s, mothers on average spend well over 2 hours a day breastfeeding their infants.

- (i) subject to the budget constraint;
(ii) yields:

$$e_i = \begin{cases} 0 & \text{if } h_i \leq \left(\frac{\eta \bar{h}}{\theta \varphi \underline{h}^{1-\phi}} \right)^{\frac{1}{\phi}} \equiv h_e \\ \frac{\theta \varphi \underline{h}^{1-\phi} h_i^\phi - \eta \bar{h}}{\bar{h}(1-\theta)} & \text{otherwise.} \end{cases} \quad (12)$$

Notice that for a parent with low human capital, η could be large enough that the optimal level of education is zero. We ignore, henceforth, this corner solution by assuming that the lowest level of parental human capital is above h_e . Consequently, the optimal level of fertility is given by:

$$n_i = \frac{h_i(1-\theta)}{2(\varphi \underline{h}^{1-\phi} h_i^\phi - \eta \bar{h})} \quad (13)$$

and

$$c_i = \frac{\omega}{2} h_i^{1-\alpha} \left[1 + \left(\frac{h_i}{\hat{h}} \right)^{\frac{\alpha}{1-\sigma}} \right]^{\alpha(\frac{1}{\sigma}-1)}. \quad (14)$$

Equations (6), (7), (9), (10), (12), (13) and (14) yield the following seven propositions.

PROPOSITION 1. *The educational choice, e , strictly increases with h_i for all $h_i > h_e$.*

Proof. Follows directly from differentiating (12) with respect to h_i .

The intuition behind this result is straightforward. With a log linear utility function from consumption and full income of the children, the optimal level of education is independent of the parent's human capital since any additional unit of education is given to all children equally. Moreover, since any additional child will be given the same education as her siblings, the optimal level of education depends negatively on the price of education (quality) relative to fertility (quantity).

The value of parental time is equal to her human capital. While quality is bought in the market at a given cost, independently of the parent's human capital, quantity requires some of the parent's time and, thus, its price positively depends on the parent's human capital. Consequently, the relative price of quality declines in the parent's human capital, yielding a higher investment in education.

Notice that as the parent's human capital increases, the share of income that is allocated to the quality of each child increases at the expense of the share of income allocated to quantity. The intuition for this is simple. For low income parents, the basic skill, η , which is equivalent to $\eta \bar{h}$ in terms of income, is relatively important. As a result, parents find it optimal to invest a large share of income in quantity and a low share in quality. In contrast, for high income parents, the value of the basic skill in terms of

income, $\eta\bar{h}$, is relatively small, which induces parents to allocate a higher share of income for quality at the expense of quantity.

PROPOSITION 2. *The fertility choice, n , is U-shaped as a function of h_i .*

Proof. Differentiating (13) with respect to h_i yields:

$$\frac{\partial n}{\partial h_i} = \frac{(1 - \theta)[(1 - \phi)\varphi\underline{h}^{1-\phi}h_i^\phi - \eta\bar{h}]}{2(\varphi\underline{h}^{1-\phi}h_i^\phi - \eta\bar{h})^2}.$$

Thus,

$$\frac{\partial n}{\partial h_i} \begin{cases} < 0, & \text{for } h_i < \tilde{h} \\ = 0, & \text{for } h_i = \tilde{h} \\ > 0, & \text{for } h_i > \tilde{h} \end{cases}$$

where

$$\tilde{h} = \left[\frac{\eta\bar{h}}{(1 - \phi)\varphi\underline{h}^{(1-\phi)}} \right]^{\frac{1}{\phi}}.$$

The intuition behind this result is as follows. As described above, the optimal level of education depends on the relative price of quality and the basic skill. Fertility, however, depends on the share of income allocated to quantity and the price of an additional child. As explained above, the share of income allocated to quantity decreases with the parent's human capital. We now analyse how the price for quantity changes with the parent's human capital to determine the optimal level of quantity.

Marketisation is an essential element in our mechanism that yields the U-shaped fertility pattern. Let us ignore for the moment the marketisation channel and assume that quantity requires parents' time only. In this case, with an increase in the parent's human capital, both the parent's income and the price for quantity increase by the same proportion. Since parents allocate a lower share of their income to quantity, the optimal number of children monotonically declines.

Marketisation, however, affects the price for quantity that parents face. For parents with low levels of human capital (i.e. low income), marketisation is low and the parents do most of the child-raising. Thus, the intuition above holds. Parents with high levels of human capital, in contrast, outsource a major part of child-raising, which, in turn, reduces the price of children from the parents' point of view. This reduction might be sufficiently large to induce an increase in fertility.

Notice from (8) that the price of quantity is $\varphi\underline{h}^{1-\phi}h_i^\phi$. Thus, although it increases with the parents' human capital, marketisation causes this price to increase at a lower pace than income.²¹ Thus, for all $h_i > \tilde{h}$, marketisation implies that the share of income allocated to quantity decreases at a lower pace than price does, causing fertility to increase.

²¹ Notice that the Cobb–Douglas production function for quantity is not crucial for this result. It can be easily shown that this result holds for any CES production function.

PROPOSITION 3. *Mother's time spent on raising children (quantity), t_M^n , strictly decreases with income, h_i .*

Proof. Substituting (13) into (6) gives:

$$t_M^n = \frac{(1-\theta)}{2} \left(\frac{\phi}{1-\phi} \right)^{1-\phi} \frac{\underline{h}^{1-\phi} h_i^\phi}{(\phi \underline{h}^{1-\phi} h_i^\phi - \eta \bar{h})}, \quad (15)$$

differentiating (15) with respect to h_i yields:

$$\frac{\partial t_M^n}{\partial h_i} = -\phi \left(\frac{\phi}{1-\phi} \right)^{1-\phi} \frac{(1-\theta)}{2} \frac{\eta \bar{h} (\underline{h}/h_i)^{1-\phi}}{(\phi \underline{h}^{1-\phi} h_i^\phi - \eta \bar{h})^2} < 0.$$

The intuition here is straightforward. First, with a log-linear utility function as given in (1), the share of resources allocated to children is one-half. Second, as discussed above, the share of income allocated to quantity is declining in h_i . Finally, since childcare and the mother's time are aggregated using a homothetic production function, the share of income allocated to each one of these two factors is independent of h_i . Thus, the parents' time that is allocated to quantity declines with the mother's education. In subsection 2.3 below, we discuss an extension to the model in which mother's time is also used for producing child quality. This allows the mother's total time spent on children to increase, which is consistent with the empirical findings from the time-use data (Guryan *et al.*, 2008; Ramey and Ramey, 2010).

PROPOSITION 4. *Mother's time spent on home production, t_M^c , strictly decreases with income, h_i .*

Proof. Substituting (14) into (9) yields

$$t_M^c = \frac{\alpha}{2[1 + (h_i/\hat{h})^{\frac{\sigma}{1-\sigma}}]}, \quad (16)$$

which is, unambiguously, decreasing in h_i

Since the consumption good is a Cobb–Douglas aggregate of the market good and time, the share of resources allocated to each one of these factors is independent of h_i . However, the assumed gross substitutability between a mother's time and a housekeeper's time yields a declining time spent by the mother as its price, h_i , increases.

PROPOSITION 5. *The labour supply, $l \equiv 1 - t_M^n - t_M^c$, strictly increases with mother's income, h_i .*

Proof. Follows directly from Propositions 3 and 4.

PROPOSITION 6. *The amount of baby-sitter services purchased in the market, t_B^n , is strictly increasing with income for all $h_i \geq [(1+\phi)\eta\bar{h}/\phi\underline{h}^{1-\phi}]^{1/\phi} \equiv h_B$.*

Proof. Follows directly by substituting (13) into (7) and differentiating with respect to h_i .

PROPOSITION 7. *The amount of housekeeping services purchased in the market, t_H^c , strictly increases with the mother's income, h_i .*

Proof. Follows directly from substituting (14) into (10) and differentiating with respect to h_i .

As we show in subsection 3.2, purchasing childcare services monotonically increases with women's education. Hence, we would like to verify that there is a range of h_i in which our model can concurrently generate

- (i) $\partial e_i / \partial h_i > 0$;
- (ii) $\partial t_B^n / \partial h_i > 0$; and
- (iii) n_t exhibits a U-shaped relationship with h_i .

Note that (i) requires that $h_i > h_e$, (ii) requires that $h_i > h_B$ and (iii) requires that $\tilde{h} > \max\{h_e, h_B\}$.

Comparing \tilde{h} and h_B , it follows that \tilde{h} is always larger than h_B . Thus, it is sufficient to require that h_e be smaller than h_B , a condition which is satisfied if and only if $1/(1 + \phi) < \theta$. Hence, we assume that the lowest level of parental human capital is above h_B .²²

2.3. Extensions

In this subsection we discuss two extensions to our basic model. The analysis is performed in the online Appendix. The purpose of the first extension is to show that our model can account for the positive correlation between a mother's education and time spent with children as found by Guryan *et al.* (2008).²³ However, Guryan *et al.* (2008) defined childcare as the sum of four primary time-use components: 'basic', 'educational', 'recreational' and 'travel'. Some of these components represent an investment in the children's quality, a component which, in our model, is purchased in the market.

Ramey and Ramey (2010) reconcile the seemingly paradoxical allocation of time, according to which mothers with a higher opportunity cost of time spend more, rather than less time with their children despite the availability of market substitutes. They argue that as slots in elite post-secondary institutions have become scarcer, parents responded by investing more in their children's quality so that they appear more desirable to college admissions officers. This implies that parental time and market goods and services are strong complements in the production of the children's quality. In the online Appendix we incorporate such complementarity and show that our model preserves all of its results while being consistent with this stylised fact as well.

²² Finally, since h_B is a function of \underline{h} and \bar{h} , we should ensure that \underline{h} and \bar{h} are larger than h_B . From the definition of h_B , it follows that if $\bar{h}/\underline{h}^{1-\phi}$ is constant, then h_B is independent of \underline{h} and \bar{h} .

²³ Table 2 in Guryan *et al.* (2008) reports that the hours per week spent in total childcare are 12.1, 12.6, 13.3, 16.5 and 17 for mothers with <12, 12, 13–15, 16 and 16+ years of schooling respectively.

The second extension incorporates husbands into our unitary household framework. We do so because we want to examine the extent to which husbands of more educated wives could substitute for their wives in raising children. Cherchye *et al.* (2012) studied the allocation of time between labour supply, leisure, home production and childcare in a collective model and found that the time of the husband that is allocated to childcare increases with his wife's wage. In our extended model, the husband's time is optimally allocated between child raising and labour supply. It turns out that positive assortative matching is sufficient to ensure that mothers with higher human capital will purchase more baby-sitting services.²⁴ Consistent with Cherchye *et al.* (2012) and our model, we present evidence below that, indeed, husbands of highly educated women spend more time on childcare but that these households also purchase more childcare services.

3. Fertility and Childcare Over Time

In the previous Section we showed that our theory, which rests on the marketisation hypothesis accounts quite well for the qualitative features of the period 2001–11. Accordingly, childcare and housekeeping services, which are relatively cheaper for highly educated women, enabled these women to have more children and work more than women with intermediate levels of education. However, these services were available in earlier periods as well, when the relationship between fertility and education was monotonically decreasing. Thus, we need to explore if the key explanatory variables in our theory have changed over time in a way that can account for the changing relationship between fertility and education.

3.1. What Drives the Change in the Relationship Between Fertility and Education?

The relationship between fertility and education in our model is governed by the cost of childcare, \underline{h} , relative to mother's income, h_i . Specifically, the lower this ratio is, the larger the optimal fertility is. To explore this idea in a systematic way, we constructed a variable to measure this ratio. Using data from the March CPS (King *et al.*, 2010) for the period 1983–2012, we estimate the average hourly wage in the 'child day-care services' industry and allow it to vary by state and year. We denote this measure by w_{st}^{cc} .²⁵ This variable should proxy for the (absolute) cost of childcare in state s and year t .²⁶ In addition, we compute the hourly wage of all women in the 25–50 year-old age group who reported a positive salary income and denote it by w_{ist} . We then compute the relative cost of childcare by taking the ratio between the two variables.²⁷ Figure 6 presents the fitted values of the average of this variable for each of our five educational

²⁴ Assortative matching is a stylised fact of the marriage market. Charles *et al.* (2013) show assortative matching on parental wealth and Pencavel (1998) on spousal education.

²⁵ The industry 'Child day care services' is available only from 1983. In principle, we should have $51 \times 30 = 1,530$ year-state cells. In practice, we have only 1,520 because 10 state-year cells have no observations.

²⁶ We use the word proxy because it measures only the labour cost component of childcare.

²⁷ To measure the change in the probability of giving birth in response to percentage change in the relative cost of childcare, we take the log of this ratio.

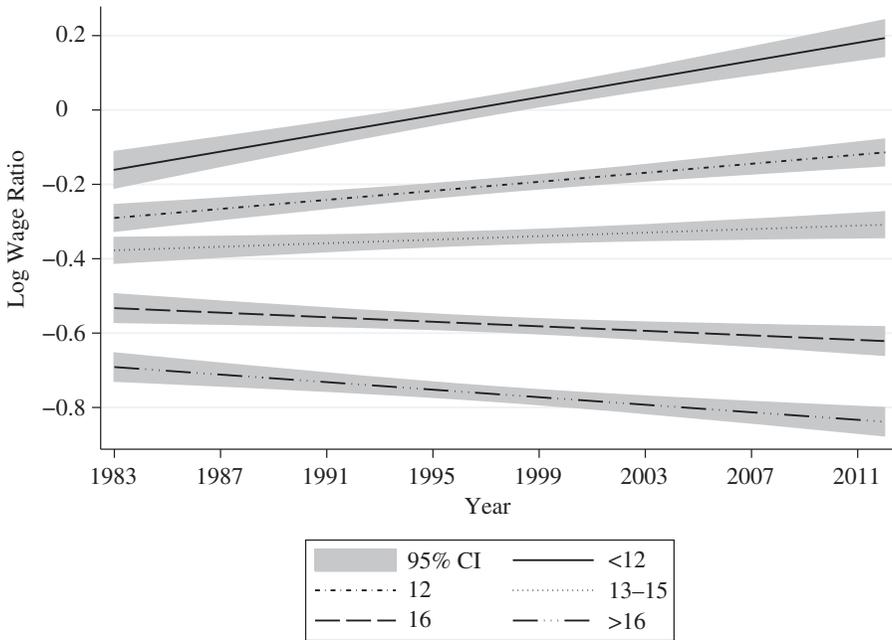


Fig. 6. *Linear Prediction of the Log of the Ratio of Average Wage in the Childcare Industry to Average Wage in the Five Educational Groups 1983–2012*
 Source. Authors' calculations using data from the March CPS.

groups. The Figure shows that childcare has become relatively more expensive for women with less than a college degree but relatively cheaper for women with a college or an advanced degree. Note that the changes are quantitatively large. Over the 30 years between 1983 and 2012, the relative childcare cost has increased by 33%, 16.5% and 5.2% for women with no high school diploma, high school degree and some college respectively. In contrast, this relative cost decreased by 9% for women with a college degree and by 15.5% for women with an advanced degree.

With this measure in hand, we can estimate models, similar to the models in subsection 1.2. Specifically, we estimate models of the form:

$$b_{ist} = \alpha + \beta \ln \left(\frac{w_{st}^{cc}}{w_{ist}} \right) + \kappa N_{ist} + \mathbf{X}'_{ist} \times \gamma + \delta_a + \delta_m + \delta_t + \delta_s + \epsilon_{ist},$$

where b_{ist} is a dummy equal to 1 if a woman i living in state s gave birth in year t and 0 otherwise, $\ln(w_{st}^{cc}/w_{ist})$ is the log of the ratio between the average wage paid to workers in the childcare industry in state s in year t and the wage of woman i , living in state s in year t . N_{ist} is the number of children woman i has, not including the current birth. \mathbf{X}'_{ist} includes total personal income, total personal income square and spouse's wage. δ_a , δ_m , δ_t and δ_s are age, marital status, year and state dummies respectively.

The key parameter of interest is β which measures the change in the probability of giving birth in response to a 1% change in the relative cost of childcare. Since the log of relative cost varies at the state-year level, we cluster the standard errors at the state level. Table 2 shows the result of estimating these models. As can be seen from models

Table 2
The Association Between Giving a Birth and Childcare Relative Cost: 1983–2012

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable: birth in the past 12 months							
Childcare relative cost	-0.008*** (0.001)	-0.012*** (0.001)	-0.011*** (0.001)	-0.010*** (0.000)	-0.011*** (0.000)	-0.023*** (0.001)	-0.032*** (0.002)
Number of children	-0.008*** (0.000)	-0.003*** (0.001)	-0.014*** (0.000)	-0.014*** (0.000)	-0.015*** (0.000)	-0.015*** (0.000)	-0.020*** (0.001)
Total personal income						-0.075*** (0.008)	-0.108*** (0.012)
Total personal income ²						0.025*** (0.006)	0.030*** (0.010)
Spouse's wage							0.402*** (0.053)
Age dummies	No	Yes	Yes	Yes	Yes	Yes	Yes
Marital status dummies	No	No	Yes	Yes	Yes	Yes	–
Year dummies	No	No	No	Yes	Yes	Yes	Yes
State dummies	No	No	No	No	Yes	Yes	Yes
Observations	514,829	514,829	514,829	514,829	514,829	514,829	305,847
R ²	0.003	0.038	0.064	0.065	0.066	0.068	0.079

Notes: *p < 0.05, **p < 0.01, ***p < 0.001. Robust standard errors adjusted for heteroscedasticity and clustered at the state level are reported in parentheses. Linear probability models. All models are weighted by CPS sampling weights. Childcare relative cost is the log of the average wage in the child day care services, varied at the state-year level, relative to mother's wage.

(1) to (5), the coefficient is nearly unchanged by the inclusion of age, marital status, year and state dummies. In model (6), we include total personal income and total personal income square, measured in hundreds of thousands of 1999 dollars. Notice that controlling for total personal income roughly doubles β . Finally, model (7), which controls for spouse's wage, expressed in thousands of 1999 dollars, further increases the magnitude of β by another 50% (in absolute terms).

While the results in Table 2 strongly support our theory, there are several potential problems. First, the fact that wages are observed only for working women raises a selection bias problem.²⁸ Second, the wage we observe may be endogenous to the decision to have a baby. For example, the hourly wage during the year a woman is giving birth may be lower than her wages in other years because of a weaker attachment to the labour market or poorer health due to the pregnancy. In the online Appendix, we explain how we correct for selection bias and endogeneity of wages and report our estimates in Table B3. The Table shows that the estimates are all highly statistically significant and similar in magnitude to the estimate in column (7) in Table 2.

Another potential concern might be that we pool data for 30 years and that the relationship between the probability of giving birth and the relative cost of childcare may be driven by a sub-period. Table 3 shows the results of estimating the model in column (7) in Table 2 separately for each three consecutive years from 1983–85 to 2010–12. As can be seen from the Table, the estimates of β are all highly statistically significant and highly stable over these 30 years.²⁹

We can use the estimates of β to estimate the counterfactual hybrid fertility rate in 2001–11 under the 1983–85 relative childcare cost. The change in the hybrid fertility rate for each educational group j that is due to the change in the relative cost of childcare for this group is given by:

$$\Delta F_j = \beta [\ln(w^{cc}/w)_{j_{t_1}} - \ln(w^{cc}/w)_{j_{t_0}}] \times 26,$$

where ΔF_j is the change in hybrid fertility rate, t_1 is 2010–12 and t_0 is 1983–85. Recall that b_{ist} is the probability of giving birth at a given age over a horizon of 26 years of a woman's fertile period.

Figure 7 shows our baseline hybrid fertility (the dark solid line) and adds the counterfactual hybrid fertility measure obtained by subtracting ΔF_j using the estimate of β from model (7) in Table 2 (the dark dashed line). The Figure shows that the counterfactual fertility curve is obtained by a clockwise rotation of the hybrid fertility curve around the some college education group.³⁰ Specifically, had childcare costs for women with a college degree and women with advanced degrees been constant, their fertility would have been lower by 0.07 and 0.13 respectively. Notice that while the counterfactual fertility is still U-shaped, it is less pronounced.

Our discussion above assumes that the impact of the relative childcare cost on a woman's decision to give birth is independent of her level of education. However, this

²⁸ Mulligan and Rubinstein (2008) found a positive selection in the female workforce since the 1990s.

²⁹ We repeat the results reported in Table 3 using the measure of the relative cost of childcare used in Column 4 of Table B2 that corrects for selection bias and wage endogeneity and found a negative and statistically significant coefficient in each three-year sample. These results are reported in Table B3.

³⁰ Note that this clockwise rotation is a mirror image of the counter-clockwise rotation we observed in completed fertility shown in Figure 5.

Table 3
The Association Between Giving a Birth and Childcare Relative Cost: Each Three Consecutive Years 1983–2012

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	1983–85	1986–88	1989–91	1992–94	1995–97	1998–00	2001–3	2004–6	2007–9	2010–12
Dependent variable: birth in the past 12 months										
Childcare relative cost	-0.030*** (0.005)	-0.034*** (0.003)	-0.038*** (0.005)	-0.037*** (0.004)	-0.032*** (0.004)	-0.037*** (0.005)	-0.042*** (0.003)	-0.043*** (0.003)	-0.039*** (0.005)	-0.037*** (0.004)
Number of children	-0.019*** (0.002)	-0.019*** (0.001)	-0.020*** (0.001)	-0.020*** (0.001)	-0.020*** (0.002)	-0.020*** (0.002)	-0.016*** (0.001)	-0.021*** (0.002)	-0.026*** (0.002)	-0.024*** (0.002)
Total personal income	-0.218*** (0.028)	-0.197*** (0.025)	-0.218*** (0.028)	-0.208*** (0.026)	-0.165*** (0.025)	-0.135*** (0.028)	-0.140*** (0.012)	-0.175*** (0.017)	-0.114*** (0.018)	-0.069*** (0.016)
Total personal income ²	0.143*** (0.026)	0.131*** (0.024)	0.146*** (0.028)	0.157*** (0.029)	0.106*** (0.021)	0.049** (0.017)	0.037*** (0.007)	0.079*** (0.013)	0.029*** (0.008)	0.009* (0.004)
Spouse's wage	0.426* (0.160)	0.415** (0.123)	0.035 (0.213)	0.533** (0.155)	0.428** (0.124)	0.435* (0.173)	0.247** (0.083)	0.187 (0.114)	0.439** (0.139)	0.283* (0.108)
Observations	26,431	28,114	29,702	29,590	26,602	23,789	41,371	36,873	33,945	30,018
R ²	0.072	0.073	0.075	0.075	0.080	0.080	0.075	0.096	0.103	0.098

Notes. *p < 0.05, **p < 0.01, ***p < 0.001. Linear probability models. All models are weighted by CPS sampling weights. Robust standard errors adjusted for heteroscedasticity and clustered at the state level are reported in parentheses. All models include age, year and state dummies. See note to Table 2 for further details.

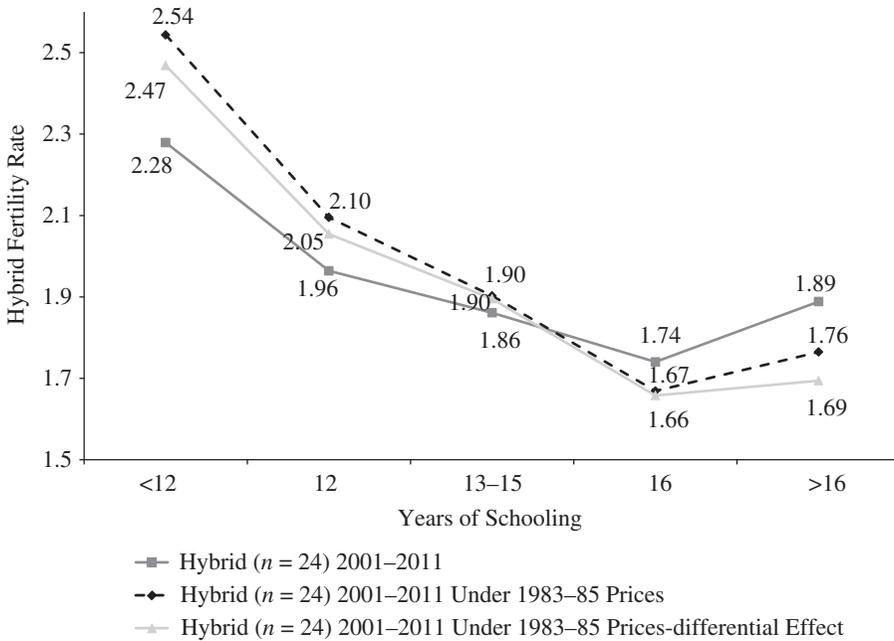


Fig. 7. Hybrid Fertility 2001–11, Counterfactual: Hybrid Fertility 2001–11 under 1983–85 Prices, Counterfactual: Hybrid Fertility 2001–11 under 1983–85 Prices – Differential Effects for Each Educational Group

Note. See text for more details.

restricted model ignores other dimensions that may affect the relationship between the decision to give birth and childcare costs. Indeed, one may assume that women care about pursuing a career and that this aspiration increases with women’s education. To illustrate this, assume that there are two types of women: uneducated women who do not care about pursuing a career and educated women who do. For the first type, the reduction in the relative cost of childcare has a pure price effect. For the second type, there is an additional effect that stems from a reduction in the rivalry between children and career. Thus, a reduction in the childcare cost should have a larger effect on the probability of more educated women giving birth. To explore this possibility, we estimate models that allow for differential effects of childcare cost of the following form:

$$b_{ist} = \alpha + \sum_{j=2}^5 \pi_j e_{ist}^j + \beta \ln \left(\frac{w_{st}^{cc}}{w_{ist}} \right) + \sum_{j=2}^5 \gamma_j e_{ist}^j \ln \left(\frac{w_{st}^{cc}}{w_{ist}} \right) + \kappa N_{ist} + \delta_a + \delta_m + \delta_t + \delta_s + \varepsilon_{ist},$$

where e_{ist}^j are educational group dummies equal to 1 if woman i is in the j educational group and 0 otherwise. Now the partial association between the relative cost of childcare and the probability of giving birth equals $\beta + \gamma_j$. Table 4 repeats Table 2. The only difference is the inclusion of the educational dummies and their interaction with the relative cost. As can be seen from the Table, the effect increases with the level of education (in absolute terms) and the differences are quantitatively large. Column (7)

Table 4
The Association Between Giving a Birth and Childcare Relative Cost: 1983–2012

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Birth						
Dependent variable: birth in the past 12 months							
Childcare relative cost	-0.002 (0.001)	-0.006*** (0.001)	-0.003** (0.001)	-0.004*** (0.001)	-0.005*** (0.001)	-0.015*** (0.001)	-0.023*** (0.002)
Childcare relative cost × high school graduates	0.005** (0.002)	0.003 (0.002)	0.003 (0.002)	0.003* (0.001)	0.003* (0.001)	0.001 (0.001)	0.001 (0.003)
Childcare relative cost × some college	0.002 (0.002)	-0.000 (0.001)	-0.000 (0.001)	-0.000 (0.001)	-0.000 (0.001)	-0.003* (0.001)	-0.003 (0.002)
Childcare relative cost × college graduates	-0.007*** (0.002)	-0.008*** (0.001)	-0.008*** (0.001)	-0.007*** (0.001)	-0.008*** (0.001)	-0.014*** (0.001)	-0.014*** (0.002)
Childcare relative cost × advanced degrees	-0.011*** (0.003)	-0.014*** (0.003)	-0.014*** (0.002)	-0.013*** (0.002)	-0.013*** (0.002)	-0.021*** (0.003)	-0.027*** (0.003)
Number of children	-0.008*** (0.000)	-0.003*** (0.001)	-0.014*** (0.000)	-0.014*** (0.000)	-0.014*** (0.000)	-0.015*** (0.000)	-0.020*** (0.001)
Total personal income							
Total personal income ²							
Spouse's wage							
Age dummies	No	Yes	Yes	Yes	Yes	Yes	Yes
Marital status dummies	No	No	Yes	Yes	Yes	Yes	Yes
Year dummies	No	No	No	Yes	Yes	Yes	Yes
State dummies	No	No	No	No	Yes	Yes	Yes
Observations	514,829	514,829	514,829	514,829	514,829	514,829	305,847
R ²	0.005	0.039	0.066	0.067	0.068	0.070	0.082

Notes: *p < 0.05, **p < 0.01, ***p < 0.001. Robust standard errors adjusted for heteroscedasticity and clustered at the state level are reported in parentheses. Linear probability models. All models are weighted by CPS sampling weights. See note to Table 2 for further details.

of Table 4 suggests that the effect for women with advanced degrees is more than double the effect for women with up to some college education.

Figure 7 visualises these estimates by translating them into the counterfactual hybrid fertility rate in 2001–11 under the 1983–85 relative childcare cost. As can be seen from the Figure, the counterfactual fertility of women with college education is largely unchanged when we allow the effect to differ by educational groups. For women with advanced degrees, however, the drop increases by nearly 50%, making the cross-sectional relationship between fertility and education almost monotonically decreasing.

These results provide strong support for the marketisation hypothesis. Accounting only for the change in the relative cost of childcare can nearly eliminate the U-shaped fertility pattern. Plausibly, if we could take into account changes in the relative cost of other services such as housekeeping, laundry and takeouts the counterfactual fertility would have looked even more like the cross-section prior to the 2000s.

3.2. *Purchase of Childcare Services*

The previous subsection shows the response of fertility to the change in the relative cost of childcare. In this subsection we utilise the childcare module in the Survey of Income and Program Participation (SIPP), to show how the purchase of childcare services has changed over time across the five educational groups.³¹

We use the topical module of the micro data of the SIPP for the years 1990, 1996, 2001, 2004 and 2008.³² In 1990, all women with children under 5 specified a main arrangement for childcare and only 4% did not specify any childcare hours. In contrast, in 1996, 2001, 2004 and 2008, between 26% and 28% did not specify childcare hours. In all these years, the fraction of women with children under 5 who did not specify childcare hours decreased with education. With these caveats in mind, we now describe the evolution of the cross-sectional relationship between purchased childcare hours and women's education.

Figure 8 shows the average weekly hours of paid childcare by all women in the 25–50 age group. The Figure presents two important features that are worth mentioning. First, the cross-sectional relationship monotonically increases with education in all years. Second, while there has been a large increase in paid childcare hours by women with college and advanced degrees, there is no clear trend over time for lower educational groups.³³

³¹ Besharov *et al.* (2006) list the major shortcoming of the childcare module in the SIPP. Perhaps the most severe problem is that the SIPP is supposed to interview at least one parent of each child 15 years old and younger in the household. But if a parent is not available, the SIPP allows proxy responses in order to reduce the 'person non-response' rate. Proxy responses, however, are probably less complete and less accurate than those from the child's mother. Besharov *et al.* (2006) calculate that proxy respondents constituted between 30% and 40% of respondents during the 1990s and early 2000s.

³² Data were downloaded from: <http://www.nber.org/data/survey-of-income-and-program-participation-sipp-data.html>

³³ We also calculated expenditures on childcare across the educational groups for these years and found very similar patterns.

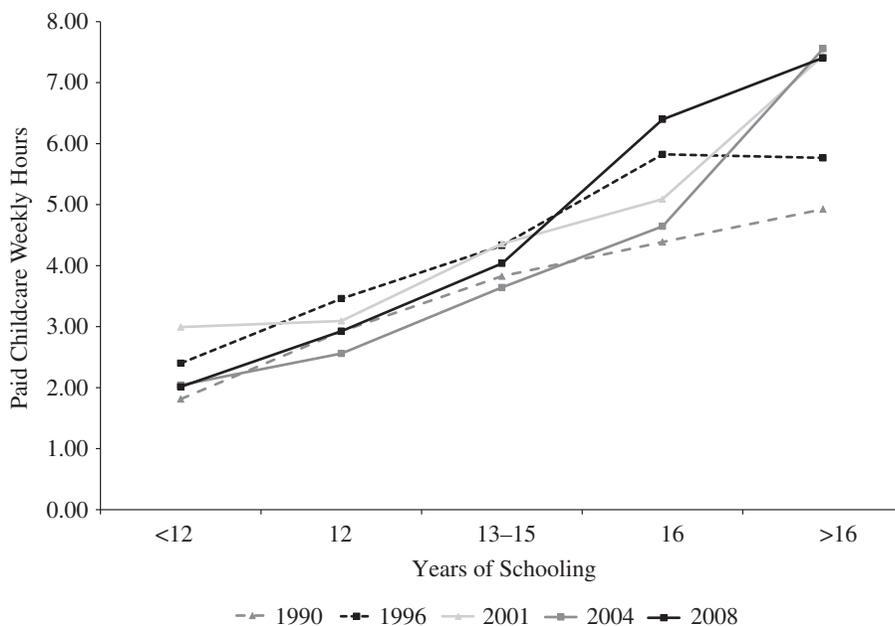


Fig. 8. *Paid Childcare Weekly Hours Per Woman Aged 25-50*

4. Supportive Evidence and Alternative Hypotheses

In this Section, we provide supportive evidence for our theory and rule out alternative hypotheses. We begin by showing that the number of average hours worked increases monotonically with women's education and that this pattern is true for all women and mothers of newborns regardless of marital status. We then discuss several competing hypotheses related to marriage rates, the role of husbands and improvements in reproductive technologies.

4.1. *Labour Supply and Marriage Rates*

In Section 1 we established that the association between fertility and women's education is U-shaped. Using the ACS sample for the years 2001-11, we present here evidence in support of our model. We begin with labour supply. It is well established that the cross-sectional relationship between female labour supply and education is upward sloping. Figure 9 shows that the usual hours women aged 25-50 worked per week during the past 12 months monotonically increases with education.³⁴ Notice that the difference across the educational groups is quantitatively large. Among all women aged 25-50, women lacking a high school diploma work somewhat less than 21 hours per week, while women with advanced degrees work more than 36 hours per week (Figure 10).

The positive correlation between fertility and labour supply for women with at least a college degree, however, does not necessarily imply that highly educated women work

³⁴ We restrict the minimum age to 25 because women with advanced degrees might still be out of the labour market at younger ages.

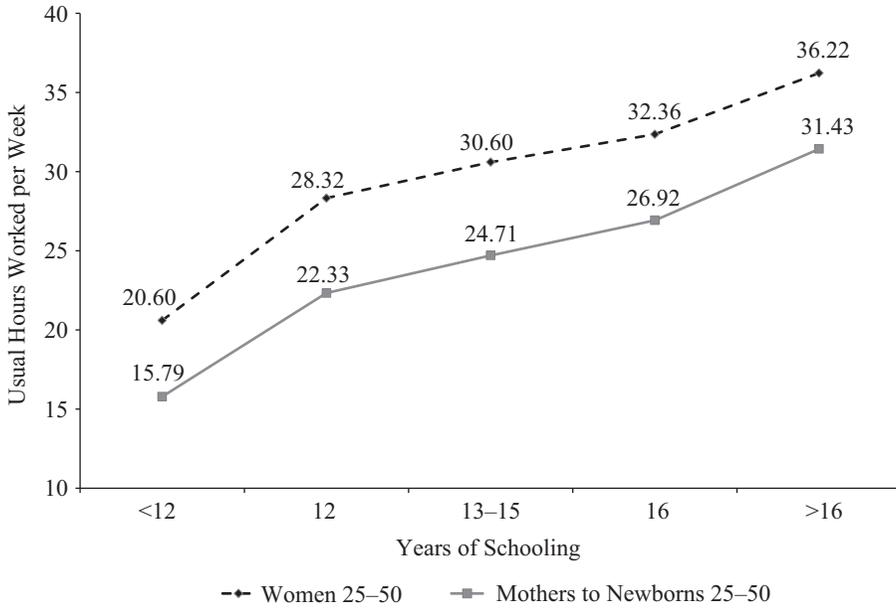


Fig. 9. Usual Hours Worked by Women Aged 25-50 and Women with Newborns, 2001-11
 Source. Authors' calculations using data from the American Community Survey.

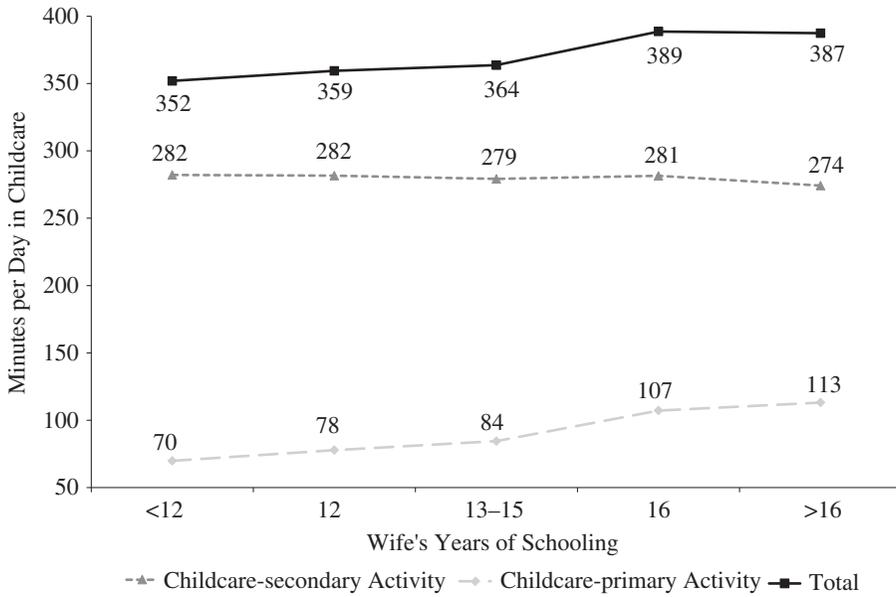


Fig. 10. Minutes Per Day in Childcare, 2003-11
 Note. Married men according to the educational group of their wives.
 Source. Authors' calculation using American Time Use Survey data.

more and have more children. Since only a small fraction of women give birth in each year, it could be, for example, that women who gave birth in a given year do not work at all during that same year. To address this, Figure 9 also shows the cross-sectional relationship between education and usual hours worked for the sub-sample of women age 25–50 who gave birth during the reference period.³⁵ As can be seen from the Figure, highly educated mothers of newborns work more hours per week than less educated mothers with newborns.

So far we have shown that highly educated women have higher fertility rates and work more hours and that, among mothers of newborns, the number of usual hours worked increases with education. However, a potential threat to the dominance of marketisation might derive from spouses who are married to highly educated women who allocate more time to childcare. To investigate this, we use data from the American Time Use Survey (ATUS) for the years 2003–11. Our sample consists of all men who have a white, non-Hispanic spouse and at least one child below the age of 13.³⁶ Using this sample we estimate the total time spent during diary day providing primary and secondary childcare for household children. Figure 10 shows that the total time that fathers spent on childcare increases with mothers' education until college graduates and then becomes flat. We conclude from this Figure that fathers' time is probably an important input into the fertility decision of highly educated women. Nevertheless, it seems to play a lesser role in accounting for the gap in fertility between women with exactly a college degree and women with an advanced degree.

Another concern our model may arouse is that marriage rates differ across different educational groups. If married women have higher fertility rates and if more educated women have higher marriage rates, more educated women's higher fertility rates may not be caused by marketisation but rather simply by their higher marriage rates. Figure 11 shows the fraction of currently married women by age-group and education.

As can be seen, the fraction of currently married women increases with age at any level of education; for women above the age of 30, it increases with educational attainment only through college degrees. Notice that the fraction of women with advanced degrees who are currently married is somewhat lower than that of women with a college degree. Thus, the increase in fertility between women with college degrees and advanced degrees cannot be attributed to marriage rates.

Another concern might be related to the mechanisms that govern these outcomes. For example, it might be that the increase in labour supply of mothers of newborns along the educational gradient, as shown in Figure 9, is driven by the pattern of unmarried mothers, while the reverse is true among married mothers. Figure 12 presents the number of usual hours worked for women aged 25–50 with a newborn by marital status.³⁷

Two features emerge from the Figure. First, at any level of education, unmarried mothers work more than married ones. Second, and more importantly for our theory, is the fact that regardless of marital status, the usual hours worked increase with women's education.

³⁵ Both curves remain intact if we use the age group 15–50.

³⁶ Information on secondary childcare is collected only for adults with children under the age of 13. We define childcare as a primary and secondary activity as in Altintas (2012).

³⁷ Both curves remain intact if we use the age group 15–50.

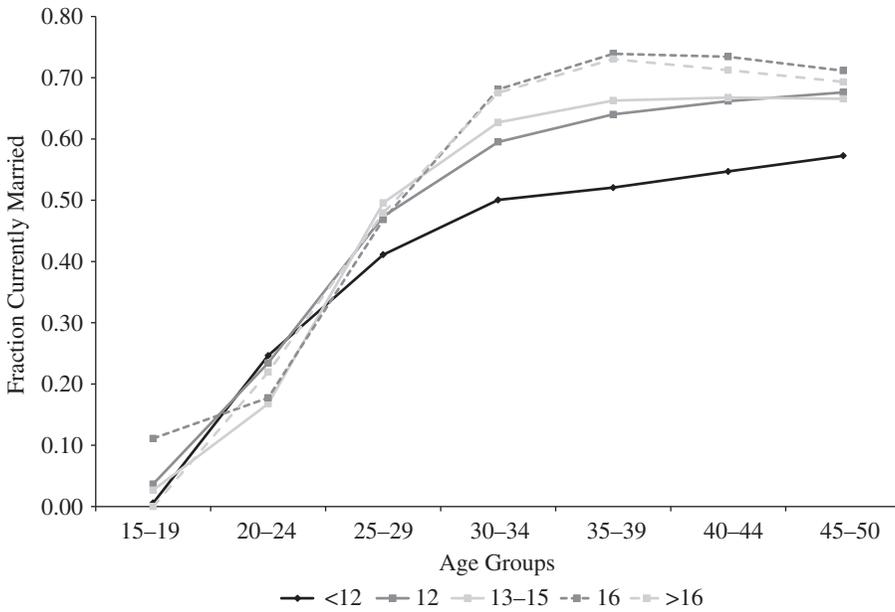


Fig. 11. *Fraction of Currently Married Women by Age and Education, 2001–11*
 Source. Authors' calculations using data from the American Community Survey.

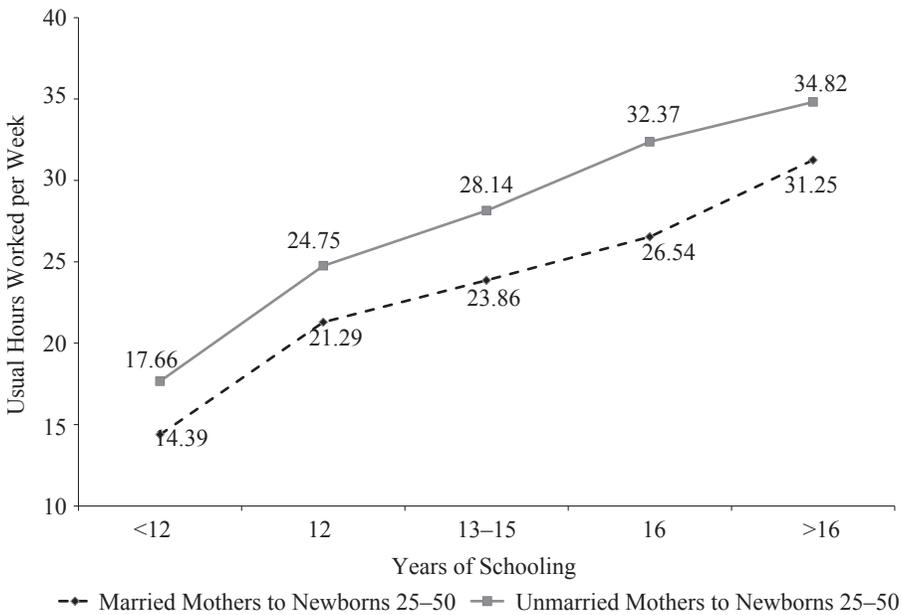


Fig. 12. *Usual Hours Worked of Women with Newborns by Marital Status, 2001–11*
 Source. Authors' calculations using data from the American Community Survey.

4.2. *Improvement in Reproductive Technology*

One possible hypothesis for the rise in fertility among highly educated women is that current reproductive technology allows women today to spend much of their fertile period in school and to postpone fertility to relatively older ages, an option that was not available in the past. During the 2000s, the number of births per 1,000 white American women with advanced degrees in the age groups 35–39, 40–44 and 45–49 were 97.3, 24.4 and 4.2 respectively. Are these unprecedentedly high levels of fertility rates for women in these age groups? Figure 13 suggests this is not the case. In 1920, the number of live births per 1,000 white American women in the age groups 35–39, 40–44 and 45–49 were 79.7, 31.9 and 3.8 respectively. For foreign born whites, the corresponding numbers were 107.4, 42.8 and 5.8, well above the current rates among highly educated women. In several states, fertility rates in 1920 among all white women were even higher. For example, in North Carolina, the number of live births per 1,000 white women in the age groups 35–39, 40–44 and 45–49 were 144.3, 62.1, and 9.9 respectively. The corresponding numbers for Utah were 128.4, 68.2 and 10.8; for South Carolina 114.5, 49.6 and 5.8; for Virginia 114.1, 43.4 and 6; and for Kentucky 100.5, 44.9 and 5.3 respectively.³⁸ These historical levels of fertility rates among women above age 35 suggest that the current level of fertility among the highly educated is not likely to be driven by reproductive technology which was not available for women at the time when the cross-sectional relationship between fertility and education was monotonically declining (Figure 13).

Alternatively, let us assume that highly educated women have lower desired fertility rates than middle-education women. However, the highly educated can afford assisted reproductive technology (ART) while middle-education women cannot. This may make fertility rates higher among the highly educated than among middle-education women. Moreover, since ART often results in multiple births, the fertility of the highly educated women may exceed desired fertility.³⁹ To investigate the plausibility of the affordability of ART, we take advantage of the fact that 15 states have infertility insurance laws that provide coverage to infertile individuals.⁴⁰ Although there is some variation in coverage among these 15 states and coverage is not complete in any of them, we re-estimate the models in Table 1 only for women who live in these states. As shown in Table B4 in the online Appendix, the U-shaped patterns still holds though it is somewhat less pronounced.

The fact that ART often results in multiple births does not pose a threat to our approach because of the way we coded the data. Specifically, the ACS asks women if they have given birth during the past 12 months. We coded this to indicate the birth of one infant. Hence, whenever the birth resulted in more than one infant we in fact under-counted births. If, indeed, most of these births occur to highly educated women, our estimates of the fertility of these women are biased downward. Given the evidence, we conclude that it is unlikely that ART is an important factor behind the emergence of the U-shape pattern.

³⁸ These data are taken from the Vital Statistics Rates in the US 1900–40, Tables 47 and 48.

³⁹ According to CDC's 2011 ART Fertility Clinic Success Rates Report, there were 47,818 live births and 61,610 live born infants in the US during 2011. This amounts to 1.2% of the total number of births in the US in 2011.

⁴⁰ These states are: Arkansas, California, Connecticut, Hawaii, Illinois, Louisiana, Maryland, Massachusetts, Montana, New Jersey, New York, Ohio, Rhode Island, Texas and West Virginia.

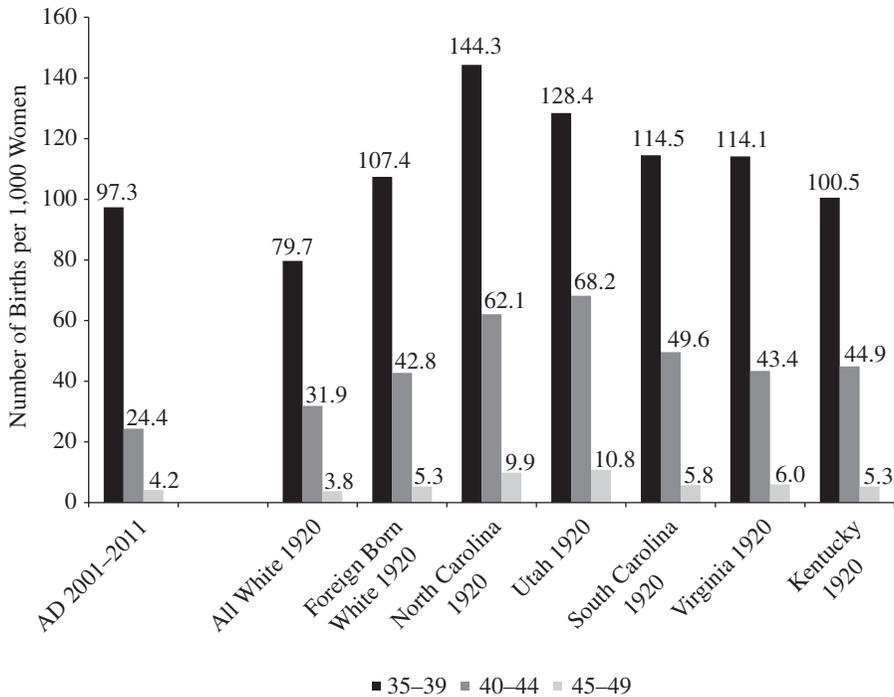


Fig. 13. *Number of Births Per 1,000 White Women in the US: Women with Advanced Degrees (2001–11) and Historical Rates*

Source. Authors' calculations using data from the American Community Survey and Vital Statistics Rates in the US 1900–40, Tables 47 and 48.

5. Concluding Remarks

We present new evidence that between 2001 and 2011, the cross-sectional relationship between fertility and women's education in the US is U-shaped. This pattern is robust to controlling for a host of covariates such as family income, marital and age dummies, year and state of residence dummies. Our analysis of earlier periods shows that this pattern is new, revealing an emerging new pattern of cohort fertility. Studying the period 1983–2012, we found that childcare has become relatively more expensive to women with less than a college degree but relatively cheaper for women with college or advanced degrees. We then show that the association between the probability of giving birth and our measure of the relative cost of childcare services is negative, highly significant and robust to the inclusions of various controls and different specifications that correct for endogeneity of women's wages and selection bias in the labour market. Moreover, we show that this structural relationship is stable over time and independent of the relative cost of childcare. Conducting a counterfactual exercise we show that the change in the relative cost of childcare over these 30 years accounts for much of the U-shaped pattern.

Our model demonstrates how parents can substitute their own parenting time for market-purchased childcare. We show that highly educated women substitute a significant part of their own parenting with childcare. This enables them to have more

children and work longer hours. Furthermore, we show that these highly educated women not only work more and have more children, they invest more in the education of each of their children. This result may have important implications for the relationship between inequality and economic growth. In particular, de la Croix and Doepke (2003) argue that because poorer individuals have more children and invest less in the education of each child, higher inequality leads to lower growth. The evidence presented here, that highly educated women choose larger families than women with intermediate levels of education, may weaken or even undo this result. This line of inquiry is beyond the scope of the current article and is left for future research.

Our model can also explain the differences in fertility and time allocation of women between the US and Europe. European women spend more time on home production and less time in labour market activities than American women (Freeman and Schettkat, 2005). They also give birth to fewer children. For example, in 2009, the gap in TFR between the US and EU members amounts to nearly one-half of a child per woman. Another noticeable difference between the US and Europe is in the degree of income inequality. For example, according to OECD statistics, the Gini coefficient after tax and transfers in the mid 2000s for the working age population was 0.37 in the US while it was 0.31 for all European OECD members. Similarly, the 90:10 ratio during that period in the US was 5.91 while for all European OECD members it was 3.84. In Hazan and Zoabi (2011) we studied the aggregate behaviour of the model presented in this article. Specifically, we computed the average fertility and time allocated to labour market and home production in our model economy. We then analysed the effect of a mean preserving spread of the distribution of women's human capital. This is the model's analogy to the higher income inequality in the US when compared to Europe. Consistent with the data, we found that an increase in inequality leads unambiguously to an increase in average fertility. The predictions of the model with respect to the average time allocated to home production and children depend on the model's parameters. We demonstrated, however, that the time allocated to the labour market and to childcare increase in inequality while the sum of time allocated to childcare and home production decrease in inequality. We believe that research investigating differences between the US and Europe along these lines in greater depth will undoubtedly prove to be very informative.

*Tel-Aviv University and CEPR
The New Economic School*

Submitted: 17 August 2012

Accepted: 24 January 2014

Additional Supporting Information may be found in the online version of this article:

Appendix A. Model's Extensions.

Appendix B. Additional Regression Results.

Appendix C. Estimating Fertility Rates Using Age of Youngest Child.

Data S1.

References

- Aiyagari, S.R., Greenwood, J. and Seshadri, A. (2002). 'Efficient investment in children', *Journal of Economic Theory*, vol. 102(2), pp. 290–321.
- Altintas, E. (2012). 'Parents' time with children: micro and macro perspectives', Ph.D. Thesis, Nuffield College, University of Oxford.
- Apps, P. and Rees, R. (2004). 'Fertility, taxation and family policy', *Scandinavian Journal of Economics*, vol. 106(4), pp. 745–63.
- Attanasio, O., Low, H. and Sanchez-Marcos, V. (2008). 'Explaining changes in female labour supply in a life-cycle model', *American Economic Review*, vol. 98(4), pp. 1517–52.
- Baum-Snow, N. and Neal, D. (2009). 'Mismeasurement of usual hours worked in the census and ACS', *Economics Letters*, vol. 102(1), pp. 39–41.
- Becker, G.S. (1965). 'A theory of the allocation of time', *ECONOMIC JOURNAL*, vol. 75(299), pp. 493–517.
- Besharov, D.J., Morrow, J.S. and Fengyan Shi, A. (2006). 'Child care data in the survey of income and program participation (SIPP): Inaccuracies and corrections', University of Maryland, Welfare Reform Academy.
- Black, D.A., Kolesnikova, N., Sanders, S.G. and Taylor, L.J. (2013). 'Are children "normal"?'', *Review of Economics and Statistics*, vol. 95(1), pp. 21–33.
- Charles, K.K., Hurst, E. and Killewald, A. (2013). 'Marital sorting and parental wealth', *Demography*, vol. 50(1), pp. 51–70.
- Cherchye, L., De Rock, B. and Vermeulen, F. (2012). 'Married with children: a collective labor supply model with detailed time use and intrahousehold expenditure information', *American Economic Review*, vol. 102(7), pp. 3377–405.
- Cortes, P. and Tessada, J. (2011). 'Low-skilled immigration and the labor supply of highly skilled women', *American Economic Journal: Applied Economics*, vol. 3(1), pp. 88–123.
- de la Croix, D. and Doepke, M. (2003). 'Inequality and growth: why differential fertility matters', *American Economic Review*, vol. 93(4), pp. 1091–113.
- Freeman, R.B. and Schettkat, R. (2005). 'Marketization of household production and the EU–US gap in work', *Economic Policy*, vol. 20(41), pp. 5–50.
- Furtado, D. and Hock, H. (2010). 'Low skilled immigration and work-fertility tradeoffs among high skilled US natives', *American Economic Review*, vol. 100(2), pp. 224–8.
- Galor, O. (2011). *Unified Growth Theory*, Princeton NJ: Princeton University Press.
- Galor, O. and Weil, D.N. (1996). 'The gender gap, fertility, and growth', *American Economic Review*, vol. 86(3), pp. 374–87.
- Greenwood, J., Seshadri, A. and Vandenbroucke, G. (2005). 'The baby boom and baby bust', *American Economic Review*, vol. 95(1), pp. 183–207.
- Gronau, R. (1977). 'Leisure, home production, and work – the theory of the allocation of time revisited', *Journal of Political Economy*, vol. 85(6), pp. 1099–123.
- Guinnane, T.W. (2011). 'The historical fertility transition: a guide for economists', *Journal of Economic Literature*, vol. 49(3), pp. 589–614.
- Guryan, J., Hurst, E. and Kearney, M. (2008). 'Parental education and parental time with children', *Journal of Economic Perspectives*, vol. 22(3), pp. 23–46.
- Hazan, M. and Berdugo, B. (2002). 'Child labor, fertility, and economic growth', *ECONOMIC JOURNAL*, vol. 112(482), pp. 810–28.
- Hazan, M. and Zoabi, H. (2011). 'Do highly educated women choose smaller families?', Discussion Paper 8590, CEPR.
- Isen, A. and Stevenson, B. (2010). 'Women's education and family behavior: trends in marriage, divorce and fertility', in (J. Shoven, ed.), *Demography and the Economy*, pp. 107–140, Cambridge: University of Chicago Press.
- Jones, L.E. and Tertilt, M. (2008). 'An economic history of fertility in the US: 1826–1960', in (P. Rupert, ed.), *Frontiers of Family Economics*, pp. 165–230, Bingley: Emerald.
- King, M., Ruggles, S.J., Alexander, T., Flood, S., Genadek, K., Schroeder, M.B., Trampe, B. and Vick, R. (2010). *Integrated Public Use Microdata Series, Current Population Survey: Version 3.0. [Machine-readable database]*, Minneapolis, MN: University of Minnesota.
- Kremer, M. and Chen, D.L. (2002). 'Income distribution dynamics with endogenous fertility', *Journal of Economic Growth*, vol. 7(3), pp. 227–58.
- Manning, A. (2004). 'We can work it out: the impact of technological change on the demand for low-skill workers', *Scottish Journal of Political Economy*, vol. 51(5), pp. 581–608.
- Mazzolari, F. and Ragusa, G. (2013). 'Spillovers from high-skill consumption to low-skill labor markets', *Review of Economics and Statistics*, vol. 95(1), pp. 74–86.
- Moav, O. (2005). 'Cheap children and the persistence of poverty', *ECONOMIC JOURNAL*, vol. 115(500), pp. 88–110.

- Mulligan, C.B. and Rubinstein, Y. (2008). 'Selection, investment, and women's relative wages over time', *Quarterly Journal Economics*, vol. 123(3), pp. 1061–110.
- Pencavel, J. (1998). 'Assortative mating by schooling and the work behavior of wives and husbands', *American Economic Review*, vol. 88(2), pp. 326–9.
- Preston, S.H. and Hartnett, C. (2008). 'The future of american fertility', Working Paper 14498, NBER.
- Ramey, G. and Ramey, V.A. (2010). 'The rug rat race', *Brookings Papers on Economic Activity*, vol. Spring, pp. 129–76.
- Ruggles, S.J., Alexander, T., Genadek, K., Goeken, R., Schroeder, M.B. and Sobek, M. (2010). *Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]*, Minneapolis, MN: University of Minnesota.
- Sacks, D.W. and Stevenson, B. (2010). 'What explains the rapid increase in childcare time: a discussion of Gary Ramey and Valerie A. Ramey's "the rug rat race"', *Brookings Papers on Economic Activity*, vol. Spring, pp. 184–96.
- Shang, Q. and Weinberg, B.A. (2013). 'Opting for families: recent trends in the fertility of highly educated women', *Journal of Population Economics*, vol. 26(1), pp. 5–32.
- Weil, D.N. (2005). *Economic Growth*, Boston MA: Pearson.